

MPI: A Message-Passing Interface Standard

Version 4.1

(Draft)

Unofficial, for comment only

Message Passing Interface Forum

October 17, 2023

DRAFT

1 This document describes the Message-Passing Interface (MPI) standard, version 4.1.

2 The MPI standard includes point-to-point message-passing, collective communications,
3 group and communicator concepts, process topologies, environmental management, process
4 creation and management, one-sided communications, extended collective operations, ex-
5 ternal interfaces, I/O, some miscellaneous topics, and multiple tool interfaces. Language
6 bindings for C and Fortran are defined.

7 Historically, the evolution of the standard is:

- 8
- 9 • MPI-1.0 (May 5, 1994): Initial release.
- 10
- 11 • MPI-1.1 (June 12, 1995): Minor updates and bug fixes.
- 12
- 13 • MPI-1.2 (July 18, 1997): Several clarifications and additions.
- 14
- 15 • MPI-2.0 (July 18, 1997): New functionality and all the clarifications and additions
16 from MPI-1.2.
- 17
- 18 • MPI-1.3 (May 30, 2008): For historical reasons, combining the MPI-1.1, MPI-1.2, and
19 several errata documents into one combined document.
- 20
- 21 • MPI-2.1 (June 23, 2008): Combining the previous documents.
- 22
- 23 • MPI-2.2 (September 4, 2009): Additional clarifications and seven new routines.
- 24
- 25 • MPI-3.0 (September 21, 2012): Extension of MPI-2.2.
- 26
- 27 • MPI-3.1 (June 4, 2015): Clarifications and minor extensions to MPI-3.0.
- 28
- 29 • MPI-4.0 (June 9, 2021): Significant new features beyond MPI-3.1.
- 30
- 31 • MPI-4.1 (XX X, 2023): Clarifications and minor extensions to MPI-4.0.

32 Comments. Please send comments on MPI to the MPI Forum as follows:

- 33 1. Subscribe to <https://lists.mpi-forum.org/mailman/listinfo/mpi-comments>
- 34 2. Send your comment to: mpi-comments@lists.mpi-forum.org, together with the version
35 of the MPI standard and the page and line numbers on which you are commenting.

36 Your comment will be forwarded to MPI Forum committee members for consideration.
37 Messages sent from an unsubscribed e-mail address will not be considered.

38

39

40

41

42

43

44

45 ©1993, 1994, 1995, 1996, 1997, 2008, 2009, 2012, 2015, 2021, 2023 University of Tennessee,
46 Knoxville, Tennessee. Permission to copy without fee all or part of this material is granted,
47 provided the University of Tennessee copyright notice and the title of this document appear,
48 and notice is given that copying is by permission of the University of Tennessee.

Version 4.1: XXXX, 2023. This version of the MPI-4.1 standard contains mostly corrections and clarifications to the MPI-4.0 document. Several routines, the attribute key MPI_HOST, and the mpif.h Fortran include file are deprecated.

Version 4.0: June 9, 2021. This version of the MPI-4 standard is a major update and includes significant new functionality. The largest changes are the addition of large-count versions of many routines to address the limitations of using an int or INTEGER for the count parameter, persistent collectives, partitioned communications, an alternative way to initialize MPI, application info assertions, and improvements to the definitions of error handling. In addition, there are a number of smaller improvements and corrections.

Version 3.1: June 4, 2015. This document contains mostly corrections and clarifications to the MPI-3.0 document. The largest change is a correction to the Fortran bindings introduced in MPI-3.0. Additionally, new functions added include routines to manipulate MPI_Aint values in a portable manner, nonblocking collective I/O routines, and routines to get the index value by name for MPI_T performance and control variables.

Version 3.0: September 21, 2012. Coincident with the development of MPI-2.2, the MPI Forum began discussions of a major extension to MPI. This document contains the MPI-3 standard. This version of the MPI-3 standard contains significant extensions to MPI functionality, including nonblocking collectives, new one-sided communication operations, and Fortran 2008 bindings. Unlike MPI-2.2, this standard is considered a major update to the MPI standard. As with previous versions, new features have been adopted only when there were compelling needs for the users. Some features, however, may have more than a minor impact on existing MPI implementations.

Version 2.2: September 4, 2009. This document contains mostly corrections and clarifications to the MPI-2.1 document. A few extensions have been added; however all correct MPI-2.1 programs are correct MPI-2.2 programs. New features were adopted only when there were compelling needs for users, open source implementations, and minor impact on existing MPI implementations.

Version 2.1: June 23, 2008. This document combines the previous documents MPI-1.3 (May 30, 2008) and MPI-2.0 (July 18, 1997). Certain parts of MPI-2.0, such as some sections of Chapter 4, Miscellany, and Chapter 7, Extended Collective Operations, have been merged into the chapters of MPI-1.3. Additional errata and clarifications collected by the MPI Forum are also included in this document.

Version 1.3: May 30, 2008. This document combines the previous documents MPI-1.1 (June 12, 1995) and the MPI-1.2 chapter in MPI-2 (July 18, 1997). Additional errata collected by the MPI Forum referring to MPI-1.1 and MPI-1.2 are also included in this document.

Version 2.0: July 18, 1997. Beginning after the release of MPI-1.1, the MPI Forum began meeting to consider corrections and extensions. MPI-2 has been focused on process creation and management, one-sided communications, extended collective communications, external

1 interfaces and parallel I/O. A miscellany chapter discusses items that do not fit elsewhere,
2 in particular language interoperability.

3
4 Version 1.2: July 18, 1997. The MPI-2 Forum introduced MPI-1.2 as Chapter 3 in the
5 standard “MPI-2: Extensions to the Message-Passing Interface”, July 18, 1997. This section
6 contains clarifications and minor corrections to Version 1.1 of the MPI standard. The only
7 new function in MPI-1.2 is one for identifying to which version of the MPI standard the
8 implementation conforms. There are small differences between MPI-1 and MPI-1.1. There
9 are very few differences between MPI-1.1 and MPI-1.2, but large differences between MPI-1.2
10 and MPI-2.

11
12 Version 1.1: June, 1995. Beginning in March, 1995, the Message-Passing Interface Forum
13 reconvened to correct errors and make clarifications in the MPI document of May 5, 1994,
14 referred to below as Version 1.0. These discussions resulted in Version 1.1. The changes
15 from Version 1.0 are minor. A version of this document with all changes marked is available.

16
17 Version 1.0: May, 1994. The Message-Passing Interface Forum, with participation from
18 over 40 organizations, has been meeting since January 1993 to discuss and define a set of
19 library interface standards for message passing. The Message-Passing Interface Forum is
20 not sanctioned or supported by any official standards organization.

21 The goal of the Message-Passing Interface, simply stated, is to develop a widely used
22 standard for writing message-passing programs. As such the interface should establish a
23 practical, portable, efficient, and flexible standard for message-passing.

24 This is the final report, Version 1.0, of the Message-Passing Interface Forum. This
25 document contains all the technical features proposed for the interface. This copy of the
26 draft was processed by L^AT_EX on May 5, 1994.

Contents

List of Figures	xix
List of Tables	xxi
Acknowledgments	xxii
1 Introduction to MPI	1
1.1 Overview and Goals	1
1.2 Background of MPI-1.0	2
1.3 Background of MPI-1.1, MPI-1.2, and MPI-2.0	2
1.4 Background of MPI-1.3 and MPI-2.1	3
1.5 Background of MPI-2.2	4
1.6 Background of MPI-3.0	4
1.7 Background of MPI-3.1	4
1.8 Background of MPI-4.0	4
1.9 Background of MPI-4.1	5
1.10 Who Should Use This Standard?	5
1.11 What Platforms Are Targets for Implementation?	5
1.12 What Is Included in the Standard?	6
1.13 Side-documents	6
1.14 Organization of This Document	7
2 MPI Terms and Conventions	9
2.1 Document Notation	9
2.2 Naming Conventions	9
2.3 Procedure Specification	10
2.4 Semantic Terms	11
2.4.1 MPI Operations	11
2.4.2 MPI Procedures	14
2.4.3 MPI Datatypes	16
2.5 Datatypes	17
2.5.1 Opaque Objects	17
2.5.2 Array Arguments	19
2.5.3 State	19
2.5.4 Named Constants	19
2.5.5 Choice	20
2.5.6 Absolute Addresses and Relative Address Displacements	21
2.5.7 File Offsets	21
2.5.8 Counts	21
2.6 Language Binding	22
2.6.1 Deprecated and Removed Interfaces	22
2.6.2 Fortran Binding Issues	23

2.6.3	C Binding Issues	24
2.6.4	Functions and Macros	25
2.7	Processes	25
2.8	Error Handling	26
2.9	Progress	27
2.10	Implementation Issues	29
2.10.1	Independence of Basic Runtime Routines	29
2.10.2	Interaction with Signals	29
2.11	Examples	30
3	Point-to-Point Communication	31
3.1	Introduction	31
3.2	Blocking Send and Receive Operations	32
3.2.1	Blocking Send	32
3.2.2	Message Data	33
3.2.3	Message Envelope	35
3.2.4	Blocking Receive	36
3.2.5	Return Status	39
3.2.6	Passing MPI_STATUS_IGNORE for Status	42
3.2.7	Blocking Send-Receive	43
3.3	Datatype Matching and Data Conversion	46
3.3.1	Type Matching Rules	46
	Type MPI_CHARACTER	48
3.3.2	Data Conversion	48
3.4	Communication Modes	50
3.5	Semantics of Point-to-Point Communication	55
3.6	Buffer Allocation and Usage	58
3.6.1	Model Implementation of Buffered Mode	68
3.7	Nonblocking Communication	69
3.7.1	Communication Request Objects	70
3.7.2	Communication Initiation	71
3.7.3	Communication Completion	78
3.7.4	Semantics of Nonblocking Communication Operations	82
3.7.5	Multiple Completions	83
3.7.6	Non-Destructive Test of status	90
3.8	Probe and Cancel	94
3.8.1	Probe	94
3.8.2	Matching Probe	97
3.8.3	Matched Receives	100
3.8.4	Cancel	102
3.9	Persistent Communication Requests	104
3.10	Null MPI Processes	111
4	Partitioned Point-to-Point Communication	113
4.1	Introduction	113
4.2	Semantics of Partitioned Point-to-Point Communication	114
4.2.1	Communication Initialization and Starting with Partitioning	116
4.2.2	Communication Completion under Partitioning	120

4.2.3	Semantics of Communications in Partitioned Mode	121
4.3	Partitioned Communication Examples	121
4.3.1	Partition Communication with Threads/Tasks Using OpenMP 4.0 or later	122
4.3.2	Send-only Partitioning Example with Tasks and OpenMP version 4.0 or later	123
4.3.3	Send and Receive Partitioning Example with OpenMP version 4.0 or later	124
5	Datatypes	127
5.1	Derived Datatypes	127
5.1.1	Type Constructors with Explicit Addresses	129
5.1.2	Datatype Constructors	129
5.1.3	Subarray Datatype Constructor	141
5.1.4	Distributed Array Datatype Constructor	143
5.1.5	Address and Size Procedures	148
5.1.6	Lower-Bound and Upper-Bound Markers	151
5.1.7	Extent and Bounds of Datatypes	153
5.1.8	True Extent of Datatypes	155
5.1.9	Commit and Free	156
5.1.10	Duplicating a Datatype	158
5.1.11	Use of General Datatypes in Communication	158
5.1.12	Correct Use of Addresses	161
5.1.13	Decoding a Datatype	162
5.1.14	Examples	170
5.2	Pack and Unpack	178
5.3	Canonical MPI_PACK and MPI_UNPACK	185
6	Collective Communication	189
6.1	Introduction and Overview	189
6.2	Communicator Argument	192
6.2.1	Specifics for Intra-Communicator Collective Operations	192
6.2.2	Applying Collective Operations to Inter-Communicators	193
6.2.3	Specifics for Inter-Communicator Collective Operations	194
6.3	Barrier Synchronization	196
6.4	Broadcast	196
6.4.1	Example using MPI_BCAST	197
6.5	Gather	198
6.5.1	Examples using MPI_GATHER, MPI_GATHERV	202
6.6	Scatter	208
6.6.1	Examples using MPI_SCATTER, MPI_SCATTERV	212
6.7	Gather-to-all	214
6.7.1	Example using MPI_ALLGATHER	217
6.8	All-to-All Scatter/Gather	218
6.9	Global Reduction Operations	224
6.9.1	Reduce	224
6.9.2	Predefined Reduction Operations	227
6.9.3	Signed Characters and Reductions	229

6.9.4	MINLOC and MAXLOC	229
6.9.5	User-Defined Reduction Operations	235
	Example of User-Defined Reduce	239
6.9.6	All-Reduce	240
6.9.7	MPI Process-Local Reduction	242
6.10	Reduce-Scatter	243
6.10.1	MPI_REDUCE_SCATTER_BLOCK	243
6.10.2	MPI_REDUCE_SCATTER	245
6.11	Scan	247
6.11.1	Inclusive Scan	247
6.11.2	Exclusive Scan	248
6.11.3	Example using MPI_SCAN	249
6.12	Nonblocking Collective Operations	250
6.12.1	Nonblocking Barrier Synchronization	253
6.12.2	Nonblocking Broadcast	253
	Example using MPI_IBCAST	254
6.12.3	Nonblocking Gather	255
6.12.4	Nonblocking Scatter	257
6.12.5	Nonblocking Gather-to-all	260
6.12.6	Nonblocking All-to-All Scatter/Gather	262
6.12.7	Nonblocking Reduce	267
6.12.8	Nonblocking All-Reduce	268
6.12.9	Nonblocking Reduce-Scatter with Equal Blocks	269
6.12.10	Nonblocking Reduce-Scatter	270
6.12.11	Nonblocking Inclusive Scan	271
6.12.12	Nonblocking Exclusive Scan	272
6.13	Persistent Collective Operations	273
6.13.1	Persistent Barrier Synchronization	275
6.13.2	Persistent Broadcast	276
6.13.3	Persistent Gather	277
6.13.4	Persistent Scatter	279
6.13.5	Persistent Gather-to-all	282
6.13.6	Persistent All-to-All Scatter/Gather	285
6.13.7	Persistent Reduce	289
6.13.8	Persistent All-Reduce	291
6.13.9	Persistent Reduce-Scatter with Equal Blocks	292
6.13.10	Persistent Reduce-Scatter	293
6.13.11	Persistent Inclusive Scan	294
6.13.12	Persistent Exclusive Scan	296
6.14	Correctness	297
7	Groups, Contexts, Communicators, and Caching	305
7.1	Introduction	305
7.1.1	Features Needed to Support Libraries	305
7.1.2	MPI's Support for Libraries	305
7.2	Basic Concepts	307
7.2.1	Groups	308
7.2.2	Contexts	308

7.2.3	Intra-Communicators	309
7.2.4	Predefined Intra-Communicators	309
7.3	Group Management	310
7.3.1	Group Accessors	310
7.3.2	Group Constructors	312
7.3.3	Group Destructors	318
7.4	Communicator Management	319
7.4.1	Communicator Accessors	319
7.4.2	Communicator Constructors	321
7.4.3	Communicator Destructors	339
7.4.4	Communicator Info	340
7.5	Motivating Examples	343
7.5.1	Current Practice #1	343
7.5.2	Current Practice #2	344
7.5.3	(Approximate) Current Practice #3	344
7.5.4	Communication Safety Example	345
7.5.5	Library Example #1	346
7.5.6	Library Example #2	347
7.6	Inter-Communication	349
7.6.1	Inter-Communicator Accessors	351
7.6.2	Inter-Communicator Operations	353
7.6.3	Inter-Communication Examples	356
	Example 1: Three-Group “Pipeline”	356
	Example 2: Three-Group “Ring”	358
7.7	Caching	359
7.7.1	Functionality	360
7.7.2	Communicators	361
7.7.3	Windows	366
7.7.4	Datatypes	370
7.7.5	Error Class for Invalid Keyval	373
7.7.6	Attributes Example	374
7.8	Naming Objects	376
7.9	Formalizing the Loosely Synchronous Model	380
7.9.1	Basic Statements	380
7.9.2	Models of Execution	380
	Static Communicator Allocation	381
	Dynamic Communicator Allocation	381
	The General Case	381
8	Virtual Topologies for MPI Processes	383
8.1	Introduction	383
8.2	Virtual Topologies	384
8.3	Embedding in MPI	384
8.4	Overview of the Functions	385
8.5	Topology Constructors	386
8.5.1	Cartesian Constructor	386
8.5.2	Cartesian Convenience Function: <code>MPI_DIMS_CREATE</code>	387
8.5.3	Graph Constructor	388

8.5.4	Distributed Graph Constructor	390
8.5.5	Topology Inquiry Functions	397
8.5.6	Cartesian Shift Coordinates	405
8.5.7	Partitioning of Cartesian Structures	407
8.5.8	Low-Level Topology Functions	408
8.6	Neighborhood Collective Communication	410
8.6.1	Neighborhood Gather	411
8.6.2	Neighborhood Alltoall	415
8.7	Nonblocking Neighborhood Communication	422
8.7.1	Nonblocking Neighborhood Gather	422
8.7.2	Nonblocking Neighborhood Alltoall	425
8.8	Persistent Neighborhood Communication	429
8.8.1	Persistent Neighborhood Gather	429
8.8.2	Persistent Neighborhood Alltoall	432
8.9	An Application Example	436
9	MPI Environmental Management	441
9.1	Implementation Information	441
9.1.1	Version Inquiries	441
9.1.2	Environmental Inquiries	442
	Tag Values	443
	IO Rank	443
	Clock Synchronization	444
	Inquire Processor Name	444
	Inquire Hardware Resource Information	445
9.2	Memory Allocation	446
9.3	Error Handling	449
9.3.1	Error Handlers for Communicators	452
9.3.2	Error Handlers for Windows	454
9.3.3	Error Handlers for Files	456
9.3.4	Error Handlers for Sessions	457
9.3.5	Freeing Errorhandlers and Retrieving Error Strings	459
9.4	Error Codes and Classes	460
9.5	Error Classes, Error Codes, and Error Handlers	464
9.5.1	User-Defined Error Classes and Codes	464
9.5.2	Calling Error Handlers	468
9.6	Timers and Synchronization	470
10	The Info Object	473
11	Process Initialization, Creation, and Management	481
11.1	Introduction	481
11.2	The World Model	482
11.2.1	Starting MPI Processes	482
11.2.2	Finalizing MPI	488
11.2.3	Determining Whether MPI Has Been Initialized When Using the World Model	491
11.2.4	Allowing User Functions at MPI Finalization	492

11.3	The Sessions Model	493
11.3.1	Session Creation and Destruction Methods	494
11.3.2	Processes Sets	497
11.3.3	Runtime Query Functions	498
11.3.4	Sessions Model Examples	501
11.4	Common Elements of Both Process Models	506
11.4.1	MPI Functionality that is Always Available	506
11.4.2	Aborting MPI Processes	506
11.4.3	Memory Allocation Info	508
11.5	Portable MPI Process Startup	512
11.6	MPI and Threads	514
11.6.1	General	515
11.6.2	Clarifications	516
11.7	The Dynamic Process Model	518
11.7.1	Starting Processes	518
11.7.2	The Runtime Environment	518
11.8	Process Manager Interface	519
11.8.1	Processes in MPI	519
11.8.2	Starting Processes and Establishing Communication	519
11.8.3	Starting Multiple Executables and Establishing Communication	524
11.8.4	Reserved Keys	527
11.8.5	Spawn Example	528
11.9	Establishing Communication	530
11.9.1	Names, Addresses, Ports, and All That	530
11.9.2	Server Routines	531
11.9.3	Client Routines	534
11.9.4	Name Publishing	535
11.9.5	Reserved Key Values	538
11.9.6	Client/Server Examples	538
11.10	Other Functionality	540
11.10.1	Universe Size	540
11.10.2	Singleton MPI Initialization	541
11.10.3	MPI_APPNUM	542
11.10.4	Releasing Connections	542
11.10.5	Another Way to Establish MPI Communication	544
12	One-Sided Communications	547
12.1	Introduction	547
12.2	Initialization	548
12.2.1	Window Creation	548
12.2.2	Window That Allocates Memory	552
12.2.3	Window That Allocates Shared Memory	554
12.2.4	Window of Dynamically Attached Memory	559
12.2.5	Window Destruction	562
12.2.6	Window Attributes	563
12.2.7	Window Info	564
12.3	Communication Calls	566
12.3.1	Put	567

12.3.2	Get	569
12.3.3	Examples for Communication Calls	570
12.3.4	Accumulate Functions	573
	Accumulate	573
	Get Accumulate	575
	Fetch and Op	577
	Compare and Swap	579
12.3.5	Request-based RMA Communication Operations	580
12.4	Memory Model	586
12.5	Synchronization Calls	587
12.5.1	Fence	590
12.5.2	General Active Target Synchronization	593
12.5.3	Lock	597
12.5.4	Flush and Sync	601
12.5.5	Assertions	603
12.5.6	Miscellaneous Clarifications	605
12.6	Error Handling	605
12.6.1	Error Handlers	605
12.6.2	Error Classes	605
12.7	Semantics and Correctness	606
12.7.1	Atomicity	614
12.7.2	Ordering	614
12.7.3	Progress	615
12.7.4	Registers and Compiler Optimizations	618
12.8	Examples	618
13	External Interfaces	627
13.1	Introduction	627
13.2	Generalized Requests	627
13.2.1	Examples	631
13.3	Associating Information with Status	633
14	I/O	637
14.1	Introduction	637
14.1.1	Definitions	637
14.2	File Manipulation	639
14.2.1	Opening a File	639
14.2.2	Closing a File	641
14.2.3	Deleting a File	642
14.2.4	Resizing a File	643
14.2.5	Preallocating Space for a File	644
14.2.6	Querying the Size of a File	644
14.2.7	Querying File Parameters	645
14.2.8	File Info	646
	Reserved File Hints	648
14.3	File Views	650

14.4	Data Access	653
14.4.1	Data Access Routines	653
	Positioning	654
	Synchronism	654
	Coordination	655
	Data Access Conventions	655
14.4.2	Data Access with Explicit Offsets	656
14.4.3	Data Access with Individual File Pointers	663
14.4.4	Data Access with Shared File Pointers	673
	Noncollective Operations	674
	Collective Operations	677
	Seek	679
14.4.5	Split Collective Data Access Routines	681
14.5	File Interoperability	689
14.5.1	Datatypes for File Interoperability	691
14.5.2	External Data Representation: "external32"	693
14.5.3	User-Defined Data Representations	694
	Extent Callback	697
	Datarep Conversion Functions	698
14.5.4	Matching Data Representations	700
14.6	Consistency and Semantics	701
14.6.1	File Consistency	701
14.6.2	Random Access vs. Sequential Files	704
14.6.3	Progress	704
14.6.4	Collective File Operations	705
14.6.5	Nonblocking Collective File Operations	705
14.6.6	Type Matching	705
14.6.7	Miscellaneous Clarifications	705
14.6.8	MPI_Offset Type	706
14.6.9	Logical vs. Physical File Layout	706
14.6.10	File Size	706
14.6.11	Examples	707
	Asynchronous I/O	709
14.7	I/O Error Handling	711
14.8	I/O Error Classes	712
14.9	Examples	712
14.9.1	Double Buffering with Split Collective I/O	712
14.9.2	Subarray Filetype Constructor	714
15	Tool Support	717
15.1	Introduction	717
15.2	Profiling Interface	717
15.2.1	Requirements	717
15.2.2	Discussion	718
15.2.3	Logic of the Design	718
15.2.4	Miscellaneous Control of Profiling	719
15.2.5	MPI Library Implementation	720

15.2.6	Complications	721
	Multiple Counting	721
	Linker Oddities	721
	Fortran Support Methods	722
15.2.7	Multiple Levels of Interception	722
15.3	The MPI Tool Information Interface	722
15.3.1	Verbosity Levels	724
15.3.2	Binding MPI Tool Information Interface Variables to MPI Objects	724
15.3.3	Convention for Returning Strings	725
15.3.4	Initialization and Finalization	726
15.3.5	Datatype System	727
15.3.6	Control Variables	729
	Control Variable Query Functions	729
	Handle Allocation and Deallocation	733
	Control Variable Access Functions	735
15.3.7	Performance Variables	736
	Performance Variable Classes	736
	Performance Variable Query Functions	738
	Performance Experiment Sessions	741
	Handle Allocation and Deallocation	742
	Starting and Stopping of Performance Variables	743
	Performance Variable Access Functions	744
15.3.8	Events	748
	Event Sources	749
	Callback Safety Requirements	751
	Event Type Query Functions	753
	Handle Allocation and Deallocation	755
	Handling Dropped Events	760
	Reading Event Data	761
	Reading Event Meta Data	762
15.3.9	Variable Categorization	764
	Category Query Functions	764
	Category Member Query Functions	766
15.3.10	Return Codes for the MPI Tool Information Interface	768
15.3.11	Profiling Interface	768
16	Deprecated Interfaces	771
16.1	Deprecated since MPI-2.0	771
16.2	Deprecated since MPI-2.2	774
16.3	Deprecated since MPI-4.0	774
16.4	Deprecated since MPI-4.1	777
17	Removed Interfaces	783
17.1	Removed MPI-1 Bindings	783
	17.1.1 Overview	783
	17.1.2 Removed MPI-1 Functions	783
	17.1.3 Removed MPI-1 Datatypes	783
	17.1.4 Removed MPI-1 Constants	783

17.1.5	Removed MPI-1 Callback Prototypes	783
17.2	C++ Bindings	784
18	Semantic Changes and Warnings	785
18.1	Semantic Changes	785
18.1.1	Semantic Changes Starting in MPI-4.0	785
18.2	Additional Warnings	785
18.2.1	Warnings Starting in MPI-4.1	785
18.2.2	Warnings Starting in MPI-4.0	786
19	Language Bindings	787
19.1	Support for Fortran	787
19.1.1	Overview	787
19.1.2	Fortran Support Through the <code>mpi_f08</code> Module	788
19.1.3	Fortran Support Through the <code>mpi</code> Module	791
19.1.4	Fortran Support Through the <code>mpif.h</code> Include File	793
19.1.5	Interface Specifications, Procedure Names, and the Profiling Inter- face	794
19.1.6	MPI for Different Fortran Standard Versions	799
19.1.7	Requirements on Fortran Compilers	803
19.1.8	Additional Support for Fortran Register-Memory-Synchronization	804
19.1.9	Additional Support for Fortran Numeric Intrinsic Types	805
	Parameterized Datatypes with Specified Precision and Exponent Range	806
	Support for Size-specific MPI Datatypes	810
	Communication With Size-specific Types	812
19.1.10	Problems With Fortran Bindings for MPI	813
19.1.11	Problems Due to Strong Typing	814
19.1.12	Problems Due to Data Copying and Sequence Association with Sub- script Triplets	815
19.1.13	Problems Due to Data Copying and Sequence Association with Vec- tor Subscripts	818
19.1.14	Special Constants	819
19.1.15	Fortran Derived Types	819
19.1.16	Optimization Problems, an Overview	821
19.1.17	Problems with Code Movement and Register Optimization	822
	Nonblocking Operations	822
	Persistent Operations	823
	One-sided Communication	823
	MPI_BOTTOM and Combining Independent Variables in Datatypes Solutions	823
	The Fortran ASYNCHRONOUS Attribute	825
	Calling MPI_F_SYNC_REG	826
	A User Defined Routine Instead of MPI_F_SYNC_REG	827
	Module Variables and COMMON Blocks	828
	The (Poorly Performing) Fortran VOLATILE Attribute	828
	The Fortran TARGET Attribute	828
19.1.18	Temporary Data Movement and Temporary Memory Modification	829

19.1.19	Permanent Data Movement	831
19.1.20	Comparison with C	832
19.2	Support for Large Count and Large Byte Displacement	833
19.3	Language Interoperability	834
19.3.1	Introduction	834
19.3.2	Assumptions	834
19.3.3	Initialization	835
	Concerns specific to the World Model	835
	Concerns specific to the Sessions Model	835
	Concerns common to both the World Model and the Sessions Model	836
19.3.4	Transfer of Handles	836
19.3.5	Status	838
19.3.6	MPI Opaque Objects	840
	Datatypes	840
	Callback Functions	842
	Error Handlers	842
	Reduce Operations	842
19.3.7	Attributes	843
19.3.8	Extra-State	846
19.3.9	Constants	847
19.3.10	Interlanguage Communication	847
A	Language Bindings Summary	849
A.1	Defined Values and Handles	849
A.1.1	Defined Constants	849
A.1.2	Types	865
A.1.3	Prototype Definitions	866
	C Bindings	866
	Fortran 2008 Bindings with the <code>mpi_f08</code> Module	867
	Fortran Bindings with <code>mpif.h</code> or the <code>mpi</code> Module	870
A.1.4	Deprecated Prototype Definitions	872
A.1.5	String Values	873
	Default Communicator Names	873
	Default Datatype Names	873
	Default Window Names	873
	Reserved Data Representations	873
	Process Set Names	873
	Info Keys	873
	Info Values	874
A.2	Summary of the Semantics of all Op.-Related Routines	875
A.3	C Bindings	884
A.3.1	Point-to-Point Communication C Bindings	884
A.3.2	Partitioned Communication C Bindings	888
A.3.3	Datatypes C Bindings	888
A.3.4	Collective Communication C Bindings	891
A.3.5	Groups, Contexts, Communicators, and Caching C Bindings	899
A.3.6	Virtual Topologies for MPI Processes C Bindings	902
A.3.7	MPI Environmental Management C Bindings	906

A.3.8	The Info Object C Bindings	907
A.3.9	Process Creation and Management C Bindings	907
A.3.10	One-Sided Communications C Bindings	908
A.3.11	External Interfaces C Bindings	911
A.3.12	I/O C Bindings	912
A.3.13	Language Bindings C Bindings	916
A.3.14	Tools / Profiling Interface C Bindings	917
A.3.15	Tools / MPI Tool Information Interface C Bindings	917
A.3.16	Deprecated C Bindings	919
A.4	Fortran 2008 Bindings with the <code>mpi_f08</code> Module	921
A.4.1	Point-to-Point Communication Fortran 2008 Bindings	921
A.4.2	Partitioned Communication Fortran 2008 Bindings	933
A.4.3	Datatypes Fortran 2008 Bindings	934
A.4.4	Collective Communication Fortran 2008 Bindings	941
A.4.5	Groups, Contexts, Communicators, and Caching Fortran 2008 Bindings	963
A.4.6	Virtual Topologies for MPI Processes Fortran 2008 Bindings	970
A.4.7	MPI Environmental Management Fortran 2008 Bindings	979
A.4.8	The Info Object Fortran 2008 Bindings	982
A.4.9	Process Creation and Management Fortran 2008 Bindings	983
A.4.10	One-Sided Communications Fortran 2008 Bindings	986
A.4.11	External Interfaces Fortran 2008 Bindings	993
A.4.12	I/O Fortran 2008 Bindings	994
A.4.13	Language Bindings Fortran 2008 Bindings	1007
A.4.14	Tools / Profiling Interface Fortran 2008 Bindings	1007
A.4.15	Deprecated Fortran 2008 Bindings	1008
A.5	Fortran Bindings with <code>mpif.h</code> or the <code>mpi</code> Module	1009
A.5.1	Point-to-Point Communication Fortran Bindings	1009
A.5.2	Partitioned Communication Fortran Bindings	1013
A.5.3	Datatypes Fortran Bindings	1013
A.5.4	Collective Communication Fortran Bindings	1016
A.5.5	Groups, Contexts, Communicators, and Caching Fortran Bindings	1021
A.5.6	Virtual Topologies for MPI Processes Fortran Bindings	1025
A.5.7	MPI Environmental Management Fortran Bindings	1029
A.5.8	The Info Object Fortran Bindings	1031
A.5.9	Process Creation and Management Fortran Bindings	1031
A.5.10	One-Sided Communications Fortran Bindings	1033
A.5.11	External Interfaces Fortran Bindings	1037
A.5.12	I/O Fortran Bindings	1038
A.5.13	Language Bindings Fortran Bindings	1042
A.5.14	Tools / Profiling Interface Fortran Bindings	1043
A.5.15	Deprecated Fortran Bindings	1043
B	Change-Log	1045
B.1	Changes from Version 4.0 to Version 4.1	1045
B.1.1	Fixes to Errata in Previous Versions of MPI	1045
B.1.2	Changes in MPI-4.1	1047

B.2	Changes from Version 3.1 to Version 4.0	1049
B.2.1	Fixes to Errata in Previous Versions of MPI	1049
B.2.2	Changes in MPI-4.0	1050
B.3	Changes from Version 3.0 to Version 3.1	1053
B.3.1	Fixes to Errata in Previous Versions of MPI	1053
B.3.2	Changes in MPI-3.1	1055
B.4	Changes from Version 2.2 to Version 3.0	1056
B.4.1	Fixes to Errata in Previous Versions of MPI	1056
B.4.2	Changes in MPI-3.0	1056
B.5	Changes from Version 2.1 to Version 2.2	1061
B.6	Changes from Version 2.0 to Version 2.1	1064
	Bibliography	1069
	General Index	1075
	Examples Index	1084
	MPI Constant and Predefined Handle Index	1103
	MPI Declarations Index	1110
	MPI Callback Function Prototype Index	1111
	MPI Function Index	1112

DRAFT

List of Figures

2.1	State transition diagram for blocking operations	12
2.2	State transition diagram for nonblocking operations	12
2.3	State transition diagram for persistent operations	12
6.1	Collective communications, an overview	191
6.2	Inter-communicator allgather	195
6.3	Inter-communicator reduce-scatter	195
6.4	Gather example	202
6.5	Gatherv example with strides	203
6.6	Gatherv example, 2-dimensional	204
6.7	Gatherv example, 2-dimensional, subarrays with different sizes	205
6.8	Gatherv example, 2-dimensional, subarrays with different sizes and strides	207
6.9	Scatter example	212
6.10	Scatterv example with strides	213
6.11	Scatterv example with different strides and counts	214
6.12	Race conditions with point-to-point and collective communications	299
6.13	Overlapping communicators example	303
7.1	Inter-communicator creation using <code>MPI_COMM_CREATE</code>	326
7.2	Inter-communicator construction with <code>MPI_COMM_SPLIT</code>	331
7.3	Recursive communicator creation with <code>MPI_COMM_SPLIT_TYPE</code>	337
7.4	Three-group pipeline	357
7.5	Three-group ring	358
8.1	Neighborhood gather communication example	413
8.2	Cartesian neighborhood allgather example for 3 and 1 processes in a dimension	413
8.3	Cartesian neighborhood alltoall example for 3 and 1 MPI processes in a dimension	418
8.4	Set-up of MPI process structure for two-dimensional parallel Poisson solver	437
8.5	Communication routine with local data copying and sparse neighborhood alltoall	438
8.6	Communication routine with sparse neighborhood alltoallw and without local data copying	439
8.7	Two-dimensional parallel Poisson solver with persistent sparse neighborhood alltoallw and without local data copying	440
9.1	Diagram for deciding which error handler is invoked.	450
11.1	Session handle to communicator	494
11.2	Process set examples	498

12.1 Schematic description of the public/private window operations in the MPI_WIN_SEPARATE memory model for two overlapping windows	587
12.2 Active target communication	590
12.3 Active target communication, with weak synchronization	591
12.4 Passive target communication	592
12.5 Active target communication with several MPI processes	595
12.6 Symmetric communication	615
12.7 Deadlock situation	616
12.8 No deadlock	616
14.1 Etypes and filetypes	638
14.2 Partitioning a file among parallel processes	638
14.3 Displacements	651
14.4 Example array file layout	714
14.5 Example local array filetype for process 1	714
19.1 Status conversion routines	839

DRAFT

List of Tables

2.1	Deprecated and removed constructs	24
3.1	Predefined MPI datatypes corresponding to Fortran datatypes	33
3.2	Predefined MPI datatypes corresponding to C datatypes	34
3.3	Predefined MPI datatypes corresponding to both C and Fortran datatypes	35
3.4	Predefined MPI datatypes corresponding to C++ datatypes	35
7.1	MPI_COMM_* function behavior (in inter-communication mode)	352
9.1	Error classes (Part 1)	462
9.2	Error classes (Part 2)	463
11.1	List of MPI Functions that can be called at any time within an MPI program, including prior to MPI initialization and following MPI finalization	507
12.1	C types of attribute value argument to MPI_WIN_GET_ATTR and MPI_WIN_SET_ATTR	563
12.2	Error classes in one-sided communication routines	605
14.1	Data access routines	653
14.2	"external32" sizes of predefined datatypes	695
14.3	"external32" sizes of optional datatypes	696
14.4	"external32" sizes of C++ datatypes	696
14.5	I/O error classes	712
15.1	MPI tool information interface verbosity levels	724
15.2	Constants to identify associations of variables	725
15.3	MPI datatypes that can be used by the MPI tool information interface	727
15.4	Scopes for control variables	732
15.5	Hierarchy of safety requirement levels for event callback routines	751
15.6	List of MPI functions that when called from within a callback function may not return MPI_T_ERR_NOT_ACCESSIBLE	752
15.7	Return codes used in procedures of the MPI tool information interface.	769
17.1	Removed MPI-1 functions and their replacements	783
17.2	Removed MPI-1 datatypes. The indicated routine may be used for changing the lower and upper bound respectively.	784
17.3	Removed MPI-1 constants	784
17.4	Removed MPI-1 callback prototypes and their replacements	784
19.1	Specific Fortran procedure names and related calling conventions	795
19.2	Occurrence of Fortran optimization problems	821

Acknowledgments

This document is the product of a number of distinct efforts in four distinct phases: one for each of MPI-1, MPI-2, MPI-3, and MPI-4. This section describes these in historical order, starting with MPI-1. Some efforts, particularly parts of MPI-2, had distinct groups of individuals associated with them, and these efforts are detailed separately.

This document represents the work of many people who have served on the MPI Forum. The meetings have been attended by dozens of people from many parts of the world. It is the hard and dedicated work of this group that has led to the MPI standard.

The technical development was carried out by subgroups, whose work was reviewed by the full committee. During the period of development of the Message-Passing Interface (MPI), many people helped with this effort.

Those who served as primary coordinators in MPI-1.0 and MPI-1.1 are:

- Jack Dongarra, David Walker, Conveners and Meeting Chairs
- Ewing Lusk, Bob Knighten, Minutes
- Marc Snir, William Gropp, Ewing Lusk, Point-to-Point Communication
- Al Geist, Marc Snir, Steve Otto, Collective Communication
- Steve Otto, Editor
- Rolf Hempel, Process Topologies
- Ewing Lusk, Language Binding
- William Gropp, Environmental Management
- James Cownie, Profiling
- Tony Skjellum, Lyndon Clarke, Marc Snir, Richard Littlefield, Mark Sears, Groups, Contexts, and Communicators
- Steven Huss-Lederman, Initial Implementation Subset

The following list includes some of the active participants in the MPI-1.0 and MPI-1.1 process not mentioned above.

Ed Anderson	Robert Babb	Joe Baron	Eric Barszcz	1
Scott Berryman	Rob Bjornson	Nathan Doss	Anne Elster	2
Jim Feeney	Vince Fernando	Sam Fineberg	Jon Flower	3
Daniel Frye	Ian Glendinning	Adam Greenberg	Robert Harrison	4
Leslie Hart	Tom Haupt	Don Heller	Tom Henderson	5
Alex Ho	C.T. Howard Ho	Gary Howell	John Kapenga	6
James Kohl	Susan Krauss	Bob Leary	Arthur Maccabe	7
Peter Madams	Alan Mainwaring	Oliver McBryan	Phil McKinley	8
Charles Mosher	Dan Nessett	Peter Pacheco	Howard Palmer	9
Paul Pierce	Sanjay Ranka	Peter Rigsbee	Arch Robison	10
Erich Schikuta	Ambuj Singh	Alan Sussman	Robert Tomlinson	11
Robert G. Voigt	Dennis Weeks	Stephen Wheat	Steve Zenith	12

The University of Tennessee and Oak Ridge National Laboratory made the draft available by anonymous FTP mail servers and were instrumental in distributing the document.

The work on the MPI-1 standard was supported in part by ARPA and NSF under grant ASC-9310330, the National Science Foundation Science and Technology Center Cooperative Agreement No. CCR-8809615, and by the Commission of the European Community through Esprit project P6643 (PPPE).

MPI-1.2 and MPI-2.0:

Those who served as primary coordinators in MPI-1.2 and MPI-2.0 are:

- Ewing Lusk, Convener and Meeting Chair
- Steve Huss-Lederman, Editor
- Ewing Lusk, Miscellany
- Bill Saphir, Process Creation and Management
- Marc Snir, One-Sided Communications
- William Gropp and Anthony Skjellum, Extended Collective Operations
- Steve Huss-Lederman, External Interfaces
- Bill Nitzberg, I/O
- Andrew Lumsdaine, Bill Saphir, and Jeffrey M. Squyres, Language Bindings
- Anthony Skjellum and Arkady Kanevsky, Real-Time

The following list includes some of the active participants who attended MPI-2 Forum meetings and are not mentioned above.

1	Greg Astfalk	Robert Babb	Ed Benson	Rajesh Bordawekar
2	Pete Bradley	Peter Brennan	Ron Brightwell	Maciej Brodowicz
3	Eric Brunner	Greg Burns	Margaret Cahir	Pang Chen
4	Ying Chen	Albert Cheng	Yong Cho	Joel Clark
5	Lyndon Clarke	Laurie Costello	Dennis Cottel	Jim Cownie
6	Zhenqian Cui	Suresh Damodaran-Kamal		Raja Daoud
7	Judith Devaney	David DiNucci	Doug Doeffler	Jack Dongarra
8	Terry Dontje	Nathan Doss	Anne Elster	Mark Fallon
9	Karl Feind	Sam Fineberg	Craig Fischberg	Stephen Fleischman
10	Ian Foster	Hubertus Franke	Richard Frost	Al Geist
11	Robert George	David Greenberg	John Hagedorn	Kei Harada
12	Leslie Hart	Shane Hebert	Rolf Hempel	Tom Henderson
13	Alex Ho	Hans-Christian Hoppe	Joefon Jann	Terry Jones
14	Karl Kesselman	Koichi Konishi	Susan Kraus	Steve Kubica
15	Steve Landherr	Mario Lauria	Mark Law	Juan Leon
16	Lloyd Lewins	Ziyang Lu	Bob Madahar	Peter Madams
17	John May	Oliver McBryan	Brian McCandless	Tyce McLarty
18	Thom McMahon	Harish Nag	Nick Nevin	Jarek Nieplocha
19	Ron Oldfield	Peter Ossadnik	Steve Otto	Peter Pacheco
20	Yoonho Park	Perry Partow	Pratap Pattnaik	Elsie Pierce
21	Paul Pierce	Heidi Poxon	Jean-Pierre Prost	Boris Protopopov
22	James Pruyve	Rolf Rabenseifner	Joe Rieken	Peter Rigsbee
23	Tom Robey	Anna Rounbehler	Nobutoshi Sagawa	Arindam Saha
24	Eric Salo	Darren Sanders	Eric Sharakan	Andrew Sherman
25	Fred Shirley	Lance Shuler	A. Gordon Smith	Ian Stockdale
26	David Taylor	Stephen Taylor	Greg Tensa	Rajeev Thakur
27	Marydell Tholburn	Dick Treumann	Simon Tsang	Manuel Ujaldon
28	David Walker	Jerrell Watts	Klaus Wolf	Parkson Wong
29	Dave Wright			

30
31 The MPI Forum also acknowledges and appreciates the valuable input from people via
32 e-mail and in person.

33 The following institutions supported the MPI-2 effort through time and travel support
34 for the people listed above.

- 35 Argonne National Laboratory
- 36 Bolt, Beranek, and Newman
- 37 California Institute of Technology
- 38 Center for Computing Sciences
- 39 Convex Computer Corporation
- 40 Cray Research
- 41 Digital Equipment Corporation
- 42 Dolphin Interconnect Solutions, Inc.
- 43 Edinburgh Parallel Computing Centre
- 44 General Electric Company
- 45 German National Research Center for Information Technology
- 46 Hewlett-Packard
- 47 Hitachi
- 48 Hughes Aircraft Company

Intel Corporation	1
International Business Machines	2
Khoral Research	3
Lawrence Livermore National Laboratory	4
Los Alamos National Laboratory	5
MPI Software Techology, Inc.	6
Mississippi State University	7
NEC Corporation	8
National Aeronautics and Space Administration	9
National Energy Research Scientific Computing Center	10
National Institute of Standards and Technology	11
National Oceanic and Atmospheric Administration	12
Oak Ridge National Laboratory	13
The Ohio State University	14
PALLAS GmbH	15
Pacific Northwest National Laboratory	16
Pratt & Whitney	17
San Diego Supercomputer Center	18
Sanders, A Lockheed-Martin Company	19
Sandia National Laboratories	20
Schlumberger	21
Scientific Computing Associates, Inc.	22
Silicon Graphics Incorporated	23
Sky Computers	24
Sun Microsystems Computer Corporation	25
Syracuse University	26
The MITRE Corporation	27
Thinking Machines Corporation	28
United States Navy	29
University of Colorado	30
University of Denver	31
University of Houston	32
University of Illinois	33
University of Maryland	34
University of Notre Dame	35
University of San Fransisco	36
University of Stuttgart Computing Center	37
University of Wisconsin	38

MPI-2 operated on a very tight budget (in reality, it had no budget when the first meeting was announced). Many institutions helped the MPI-2 effort by supporting the efforts and travel of the members of the MPI Forum. Direct support was given by NSF and DARPA under NSF contract CDA-9115428 for travel by U.S. academic participants and Esprit under project HPC Standards (21111) for European participants.

1 MPI-1.3 and MPI-2.1:

2 The editors and organizers of the combined documents have been:

- 3
- 4 • Richard Graham, Convener and Meeting Chair
 - 5
 - 6 • Jack Dongarra, Steering Committee
 - 7
 - 8 • Al Geist, Steering Committee
 - 9
 - 10 • William Gropp, Steering Committee
 - 11
 - 12 • Rainer Keller, Merge of MPI-1.3
 - 13
 - 14 • Andrew Lumsdaine, Steering Committee
 - 15
 - 16 • Ewing Lusk, Steering Committee, MPI-1.1-Errata (Oct. 12, 1998) MPI-2.1-Errata
 - 17 Ballots 1, 2 (May 15, 2002)
 - 18
 - 19 • Rolf Rabenseifner, Steering Committee, Merge of MPI-2.1 and MPI-2.1-Errata Ballots
 - 20 3, 4 (2008)

21 All chapters have been revisited to achieve a consistent MPI-2.1 text. Those who served as authors for the necessary modifications are:

- 22 • William Gropp, Front Matter, Introduction, and Bibliography
- 23
- 24 • Richard Graham, Point-to-Point Communication
- 25
- 26 • Adam Moody, Collective Communication
- 27
- 28 • Richard Treumann, Groups, Contexts, and Communicators
- 29
- 30 • Jesper Larsson Träff, Process Topologies, Info-Object, and One-Sided Communica-
- 31 tions
- 32
- 33 • George Bosilca, Environmental Management
- 34
- 35 • David Solt, Process Creation and Management
- 36
- 37 • Bronis R. de Supinski, External Interfaces, and Profiling
- 38
- 39 • Rajeev Thakur, I/O
- 40
- 41 • Jeffrey M. Squyres, Language Bindings and MPI-2.1 Secretary
- 42
- 43 • Rolf Rabenseifner, Deprecated Functions and Annex Change-Log
- 44
- 45 • Alexander Supalov and Denis Nagorny, Annex Language Bindings

46 The following list includes some of the active participants who attended MPI-2 Forum

47 meetings and in the e-mail discussions of the errata items and are not mentioned above.

48

Pavan Balaji	Purushotham V. Bangalore	Brian Barrett	1
Richard Barrett	Christian Bell	Robert Blackmore	2
Gil Bloch	Ron Brightwell	Jeffrey Brown	3
Darius Buntinas	Jonathan Carter	Nathan DeBardeleben	4
Terry Dontje	Gabor Dozsa	Edric Ellis	5
Karl Feind	Edgar Gabriel	Patrick Geoffray	6
David Gingold	Dave Goodell	Erez Haba	7
Robert Harrison	Thomas Herault	Steve Hodson	8
Torsten Hoefler	Joshua Hursey	Yann Kalemkarian	9
Matthew Koop	Quincey Koziol	Sameer Kumar	10
Miron Livny	Kannan Narasimhan	Mark Pagel	11
Avneesh Pant	Steve Poole	Howard Pritchard	12
Craig Rasmussen	Hubert Ritzdorf	Rob Ross	13
Tony Skjellum	Brian Smith	Vinod Tipparaju	14
Jesper Larsson Träff	Keith Underwood		15

The MPI Forum also acknowledges and appreciates the valuable input from people via e-mail and in person.

The following institutions supported the MPI-2 effort through time and travel support for the people listed above.

Argonne National Laboratory	21
Bull	22
Cisco Systems, Inc.	23
Cray Inc.	24
The HDF Group	25
Hewlett-Packard	26
IBM T.J. Watson Research	27
Indiana University	28
Institut National de Recherche en Informatique et Automatique (Inria)	29
Intel Corporation	30
Lawrence Berkeley National Laboratory	31
Lawrence Livermore National Laboratory	32
Los Alamos National Laboratory	33
Mathworks	34
Mellanox Technologies	35
Microsoft	36
Myricom	37
NEC Laboratories Europe, NEC Europe Ltd.	38
Oak Ridge National Laboratory	39
The Ohio State University	40
Pacific Northwest National Laboratory	41
QLogic Corporation	42
Sandia National Laboratories	43
SiCortex	44
Silicon Graphics Incorporated	45
Sun Microsystems, Inc.	46
University of Alabama at Birmingham	47
University of Houston	48

1 University of Illinois at Urbana-Champaign
2 University of Stuttgart, High Performance Computing Center Stuttgart (HLRS)
3 University of Tennessee, Knoxville
4 University of Wisconsin

5 Funding for the MPI Forum meetings was partially supported by award #CCF-0816909
6 from the National Science Foundation. In addition, the HDF Group provided travel support
7 for one U.S. academic.
8

9
10 **MPI-2.2:**

11 All chapters have been revisited to achieve a consistent MPI-2.2 text. Those who served as
12 authors for the necessary modifications are:
13

- 14 • William Gropp, Front Matter, Introduction, and Bibliography; MPI-2.2 Chair.
- 15 • Richard Graham, Point-to-Point Communication and Datatypes
- 16 • Adam Moody, Collective Communication
- 17 • Torsten Hoefler, Collective Communication and Process Topologies
- 18 • Richard Treumann, Groups, Contexts, and Communicators
- 19 • Jesper Larsson Träff, Process Topologies, Info-Object and One-Sided Communications
- 20 • George Bosilca, Datatypes and Environmental Management
- 21 • David Solt, Process Creation and Management
- 22 • Bronis R. de Supinski, External Interfaces, and Profiling
- 23 • Rajeev Thakur, I/O
- 24 • Jeffrey M. Squyres, Language Bindings and MPI-2.2 Secretary
- 25 • Rolf Rabenseifner, Deprecated Functions, Annex Change-Log, and Annex Language Bindings
- 26 • Alexander Supalov, Annex Language Bindings

27
28 The following list includes some of the active participants who attended MPI-2 Forum
29 meetings and in the e-mail discussions of the errata items and are not mentioned above.
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

Pavan Balaji	Purushotham V. Bangalore	Brian Barrett	1
Richard Barrett	Christian Bell	Robert Blackmore	2
Gil Bloch	Ron Brightwell	Greg Bronevetsky	3
Jeff Brown	Darius Buntinas	Jonathan Carter	4
Nathan DeBardeleben	Terry Dontje	Gabor Dozsa	5
Edric Ellis	Karl Feind	Edgar Gabriel	6
Patrick Geoffray	Johann George	David Gingold	7
David Goodell	Erez Haba	Robert Harrison	8
Thomas Herault	Marc-André Hermanns	Steve Hodson	9
Joshua Hursey	Yutaka Ishikawa	Bin Jia	10
Hideyuki Jitsumoto	Terry Jones	Yann Kalemkarian	11
Ranier Keller	Matthew Koop	Quincey Koziol	12
Manojkumar Krishnan	Sameer Kumar	Miron Livny	13
Andrew Lumsdaine	Miao Luo	Ewing Lusk	14
Timothy I. Mattox	Kannan Narasimhan	Mark Pagel	15
Avneesh Pant	Steve Poole	Howard Pritchard	16
Craig Rasmussen	Hubert Ritzdorf	Rob Ross	17
Martin Schulz	Pavel Shamis	Galen Shipman	18
Christian Siebert	Anthony Skjellum	Brian Smith	19
Naoki Sueyasu	Vinod Tipparaju	Keith Underwood	20
Rolf Vandevaart	Abhinav Vishnu	Weikuan Yu	21

The MPI Forum also acknowledges and appreciates the valuable input from people via e-mail and in person.

The following institutions supported the MPI-2.2 effort through time and travel support for the people listed above.

Argonne National Laboratory	27
Auburn University	28
Bull	29
Cisco Systems, Inc.	30
Cray Inc.	31
Forschungszentrum Jülich	32
Fujitsu	33
The HDF Group	34
Hewlett-Packard	35
International Business Machines	36
Indiana University	37
Institut National de Recherche en Informatique et Automatique (Inria)	38
Institute for Advanced Science & Engineering Corporation	39
Intel Corporation	40
Lawrence Berkeley National Laboratory	41
Lawrence Livermore National Laboratory	42
Los Alamos National Laboratory	43
Mathworks	44
Mellanox Technologies	45
Microsoft	46
Myricom	47
NEC Corporation	48

1 Oak Ridge National Laboratory
2 The Ohio State University
3 Pacific Northwest National Laboratory
4 QLogic Corporation
5 RunTime Computing Solutions, LLC
6 Sandia National Laboratories
7 SiCortex, Inc.
8 Silicon Graphics Inc.
9 Sun Microsystems, Inc.
10 Tokyo Institute of Technology
11 University of Alabama at Birmingham
12 University of Houston
13 University of Illinois at Urbana-Champaign
14 University of Stuttgart, High Performance Computing Center Stuttgart (HLRS)
15 University of Tennessee, Knoxville
16 University of Tokyo
17 University of Wisconsin

18 Funding for the MPI Forum meetings was partially supported by awards #CCF-0816909
19 and #CCF-1144042 from the National Science Foundation. In addition, the HDF Group
20 provided travel support for one U.S. academic.
21

22 MPI-3.0:

23 MPI-3.0 is a significant effort to extend and modernize the MPI standard.
24 The editors and organizers of the MPI-3.0 have been:

- 25 ● William Gropp, Steering Committee, Front Matter, Introduction, Groups, Contexts,
26 and Communicators, One-Sided Communications, and Bibliography
- 27 ● Richard Graham, Steering Committee, Point-to-Point Communication, Meeting Con-
28 vener, and MPI-3.0 Chair
- 29 ● Torsten Hoefler, Collective Communication, One-Sided Communications, and Process
30 Topologies
- 31 ● George Bosilca, Datatypes and Environmental Management
- 32 ● David Solt, Process Creation and Management
- 33 ● Bronis R. de Supinski, External Interfaces and Tool Support
- 34 ● Rajeev Thakur, I/O and One-Sided Communications
- 35 ● Darius Buntinas, Info Object
- 36 ● Jeffrey M. Squyres, Language Bindings and MPI-3.0 Secretary
- 37 ● Rolf Rabenseifner, Steering Committee, Terms and Definitions, and Fortran Bindings,
38 Deprecated Functions, Annex Change-Log, and Annex Language Bindings
- 39 ● Craig Rasmussen, Fortran Bindings

The following list includes some of the active participants who attended MPI-3 Forum meetings or participated in the e-mail discussions and who are not mentioned above.

Tatsuya Abe	Tomoya Adachi	Sadaf Alam	1
Reinhold Bader	Pavan Balaji	Purushotham V. Bangalore	2
Brian Barrett	Richard Barrett	Robert Blackmore	3
Aurélien Bouteiller	Ron Brightwell	Greg Bronevetsky	4
Jed Brown	Darius Buntinas	Devendar Bureddy	5
Arno Candel	George Carr	Mohamad Chaarawi	6
Raghunath Raja Chandrasekar	James Dinan	Terry Dontje	7
Edgar Gabriel	Balazs Gerofi	Brice Goglin	8
David Goodell	Manjunath Gorentla	Erez Haba	9
Jeff Hammond	Thomas Herault	Marc-André Hermanns	10
Jennifer Herrett-Skjellum	Nathan Hjelm	Atsushi Hori	11
Joshua Hursey	Marty Itzkowitz	Yutaka Ishikawa	12
Nysal Jan	Bin Jia	Hideyuki Jitsumoto	13
Yann Kalemkarian	Krishna Kandalla	Takahiro Kawashima	14
Chulho Kim	Dries Kimpe	Christof Klausecker	15
Alice Koniges	Quincey Koziol	Dieter Kranzlmüller	16
Manojkumar Krishnan	Sameer Kumar	Eric Lantz	17
Jay Lofstead	Bill Long	Andrew Lumsdaine	18
Miao Luo	Ewing Lusk	Adam Moody	19
Nick M. Maclaren	Amith Mamidala	Guillaume Mercier	20
Scott McMillan	Douglas Miller	Kathryn Mohror	21
Tim Murray	Tomotake Nakamura	Takeshi Nanri	22
Steve Oyanagi	Mark Pagel	Swann Perarnau	23
Sreeram Potluri	Howard Pritchard	Rolf Riesen	24
Hubert Ritzdorf	Kuninobu Sasaki	Timo Schneider	25
Martin Schulz	Gilad Shainer	Christian Siebert	26
Anthony Skjellum	Brian Smith	Marc Snir	27
Raffaele Giuseppe Solca	Shinji Sumimoto	Alexander Supalov	28
Sayantana Sur	Masamichi Takagi	Fabian Tillier	29
Vinod Tipparaju	Jesper Larsson Träff	Richard Treumann	30
Keith Underwood	Rolf Vandevaart	Anh Vo	31
Abhinav Vishnu	Min Xie	Enqiang Zhou	32

The MPI Forum also acknowledges and appreciates the valuable input from people via e-mail and in person.

The MPI Forum also thanks those that provided feedback during the public comment period. In particular, the Forum would like to thank Jeremiah Wilcock for providing detailed comments on the entire draft standard.

The following institutions supported the MPI-3 effort through time and travel support for the people listed above.

Argonne National Laboratory
Bull
Cisco Systems, Inc.
Cray Inc.
CSCS

1 ETH Zurich
2 Fujitsu Ltd.
3 German Research School for Simulation Sciences
4 The HDF Group
5 Hewlett-Packard
6 International Business Machines
7 IBM India Private Ltd
8 Indiana University
9 Institut National de Recherche en Informatique et Automatique (Inria)
10 Institute for Advanced Science & Engineering Corporation
11 Intel Corporation
12 Lawrence Berkeley National Laboratory
13 Lawrence Livermore National Laboratory
14 Los Alamos National Laboratory
15 Mellanox Technologies, Inc.
16 Microsoft Corporation
17 NEC Corporation
18 National Oceanic and Atmospheric Administration, Global Systems Division
19 NVIDIA Corporation
20 Oak Ridge National Laboratory
21 The Ohio State University
22 Oracle America
23 Platform Computing
24 RIKEN AICS
25 RunTime Computing Solutions, LLC
26 Sandia National Laboratories
27 Technical University of Chemnitz
28 Tokyo Institute of Technology
29 University of Alabama at Birmingham
30 University of Chicago
31 University of Houston
32 University of Illinois at Urbana-Champaign
33 University of Stuttgart, High Performance Computing Center Stuttgart (HLRS)
34 University of Tennessee, Knoxville
35 University of Tokyo

36 Funding for the MPI Forum meetings was partially supported by awards #CCF-0816909
37 and #CCF-1144042 from the National Science Foundation. In addition, the HDF Group
38 and Sandia National Laboratories provided travel support for one U.S. academic each.
39

40 MPI-3.1:

41 MPI-3.1 is a minor update to the MPI standard.
42 The editors and organizers of the MPI-3.1 have been:

- 43 • Martin Schulz, MPI-3.1 Chair
- 44 • William Gropp, Steering Committee, Front Matter, Introduction, One-Sided Commu-
45 nications, and Bibliography; Overall Editor

- Rolf Rabenseifner, Steering Committee, Terms and Definitions, and Fortran Bindings, Deprecated Functions, Annex Change-Log, and Annex Language Bindings 1
- Richard L. Graham, Steering Committee, Meeting Convener 2
- Jeffrey M. Squyres, Language Bindings and MPI-3.1 Secretary 3
- Daniel Holmes, Point-to-Point Communication 4
- George Bosilca, Datatypes and Environmental Management 5
- Torsten Hoefler, Collective Communication and Process Topologies 6
- Pavan Balaji, Groups, Contexts, and Communicators, and External Interfaces 7
- Jeff Hammond, The Info Object 8
- David Solt, Process Creation and Management 9
- Quincey Koziol, I/O 10
- Kathryn Mohror, Tool Support 11
- Rajeev Thakur, One-Sided Communications 12

The following list includes some of the active participants who attended MPI Forum meetings or participated in the e-mail discussions. 13

Charles Archer	Pavan Balaji	Purushotham V. Bangalore	14
Brian Barrett	Wesley Bland	Michael Blocksome	15
George Bosilca	Aurélien Bouteiller	Devendar Bureddy	16
Yohann Burette	Mohamad Chaarawi	Alexey Cheptsov	17
James Dinan	Dmitry Durnov	Thomas Francois	18
Edgar Gabriel	Todd Gamblin	Balazs Gerofi	19
Paddy Gillies	David Goodell	Manjunath Gorentla Venkata	20
Richard L. Graham	Ryan E. Grant	William Gropp	21
Khaled Hamidouche	Jeff Hammond	Amin Hassani	22
Marc-André Hermanns	Nathan Hjelm	Torsten Hoefler	23
Daniel Holmes	Atsushi Hori	Yutaka Ishikawa	24
Hideyuki Jitsumoto	Jithin Jose	Krishna Kandalla	25
Christos Kavouklis	Takahiro Kawashima	Chulho Kim	26
Michael Knobloch	Alice Koniges	Quincey Koziol	27
Sameer Kumar	Joshua Ladd	Ignacio Laguna	28
Huiwei Lu	Guillaume Mercier	Kathryn Mohror	29
Adam Moody	Tomotake Nakamura	Takeshi Nanri	30
Steve Oyanagi	Antonio J. Peña	Sreeram Potluri	31
Howard Pritchard	Rolf Rabenseifner	Nicholas Radcliffe	32
Ken Raffenetti	Raghunath Raja	Craig Rasmussen	33
Davide Rossetti	Kento Sato	Martin Schulz	34
Sangmin Seo	Christian Siebert	Anthony Skjellum	35
Brian Smith	David Solt	Jeffrey M. Squyres	36

1 Hari Subramoni Shinji Sumimoto Alexander Supalov
2 Bronis R. de Supinski Sayantan Sur Masamichi Takagi
3 Keita Teranishi Rajeev Thakur Fabian Tillier
4 Yuichi Tsujita Geoffroy Vallée Rolf vandeVaart
5 Akshay Venkatesh Jerome Vienne Venkat Vishwanath
6 Anh Vo Huseyin S. Yildiz Junchao Zhang
7 Xin Zhao

8
9 The MPI Forum also acknowledges and appreciates the valuable input from people via
10 e-mail and in person.

11 The following institutions supported the MPI-3.1 effort through time and travel support
12 for the people listed above.

13 Argonne National Laboratory
14 Auburn University
15 Cisco Systems, Inc.
16 Cray
17 EPCC, The University of Edinburgh
18 ETH Zurich
19 Forschungszentrum Jülich
20 Fujitsu
21 German Research School for Simulation Sciences
22 The HDF Group
23 International Business Machines
24 Institut National de Recherche en Informatique et Automatique (Inria)
25 Intel Corporation
26 Kyushu University
27 Lawrence Berkeley National Laboratory
28 Lawrence Livermore National Laboratory
29 Lenovo
30 Los Alamos National Laboratory
31 Mellanox Technologies, Inc.
32 Microsoft Corporation
33 NEC Corporation
34 NVIDIA Corporation
35 Oak Ridge National Laboratory
36 The Ohio State University
37 RIKEN AICS
38 Sandia National Laboratories
39 Texas Advanced Computing Center
40 Tokyo Institute of Technology
41 University of Alabama at Birmingham
42 University of Houston
43 University of Illinois at Urbana-Champaign
44 University of Oregon
45 University of Stuttgart, High Performance Computing Center Stuttgart (HLRS)
46 University of Tennessee, Knoxville
47 University of Tokyo
48

MPI-4.0:

MPI-4.0 is a major update to the MPI standard.

The editors and organizers of the MPI-4.0 have been:

- Martin Schulz, MPI-4.0 Chair, Info Object, External Interfaces
- Richard Graham, MPI-4.0 Treasurer
- Wesley Bland, MPI-4.0 Secretary, Backward Incompatibilities
- William Gropp, MPI-4.0 Editor, Steering Committee, Front Matter, Introduction, One-Sided Communications, and Bibliography
- Rolf Rabenseifner, Steering Committee, Process Topologies, Deprecated Functions, Removed Interfaces, Annex Language Bindings Summary, and Annex Change-Log
- Purushotham V. Bangalore, Language Bindings
- Claudia Blaas-Schenner, Terms and Conventions
- George Bosilca, Datatypes and Environmental Management
- Ryan E. Grant, Partitioned Communication
- Marc-André Hermanns, Tool Support
- Daniel Holmes, Point-to-Point Communication, Sessions
- Guillaume Mercier, Groups, Contexts, Communicators, Caching
- Howard Pritchard, Process Creation and Management
- Anthony Skjellum, Collective Communication, I/O

As part of the development of MPI-4.0, a number of working groups were established. In some cases, the work for these groups overlapped with multiple chapters. The following describes the major working groups and the leaders of those groups:

Collective Communication, Topology, Communicators: Torsten Hoefer, Andrew Lumsdaine, and Anthony Skjellum

Fault Tolerance: Wesley Bland, Aurélien Bouteiller, and Richard Graham

Hardware-Topologies: Guillaume Mercier

Hybrid & Accelerator: Pavan Balaji and James Dinan

Large Counts: Jeff Hammond

Persistence: Anthony Skjellum

Point to Point Communication: Daniel Holmes and Richard Graham

Remote Memory Access: William Gropp and Rajeev Thakur

Semantic Terms: Purushotham V. Bangalore and Rolf Rabenseifner

1 **Sessions:** Daniel Holmes and Howard Pritchard

2 **Tools:** Kathryn Mohror and Marc-André Hermanns

3
4 The following list includes some of the active participants who attended MPI Forum
5 meetings or participated in the e-mail discussions.

6
7 Julien Adam Abdelhalim Amer Charles Archer
8 Ammar Ahmad Awan Pavan Balaji Purushotham V. Bangalore
9 Mohammadreza Bayatpour Jean-Baptiste Besnard Claudia Blaas-Schenner
10 Wesley Bland Gil Bloch George Bosilca
11 Aurélien Bouteiller Ben Bratu Alexander Calvert
12 Nicholas Chaimov Sourav Chakraborty Steffen Christgau
13 Ching-Hsiang Chu Mikhail Chuvelev James Clark
14 Carsten Clauss Isaias Alberto Comprés Ureña
15 Giuseppe Congiu Brandon Cook James Custer
16 Anna Daly Hoang-Vu Dang James Dinan
17 Matthew Dosanjh Murali Emani Christian Engelmann
18 Noah Evans Ana Gainaru Esthela Gallardo
19 Marc Gamell Balmana Balazs Gerofi Salvatore Di Girolamo
20 Brice Goglin Manjunath Gorentla Venkata Richard Graham
21 Ryan E. Grant Stanley Graves William Gropp
22 Siegmar Gross Taylor Groves Yanfei Guo
23 Khaled Hamidouche Jeff Hammond Marc-André Hermanns
24 Nathan Hjelm Torsten Hoefler Daniel Holmes
25 Atsushi Hori Josh Hursey Ilya Ivanov
26 Julien Jaeger Emmanuel Jeannot Sylvain Jauegy
27 Jithin Jose Krishna Kandalla Takahiro Kawashima
28 Chulho Kim Michael Knobloch Alice Koniges
29 Sameer Kumar Kim Kyunghun Ignacio Laguna Peralta
30 Stefan Lankes Tonglin Li Xioyi Lu
31 Kavitha Madhu Alexey Malhanov Ryan Marshall
32 William Marts Guillaume Mercier Ali Mohammed
33 Kathryn Mohror Takeshi Nanri Thomas Naughton
34 Christoph Niethammer Takafumi Nose Lena Oden
35 Steve Oyanagi Guillaume Papauré Ivy Peng
36 Antonio Peña Simon Pickartz Artem Polyakov
37 Sreeram Potluri Howard Pritchard Martina Prugger
38 Marc Pérache Rolf Rabenseifner Nicholas Radcliffe
39 Ken Raffanetti Craig Rasmussen Soren Rasmussen
40 Hubert Ritzdorf Sergio Rivas-Gomez Davide Rossetti
41 Martin Ruefenacht Amit Ruhela Whit Schonbein
42 Joseph Schuchart Martin Schulz Sangmin Seo
43 Sameh Sharkawi Sameer Shende Min Si
44 Anthony Skjellum Brian Smith David Solt
45 Jeffrey M. Squyres Srinivas Sridharan Hari Subramoni
46 Nawrin Sultana Shinji Sumimoto Sayantan Sur
47
48

Hugo Taboada	Keita Teranishi	Rajeev Thakur	1
Keith Underwood	Geoffroy Vallee	Akshay Venkatesh	2
Jerome Vienne	Anh Vo	Justin Wozniak	3
Junchao Zhang	Dong Zhong	Hui Zhou	4

The MPI Forum also acknowledges and appreciates the valuable input from people via e-mail and in person.

The following institutions supported the MPI-4.0 effort through time and travel support for the people listed above.

ATOS	12
Argonne National Laboratory	13
Arm	14
Auburn University	15
Barcelona Supercomputing Center	16
CEA	17
Cisco Systems Inc.	18
Cray Inc.	19
EPCC, The University of Edinburgh	20
ETH Zürich	21
Fujitsu	22
Fulda University of Applied Sciences	23
German Research School for Simulation Sciences	24
Hewlett Packard Enterprise	25
International Business Machines	26
Institut National de Recherche en Informatique et Automatique (Inria)	27
Intel Corporation	28
Jülich Supercomputing Center, Forschungszentrum Jülich	29
KTH Royal Institute of Technology	30
Kyushu University	31
Lawrence Berkeley National Laboratory	32
Lawrence Livermore National Laboratory	33
Lenovo	34
Los Alamos National Laboratory	35
Mellanox Technologies, Inc.	36
Microsoft Corporation	37
NEC Corporation	38
NVIDIA Corporation	39
Oak Ridge National Laboratory	40
PAR-TEC	41
Paratools, Inc.	42
RIKEN AICS (R-CCS as of 2017)	43
RWTH Aachen University	44
Rutgers University	45
Sandia National Laboratories	46
Silicon Graphics, Inc.	47
Technical University of Munich	48

1 The HDF Group
2 The Ohio State University
3 Texas Advanced Computing Center
4 Tokyo Institute of Technology
5 University of Alabama at Birmingham
6 University of Basel, Switzerland
7 University of Houston
8 University of Illinois at Urbana-Champaign and the National Center for Supercomputing Applications
9
10 University of Innsbruck
11 University of Oregon
12 University of Potsdam
13 University of Stuttgart, High Performance Computing Center Stuttgart (HLRS)
14 University of Tennessee, Chattanooga
15 University of Tennessee, Knoxville
16 University of Texas at El Paso
17 University of Tokyo
18 VSC Research Center, TU Wien
19

20 MPI-4.1:

21 MPI-4.1 is a minor update to the MPI standard.

22 The editors and organizers of the MPI-4.1 have been:

- 23 • Martin Schulz, MPI-4.1 Chair, Info Object, External Interfaces
- 24 • Brian Smith, MPI-4.1 Treasurer
- 25 • Wes Bland, MPI-4.1 Secretary, Semantic Changes and Warnings
- 26 • William Gropp, MPI-4.1 Editor, Steering Committee, Front Matter, Introduction, One-Sided Communications, and Bibliography
- 27 • Rolf Rabenseifner, Steering Committee, Process Topologies, Deprecated Functions, Removed Interfaces, Annex Language Bindings Summary, and Annex Change-Log
- 28 • Purushotham V. Bangalore, Language Bindings
- 29 • Claudia Blaas-Schenner, Terms and Conventions
- 30 • George Bosilca, Datatypes and Environmental Management
- 31 • Ryan E. Grant, Partitioned Communication
- 32 • Marc-André Hermans, Tool Support
- 33 • Dan Holmes, Point-to-Point Communication, Sessions
- 34 • Guillaume Mercier, Groups, Contexts, Communicators, Caching
- 35 • Howard Pritchard, Process Creation and Management
- 36 • Anthony Skjellum, Collective Communication, I/O

As part of the development of MPI-4.1, a number of working groups were established or continued from MPI-4.0. In some cases, the work for these groups overlapped with multiple chapters. The following describes the major working groups and the leaders of those groups:

Application Binary Interface (ABI): Jeff Hammond and Quincey Koziol

Collective Communication, Topology, Communicators: Torsten Hoeffler, Andrew Lumsdaine, and Anthony Skjellum

Fault Tolerance: Aurélien Bouteiller and Ignacio Laguna

Hardware & Virtual Topologies: Guillaume Mercier

Hybrid & Accelerator: James Dinan

Languages: Martin Ruefenacht and Tony Skjellum

Remote Memory Access: William Gropp, Joseph Schuchart, and Rajeev Thakur

Semantic Terms: Purushotham V. Bangalore and Rolf Rabenseifner

Sessions: Dan Holmes and Howard Pritchard

Tools: Marc-André Hermanns

The following list includes some of the active participants who attended MPI Forum meetings or participated in the e-mail discussions.

Julien Adam	Charles Archer	Christian Nicole Avans
Purushotham V. Bangalore	Wolfgang Bangerth	Jean-Baptiste Besnard
Claudia Blaas-Schenner	Wes Bland	George Bosilca
Aurélien Bouteiller	Emmanuel Brelle	Patrick Bridges
Alex Brooks	Jed Brown	Oliver Thomson Brown
Simon Byrne	Paul Canat	Ludovic Capelli
Ondřej Čertík	Steffen Christgau	Brandon Cook
Alfredo Correa	Pedro Costa	Lisandro Dalcin
Ali Can Demiralp	Bronis de Supinski	James Dinan
Matthew G. F. Dosanjh	Joe Downs	Dmitry Durnov
Julien Duprat	Victor Eijkhout	Ahmed Eleliemy
Kyle Gerard Felker	Edgar Gabriel	Maria J. Garzaran
Florent Germain	Sayan Ghosh	Thomas Gillis
Ryan Grant	William Gropp	Yanfei Guo
Samuel K. Gutierrez	Tobias Haas	Jeff Hammond
Marc-André Hermanns	Thomas Hines	Mark Hoemmen
Dan Holmes	Atsushi Hori	Dominik Huber
Joshua Hursey	Dan Ibanez	Nusrat Islam
Julien Jaeger	Nysal Jan K A	Joachim Jenke
Vivek Kale	Michael Klemm	Michael Knobloch
Quincey Koziol	Donald Kruse	Ignacio Laguna
Pierre Lemarinier	Ben Lynam	W. Pepper Marts
Guillaume Mercier	Jacob Merson	

1	Pedram Mohammadalizadehbakhtevvari	Ali Mohammed	
2	Ali Omar Abdelazim Mohammed	Benson Muite	
3	Evelyn Namugwanya	Grace Nansamba	Thomas Naughton
4	Christoph Niethammer	William Okuno	Guillaume Papauré
5	Nick Papior	Sri Raj Paul	Howard Pritchard
6	Rolf Rabenseifner	Nick Radcliffe	Ken Raffetti
7	Naveen Ravichandrasekaran		Florian Reynier
8	Martin Ruefenacht	Amit Ruhela	Derek Schafer
9	Erik Schnetter	Whit Schonbein	Martin Schreiber
10	Thorsten Schütt	Joseph Schuchart	Martin Schulz
11	Amir Shehata	Nat Shineman	Riley Shipley
12	Anthony Skjellum	Brian Smith	AmirHossein Sojoodi
13	Jeff Squyres	Christopher Subich	Shinji Sumimoto
14	Hugo Taboada	Daniel Taylor	Yiltan Hassan Temucin
15	Keita Teranishi	Rajeev Thakur	Tim Thompson
16	Jesper Larsson Träff	Ehsan Toton	
17	Isaías Alberto Comprés Ureña		Josef Weidendorfer
18	Bill Williams	Andrew Worley	Hui Zhou

21 The MPI Forum also acknowledges and appreciates the valuable input from people via
22 e-mail and in person.

23 The following institutions supported the MPI-4.1 effort through time and travel support
24 for the people listed above.

- 25 Advanced Micro Devices, Inc.
- 26 Amazon.com, Inc
- 27 Argonne National Laboratory
- 28 Atos
- 29 CEA
- 30 Cisco Systems Inc.
- 31 Collis-Holmes Innovations Limited
- 32 Cornelis Networks
- 33 EPCC, The University of Edinburgh
- 34 Forschungszentrum Jülich
- 35 Fujitsu
- 36 HLRS, University of Stuttgart
- 37 Hewlett Packard Enterprise
- 38 Institut National de Recherche en Informatique et Automatique (Inria)
- 39 Intel Corporation
- 40 International Business Machines
- 41 Kichakato Kizito
- 42 Lawrence Berkeley National Laboratory
- 43 Lawrence Livermore National Laboratory
- 44 Leibniz Supercomputing Centre
- 45 Los Alamos National Laboratory
- 46 Meta Platforms Inc.
- 47 NVIDIA Corporation
- 48

Oak Ridge National Laboratory	1
ParaTools SAS	2
Queen’s University	3
RWTH Aachen University	4
Sandia National Laboratory	5
Technical University of Munich	6
Tennessee Technological University	7
Texas Advanced Computing Center	8
The Ohio State University	9
University of Alabama	10
University of Alabama at Birmingham	11
University of Basel	12
University of Illinois Urbana-Champaign	13
University of New Mexico	14
University of Tennessee, Chattanooga	15
University of Tennessee, Knoxville	16
University of Tokyo	17
Université Grenoble Alpes	18
VSC Research Center, TU Wien	19
ZIH, TU Dresden	20
	21
	22
	23
	24
	25
	26
	27
	28
	29
	30
	31
	32
	33
	34
	35
	36
	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48

DRAFT

DRAFT

Chapter 1

Introduction to MPI

1.1 Overview and Goals

MPI (Message-Passing Interface) is a *message-passing library interface specification*. All parts of this definition are significant. MPI addresses primarily the message-passing parallel programming model, in which data is moved from the address space of one process to that of another process through cooperative operations on each process. Extensions to the “classical” message-passing model are provided in collective operations, remote-memory access operations, dynamic process creation, and parallel I/O. MPI is a *specification*, not an implementation; there are multiple implementations of MPI. This specification is for a *library interface*; MPI is not a language, and all MPI operations are expressed as functions, subroutines, or methods, according to the appropriate language bindings that, for C and Fortran, are part of the MPI standard. The standard has been defined through an open process by a community of parallel computing vendors, computer scientists, and application developers. The next few sections provide an overview of the history of MPI’s development.

The main advantages of establishing a message-passing standard are portability and ease of use. In a distributed memory communication environment in which the higher level routines and/or abstractions are built upon lower level message-passing routines, the benefits of standardization are particularly apparent. Furthermore, the definition of a message-passing standard, such as that proposed here, provides vendors with a clearly defined base set of routines that they can implement efficiently, or in some cases for which they can provide hardware support, thereby enhancing scalability.

The goal of the Message-Passing Interface, simply stated, is to develop a widely used standard for writing message-passing programs. As such the interface should establish a practical, portable, efficient, and flexible standard for message passing.

A complete list of goals follows.

- Design an application programming interface (not necessarily for compilers or a system implementation library).
- Allow efficient communication: Avoid memory-to-memory copying, allow overlap of computation and communication, and offload to communication co-processors, where available.
- Allow for implementations that can be used in a heterogeneous environment.
- Allow convenient C and Fortran bindings for the interface.
- Assume a reliable communication interface: the user need not cope with communication failures. Such failures are dealt with by the underlying communication subsystem.

- 1 • Define an interface that can be implemented on many vendor’s platforms, with no
2 significant changes in the underlying communication and system software.
- 3
- 4 • Semantics of the interface should be language independent.
- 5
- 6 • The interface should be designed to allow for thread safety.
- 7

8 1.2 Background of MPI-1.0

9
10 MPI sought to make use of the most attractive features of a number of existing message-
11 passing systems, rather than selecting one of them and adopting it as the standard. Thus,
12 MPI was strongly influenced by work at the IBM T. J. Watson Research Center [3, 4], Intel’s
13 NX/2 [58], Express [15], nCUBE’s Vertex [54], p4 [10, 11], and PARMACS [7, 12]. Other
14 important contributions have come from Zipcode [61, 62], Chimp [21, 22], PVM [6, 19],
15 Chameleon [32], and PICL [27].

16 The MPI standardization effort involved about 60 people from 40 organizations mainly
17 from the United States and Europe. Most of the major vendors of concurrent computers
18 were involved in MPI, along with researchers from universities, government laboratories, and
19 industry. The standardization process began with the Workshop on Standards for Message-
20 Passing in a Distributed Memory Environment, sponsored by the Center for Research on
21 Parallel Computing, held April 29–30, 1992, in Williamsburg, Virginia [70]. At this work-
22 shop the basic features essential to a standard message-passing interface were discussed,
23 and a working group established to continue the standardization process.

24 A preliminary draft proposal, known as MPI-1, was put forward by Dongarra, Hempel,
25 Hey, and Walker in November 1992, and a revised version was completed in February
26 1993 [20]. MPI-1 embodied the main features that were identified at the Williamsburg
27 workshop as being necessary in a message passing standard. Since MPI-1 was primarily
28 intended to promote discussion and “get the ball rolling,” it focused mainly on point-to-point
29 communications. MPI-1 brought to the forefront a number of important standardization
30 issues, but did not include any collective communication routines and was not thread-safe.

31 In November 1992, a meeting of the MPI working group was held in Minneapolis, at
32 which it was decided to place the standardization process on a more formal footing, and to
33 generally adopt the procedures and organization of the High Performance Fortran Forum.
34 Subcommittees were formed for the major component areas of the standard, and an email
35 discussion service established for each. In addition, the goal of producing a draft MPI
36 standard by the Fall of 1993 was set. To achieve this goal the MPI working group met every
37 6 weeks for two days throughout the first 9 months of 1993, and presented the draft MPI
38 standard at the Supercomputing 93 conference in November 1993. These meetings and the
39 email discussion together constituted the MPI Forum, membership of which has been open
40 to all members of the high performance computing community.

41 42 1.3 Background of MPI-1.1, MPI-1.2, and MPI-2.0

43
44 Beginning in March 1995, the MPI Forum began meeting to consider corrections and exten-
45 sions to the original MPI standard document [24]. The first product of these deliberations
46 was Version 1.1 of the MPI specification, released in June of 1995 [25] (see
47 <http://www.mpi-forum.org> for official MPI document releases). At that time, effort focused
48 in five areas.

1. Further corrections and clarifications for the MPI-1.1 document.
2. Additions to MPI-1.1 that do not significantly change its types of functionality (new datatype constructors, language interoperability, etc.).
3. Completely new types of functionality (dynamic processes, one-sided communication, parallel I/O, etc.) that are what everyone thinks of as “MPI-2 functionality.”
4. Bindings for Fortran 90 and C++. MPI-2 specifies C++ bindings for both MPI-1 and MPI-2 functions, and extensions to the Fortran 77 binding of MPI-1 and MPI-2 to handle Fortran 90 issues.
5. Discussions of areas in which the MPI process and framework seem likely to be useful, but where more discussion and experience are needed before standardization (e.g., zero-copy semantics on shared-memory machines, real-time specifications).

Corrections and clarifications (items of type 1 in the above list) were collected in Chapter 3 of the MPI-2 document: “Version 1.2 of MPI.” That chapter also contains the function for identifying the version number. Additions to MPI-1.1 (items of types 2, 3, and 4 in the above list) are in the remaining chapters of the MPI-2 document, and constitute the specification for MPI-2. Items of type 5 in the above list have been moved to a separate document, the “MPI Journal of Development” (JOD), and are not part of the MPI-2 standard.

This structure makes it easy for users and implementors to understand what level of MPI compliance a given implementation has:

- MPI-1 compliance will mean compliance with MPI-1.3. This is a useful level of compliance. It means that the implementation conforms to the clarifications of MPI-1.1 function behavior given in Chapter 3 of the MPI-2 document. Some implementations may require changes to be MPI-1 compliant.
- MPI-2 compliance will mean compliance with all of MPI-2.1.
- The MPI Journal of Development is not part of the MPI standard.

It is to be emphasized that forward compatibility is preserved. That is, a valid MPI-1.1 program is both a valid MPI-1.3 program and a valid MPI-2.1 program, and a valid MPI-1.3 program is a valid MPI-2.1 program.

1.4 Background of MPI-1.3 and MPI-2.1

After the release of MPI-2.0, the MPI Forum kept working on errata and clarifications for both standard documents (MPI-1.1 and MPI-2.0). The short document “Errata for MPI-1.1” was released October 12, 1998. On July 5, 2001, a first ballot of errata and clarifications for MPI-2.0 was released, and a second ballot was voted on May 22, 2002. Both votes were done electronically. Both ballots were combined into one document: “Errata for MPI-2,” May 15, 2002. This errata process was then interrupted, but the Forum and its e-mail reflectors kept working on new requests for clarification.

Restarting regular work of the MPI Forum was initiated in three meetings, at EuroPVM/MPI’06 in Bonn, at EuroPVM/MPI’07 in Paris, and at SC’07 in Reno. In December 2007, a steering committee started the organization of new MPI Forum meetings at regular 8-weeks intervals. At the January 14–16, 2008 meeting in Chicago, the MPI Forum

1 decided to combine the existing and future MPI documents to one document for each ver-
2 sion of the MPI standard. For technical and historical reasons, this series was started with
3 MPI-1.3. Additional Ballots 3 and 4 solved old questions from the errata list started in 1995
4 up to new questions from the last years. After all documents (MPI-1.1, MPI-2, Errata for
5 MPI-1.1 (Oct. 12, 1998), and MPI-2.1 Ballots 1–4) were combined into one draft document,
6 for each chapter, a chapter author and review team were defined. They cleaned up the
7 document to achieve a consistent MPI-2.1 document. The final MPI-2.1 standard document
8 was finished in June 2008, and finally released with a second vote in September 2008 in the
9 meeting at Dublin, just before EuroPVM/MPI’08.

11 1.5 Background of MPI-2.2

13 MPI-2.2 is a minor update to the MPI-2.1 standard. This version addresses additional errors
14 and ambiguities that were not corrected in the MPI-2.1 standard as well as a small number
15 of extensions to MPI-2.1 that met the following criteria:

- 17 • Any correct MPI-2.1 program is a correct MPI-2.2 program.
- 18 • Any extension must have significant benefit for users.
- 19 • Any extension must not require significant implementation effort. To that end, all
20 such changes are accompanied by an open source implementation.

23 The discussions of MPI-2.2 proceeded concurrently with the MPI-3 discussions; in some
24 cases, extensions were proposed for MPI-2.2 but were later moved to MPI-3.

26 1.6 Background of MPI-3.0

28 MPI-3.0 is a major update to the MPI standard. The updates include the extension of
29 collective operations to include nonblocking versions, extensions to the one-sided operations,
30 and a new Fortran 2008 binding. In addition, the deprecated C++ bindings have been
31 removed, as well as many of the deprecated routines and MPI objects (such as the MPI_UB
32 datatype). Any valid MPI-2.2 program not using any of these removed MPI procedures or
33 objects is a valid MPI-3.0 program.

35 1.7 Background of MPI-3.1

37 MPI-3.1 is a minor update to the MPI standard. Most of the updates are corrections
38 and clarifications to the standard, especially for the Fortran bindings. New functions added
39 include routines to manipulate MPI_Aint values in a portable manner, nonblocking collective
40 I/O routines, and routines to get the index value by name for MPI_T performance and
41 control variables. A general index was also added. Any valid MPI-3.0 program is a valid
42 MPI-3.1 program.

44 1.8 Background of MPI-4.0

47 MPI-4.0 is a major update to the MPI standard. The largest changes are the addition of
48 large-count versions of many routines to address the limitations of using an int or INTEGER

for the count parameter, persistent collectives, partitioned communications, an alternative way to initialize MPI, application info assertions, and improvements to the definitions of error handling. In addition, there are a number of smaller improvements and corrections. Any valid MPI-3.1 program is a valid MPI-4.0 program with the exception of semantic changes listed in Chapter 18.

1.9 Background of MPI-4.1

MPI-4.1 is a minor update to the MPI standard. It contains mostly corrections and clarifications to the MPI-4.0 document. Several routines, the attribute key `MPI_HOST`, and the `mpif.h` Fortran include file are deprecated. A new routine provides a way to inquire about the hardware on which the MPI program is running. Any valid MPI-4.0 program is a valid MPI-4.1 program with the exception of semantic changes listed in Chapter 18.

1.10 Who Should Use This Standard?

This standard is intended for use by all those who want to write portable message-passing programs in Fortran and C (and access the C bindings from C++). This includes individual application programmers, developers of software designed to run on parallel machines, and creators of environments and tools. In order to be attractive to this wide audience, the standard must provide a simple, easy-to-use interface for the basic user while not semantically precluding the high-performance message-passing operations available on advanced machines.

1.11 What Platforms Are Targets for Implementation?

The attractiveness of the message-passing paradigm at least partially stems from its wide portability. Programs expressed this way may run on distributed-memory multiprocessors, networks of workstations, and combinations of all of these. In addition, shared-memory implementations, including those for multi-core processors and hybrid architectures, are possible. The paradigm will not be made obsolete by architectures combining the shared- and distributed-memory views, or by increases in network speeds. It thus should be both possible and useful to implement this standard on a great variety of machines, including those “machines” consisting of collections of other machines, parallel or not, connected by a communication network.

The interface is suitable for use by fully general MIMD (Multiple Instruction, Multiple Data) programs, as well as those written in the more restricted style of SPMD (Single Program, Multiple Data). MPI provides many features intended to improve performance on scalable parallel computers with specialized interprocessor communication hardware. Thus, we expect that native, high-performance implementations of MPI will be provided on such machines. At the same time, implementations of MPI on top of standard Unix interprocessor communication protocols will provide portability to workstation clusters and heterogenous networks of workstations.

1.12 What Is Included in the Standard?

The standard includes:

- Point-to-point communication,
- Partitioned communication,
- Datatypes,
- Collective operations,
- Process groups,
- Communication contexts,
- Virtual Topologies for MPI Processes,
- Environmental management and inquiry,
- The Info object,
- Process initialization, creation, and management,
- One-sided communication,
- External interfaces,
- Parallel file I/O,
- Tool support,
- Language bindings for Fortran and C, and
- Additional topics in side-documents.

1.13 Side-documents

Side-documents extend and/or modify features, semantics, language bindings, and other aspects covered in this document. Side-documents shall not modify any aspects defined in the MPI Standard without providing a mechanism that explicitly enables these deviations. Execution of a program that does not explicitly enable deviations from the MPI Standard will comply with the MPI Standard, even when using an MPI implementation that implements a side-document that modifies any aspects.

Each side-document is versioned with a scheme that is independent from the MPI Standard version and from other side-documents. All side-documents specify compatibility and interoperability with versions of the MPI Standard and may define interoperability with features and semantics from other side-documents. Side-documents are not required to provide full coverage of all MPI concepts, but shall document which MPI concepts are affected. A compliant implementation is not required to comply with any side-documents. However, if compliance with a particular version of a side-document is claimed, the implementation must comply with the entire side-document. Side-documents will be found at the same location as the MPI Standard[1].

1.14 Organization of This Document

The following is a list of the remaining chapters in this document, along with a brief description of each.

- Chapter 2, [MPI Terms and Conventions](#), explains notational terms and conventions used throughout the MPI document.
- Chapter 3, [Point-to-Point Communication](#), defines the basic, pairwise communication subset of MPI. *Send* and *receive* are found here, along with many associated functions designed to make basic communication powerful and efficient.
- Chapter 4, [Partitioned Point-to-Point Communication](#), defines a method of performing partitioned communication in MPI. Partitioned communication allows multiple contributions of data to be made, potentially, from multiple actors (e.g., threads or tasks) in an MPI process to a single message.
- Chapter 5, [Datatypes](#), defines a method to describe any data layout, e.g., an array of structures.
- Chapter 6, [Collective Communication](#), defines process-group collective communication operations. Well known examples of this are barrier and broadcast over a group of processes (not necessarily all the processes).
- Chapter 7, [Groups, Contexts, Communicators, and Caching](#), shows how groups of processes are formed and manipulated, how unique communication contexts are obtained, and how the two are bound together into a *communicator*.
- Chapter 8, [Virtual Topologies for MPI Processes](#), explains a set of utility functions meant to assist in the mapping of MPI process groups (a linearly ordered set) to richer topological structures such as multi-dimensional grids.
- Chapter 9, [MPI Environmental Management](#), explains how the programmer can manage and make inquiries of the current MPI environment. These functions are needed for the writing of correct, robust programs, and are especially important for the construction of highly-portable message-passing programs.
- Chapter 10, [The Info Object](#), defines an opaque object that is used as input in several MPI routines.
- Chapter 11, [Process Initialization, Creation, and Management](#), defines several approaches to MPI initialization, process creation, and process management while placing minimal restrictions on the execution environment.
- Chapter 12, [One-Sided Communications](#), defines communication routines that can be completed by a single process. These include shared-memory operations (*put/get*) and remote accumulate operations.
- Chapter 13, [External Interfaces](#), defines routines designed to allow developers to layer on top of MPI.
- Chapter 14, [I/O](#), defines MPI support for parallel I/O.

- 1 • Chapter 15, [Tool Support](#), covers interfaces that allow debuggers, performance ana-
2 lyzers, and other tools to obtain data about the operation of MPI processes.
- 3
- 4 • Chapter 16, [Deprecated Interfaces](#), describes routines that are kept for reference.
5 However usage of these functions is discouraged, as they may be deleted in future
6 versions of the standard.
- 7
- 8 • Chapter 17, [Removed Interfaces](#), describes routines and constructs that have been
9 removed from MPI.
- 10
- 11 • Chapter 18, [Semantic Changes and Warnings](#), describes semantic changes from pre-
12 vious versions of MPI.
- 13
- 14 • Chapter 19, [Language Bindings](#), discusses Fortran issues, and describes language in-
15 teroperability aspects between C and Fortran.

16 The Appendices are:

- 17 • Annex A, [Language Bindings Summary](#), gives specific syntax in C and Fortran, for
18 all MPI functions, constants, and types.
- 19
- 20 • Annex B, [Change-Log](#), summarizes some changes since the previous version of the
21 standard.
- 22
- 23 • Several Index pages show the locations of [general terms and definitions](#), [examples](#), [con-](#)
24 [stants and predefined handles](#), [declarations of C and Fortran types](#), [callback routine](#)
25 [prototypes](#), and all [MPI functions](#).

26
27 MPI provides various interfaces to facilitate interoperability of distinct MPI imple-
28 mentations. Among these are the canonical data representation for MPI I/O and for
29 MPI_PACK_EXTERNAL and MPI_UNPACK_EXTERNAL. The definition of an actual binding
30 of these interfaces that will enable interoperability is outside the scope of this document.

31 A separate document consists of ideas that were discussed in the MPI Forum during the
32 MPI-2 development and deemed to have value, but were not included in the MPI standard.
33 They are part of the “Journal of Development” (JOD), which was created to capture these
34 ideas and discussions. The JOD is available at <https://www.mpi-forum.org/docs>.

Chapter 2

MPI Terms and Conventions

This chapter explains notational terms and conventions used throughout the MPI document, some of the choices that have been made, and the rationale behind those choices.

2.1 Document Notation

Rationale. Throughout this document, the rationale for the design choices made in the interface specification is set off in this format. Some readers may wish to skip these sections, while readers interested in interface design may want to read them carefully. (*End of rationale.*)

Advice to users. Throughout this document, material aimed at users and that illustrates usage is set off in this format. Some readers may wish to skip these sections, while readers interested in programming in MPI may want to read them carefully. (*End of advice to users.*)

Advice to implementors. Throughout this document, material that is primarily commentary to implementors is set off in this format. Some readers may wish to skip these sections, while readers interested in MPI implementations may want to read them carefully. (*End of advice to implementors.*)

2.2 Naming Conventions

In many cases MPI names for C functions are of the form `MPI_Class_action_subset`. This convention originated with MPI-1. Since MPI-2 an attempt has been made to standardize the names of MPI functions according to the following rules.

1. In C and the Fortran `mpi_f08` module, all routines associated with a particular type of MPI object should be of the form `MPI_Class_action_subset` or, if no subset exists, of the form `MPI_Class_action`. In the Fortran `mpi` module and (deprecated) `mpif.h` file, all routines associated with a particular type of MPI object should be of the form `MPI_CLASS_ACTION_SUBSET` or, if no subset exists, of the form `MPI_CLASS_ACTION`.
2. If the routine is not associated with a class, the name should be of the form `MPI_Action_subset` or `MPI_ACTION_SUBSET` in C and Fortran.
3. The names of certain actions have been standardized. In particular, **create** creates a new object, **get** retrieves information about an object, **set** sets this information, **delete** deletes information, **is** asks whether or not an object has a certain property.

C and Fortran names for some MPI functions (that were defined during the MPI-1 process) violate these rules in several cases. The most common exceptions are the omission of the **Class** name from the routine and the omission of the **Action** where one can be inferred.

2.3 Procedure Specification

MPI procedures are specified using a language-independent notation. The arguments of procedure calls are marked as IN, OUT, or INOUT. The meanings of these are:

IN: the call may use the input value but does not update the argument from the perspective of the caller at any time during the call’s execution,

OUT: the call may update the argument but does not use its input value,

INOUT: the call may both use and update the argument.

There is one special case—if an argument is a handle to an opaque object (these terms are defined in Section 2.5.1), and the object is updated by the procedure call, then the argument is marked INOUT or OUT. It is marked this way even though the handle itself is not modified—we use the INOUT or OUT attribute to denote that what the handle *references* is updated.

Rationale. The definition of MPI tries to avoid, to the largest possible extent, the use of INOUT arguments, because such use is error-prone, especially for scalar arguments. (*End of rationale.*)

MPI’s use of IN, OUT, and INOUT is intended to indicate to the user how an argument is to be used, but does not provide a rigorous classification that can be translated directly into all language bindings (e.g., INTENT in Fortran 90 bindings or const in C bindings). For instance, the “constant” MPI_BOTTOM can usually be passed to OUT buffer arguments. Similarly, MPI_STATUS_IGNORE can be passed as the OUT status argument.

A common occurrence for MPI functions is an argument that is used as IN by some processes and OUT by other processes. Such an argument is, syntactically, an INOUT argument and is marked as such, although, semantically, it is not used in one call both for input and for output on a single process.

Another frequent situation arises when an argument value is needed only by a subset of the processes. When an argument is not significant at a process then an arbitrary value can be passed as an argument.

Unless specified otherwise, an argument of type OUT or type INOUT cannot be aliased with any other argument passed to an MPI procedure. An example of argument aliasing in C appears below. If we define a C procedure like this,

```
void copyIntBuffer(int *pin, int *pout, int len)
{
    int i;
    for (i=0; i<len; ++i) *pout++ = *pin++;
}
```

then a call to it in the following code fragment has aliased arguments.

```
int a[10];
copyIntBuffer(a, a+3, 7);
```

Although the C language allows this, such usage of MPI procedures is forbidden unless otherwise specified. Note that Fortran prohibits aliasing of arguments.

All MPI functions are first specified in the language-independent notation. Immediately below this, language dependent bindings follow:

- The ISO C version(s) of the function.
- The Fortran version(s) used with `USE mpi_f08`.
- The Fortran version of the same function used with `USE mpi` or (deprecated) `INCLUDE 'mpif.h'`.

Some MPI procedures have two interfaces for a given language support; see Sections 2.5.6 and 2.5.8.

An exception is Section 15.3 “The MPI Tool Information Interface”, which only provides ISO C interfaces.

“Fortran” in this document refers to Fortran 90 or later; see Section 2.6.

The words function, routine, procedure, procedure call, and call are often used as synonyms within this standard.

2.4 Semantic Terms

When discussing MPI procedures the following semantic terms are used. The term **message data buffer** refers to the send/receive buffer used in a communication procedure. The term **file data buffer** refers to the data buffers used by MPI I/O procedures. In this section we use the term **data buffer** and depending on the MPI procedure it will refer to message data buffer or file data buffer. Annex A.2 shows how the terms defined in this section apply to all operation-related MPI procedures.

2.4.1 MPI Operations

MPI operation: An MPI operation is a sequence of steps performed by the MPI library to establish and enable data transfer and/or synchronization. It consists of four stages: initialization, starting, completion, and freeing, and it is implemented as a set of one or more MPI procedures, see Section 2.4.2.

Initialization hands over the argument list to the operation but not the content of the data buffers, if any. The specification of an operation may state that array arguments must not be changed until the operation is freed.

Starting hands over control of the data buffers, if any, to the associated operation.

Note that **initiation** refers to the combination of the initialization and starting stages.

Completion returns control of the content of the data buffers and indicates that output buffers and arguments, if any, have been updated.

Note that an MPI operation is **complete** when the MPI procedure implementing the completion stage returns.

Freeing returns control of the rest of the argument list (e.g., the data buffer address and array arguments).



Figure 2.1: State transition diagram for blocking operations




Figure 2.2: State transition diagram for nonblocking operations




Figure 2.3: State transition diagram for persistent operations

MPI operations are available in one or more of these forms: blocking, nonblocking, and persistent.

Blocking operation: For a **blocking operation**, all four stages are combined in a single procedure call (as shown in Figure 2.1 and defined in Section 2.4.2).

Nonblocking operation: For a **nonblocking operation**, the initialization and starting stages are combined into a single nonblocking procedure call and the completion and freeing stages are combined into a separate, single procedure call, which can be blocking or nonblocking (as shown in Figure 2.2 and defined in Section 2.4.2).

Persistent operation: For a **persistent operation**, there is a separate procedure for each of the four stages (as shown in Figure 2.3 and defined in Section 2.4.2). Each of these procedures may be blocking or nonblocking.

For a partitioned send operation, an additional call to activate each partition of the send buffer (see Section 4.2.1) is required to finish the starting stage. For a partitioned receive operation, before the operation is complete the user is allowed to access a partition of the output buffer after verifying that it has arrived (see Section 4.2.2).

These four stages lead to the **operation states initialized, started, complete, and freed**. A *started operation* is also named **active**, and the states *initialized* and *complete* are also named **inactive**.

Active communication and I/O operations are also named **pending** operations. Note that a *pending* operation can be a nonblocking or persistent operation that is started and not yet complete (even if the request handle has been freed), or a blocking operation that is not yet complete, such as a receive operation that is waiting for a message to be received.

Additionally, an MPI operation can be collective or noncollective.

Collective operation: A set of related operations, one per MPI process in a group or groups of MPI processes. For collective operations the completion stage may or may not finish before all processes in the group have started the operation.

Collective MPI operations are also available as blocking, nonblocking, or persistent operations.

Noncollective operation: Noncollective operations are defined as operations that are not collective.

Many MPI operations coordinate activities at multiple MPI processes: the semantics of such an operation require one or more other specific semantically-related operations to be *started* before it is guaranteed that the operation can transition to the *complete* operation state. For example, a receive operation requires a related send operation to be started before the receive can complete; or a collective operation might not complete before such operations are also started in all MPI processes of the respective group.

Enabled: An MPI operation is **enabled** at a particular MPI process when all specific semantically-related operations required to guarantee completion at that MPI process have been started.

Rationale. MPI implementations may include optimizations (for example, automatic buffering) that allow an MPI operation to complete before it is enabled. (*End of rationale.*)

Some MPI operations are **a priori enabled**, i.e., they do not require any other specific semantically-related operation for completion. For example, a buffered send operation completes independently of the related receive operation.

Once an MPI operation is *enabled*, the operation must eventually complete. An operation may already be enabled before it is started. For example, a receive operation is already enabled if it is started after the matching send operation was started.

Rationale. The definition of an operation A being *enabled* is asymmetric: *enabled* includes that all specific semantically-related operations A'_i required to guarantee completion have been started, but does not include that the operation A itself is already started.

Examples:

- A receive is enabled exactly when the related send is started.
- A standard mode send operation is enabled exactly when the related receive is started. If an MPI implementation chooses to use internal buffering, the send operation may be already completed before it is enabled, i.e., the receive is started.
- A synchronous mode send operation is enabled exactly when the related receive is started and must not complete before it is enabled.
- A buffered mode send operation is a priori enabled.

- 1 • A ready mode send can be started only when it is already enabled, i.e., the
- 2 related receive is started.
- 3 • For a collective broadcast, the operation at a particular MPI process is enabled
- 4 exactly when all other MPI processes in the group have started their related
- 5 broadcast operation.

6 Specifically, for the set of related operations on a group of MPI processes that consti-
 7 tute a collective operation that may synchronize, the operation on a particular MPI
 8 process p is enabled when all other MPI processes $p_i \neq p$ in the group have started
 9 their related operation, while the operation on p need not have started yet. (*End of*
 10 *rationale.*)

11 2.4.2 MPI Procedures

12 All MPI procedures can either be *local* or *nonlocal*—defined as follows:

13 **Nonlocal procedure:** An MPI procedure is **nonlocal** if returning may require, during its
 14 execution, some specific semantically-related MPI procedure to be called on another
 15 MPI process.

16 **Local procedure:** An MPI procedure is **local** if it is not *nonlocal*.

17 An MPI operation is implemented as a set of one or more MPI procedures. An MPI
 18 **operation-related procedure** implements at least a part of a stage of an MPI operation
 19 as described in Section 2.4.1. An MPI operation-related procedure may also implement
 20 one or more stages of one or several MPI operations. In certain cases, more than one MPI
 21 operation-related procedure may be needed to implement a single stage.

22 There are also other MPI procedures that do not implement any stage of any MPI
 23 operation.

24 The semantics of MPI operation-related procedures are described using two orthogonal
 25 (independent) concepts: completeness (depends on which stages are included) and locality.
 26 Such procedures can be either incomplete, or completing, or freeing, or completing and
 27 freeing based on the status of the associated operation at the time the procedure returns.
 28 Also, all such procedures can be described as either blocking or nonblocking, but these latter
 29 two terms refer to combinations of the completeness and locality concepts. Additionally,
 30 all MPI operation-related procedures can be collective or noncollective.

31 The following are properties of MPI operation-related procedures:

32 **Initialization procedure:** An MPI procedure is an **initialization procedure** if return
 33 from the procedure indicates that the associated operation has completed its initial-
 34 ization stage, which implies that the user has handed over control of the argument list
 35 (but not contents of the data buffers) to MPI. The user is still allowed to read or mod-
 36 ify the contents of the data buffers. If an initializing procedure is not also the freeing
 37 procedure of the associated operation (see below) then the user is not permitted to
 38 deallocate the data buffers or to modify the array arguments.

39 **Starting procedure:** An MPI procedure is a **starting procedure** if return from the
 40 procedure indicates that the associated operation has completed its starting stage,
 41 which implies that the user has handed over control of the data buffers to MPI. If a
 42 starting procedure is not also a completing procedure of the associated operation (see
 43

below) then the user is not permitted to modify input data buffers or to read output data buffers.

Initiation procedure: An MPI procedure is an **initiation procedure** if return from the procedure indicates that both the initialization and the starting stage have completed, which implies control of the entire argument list is handed over to MPI.

Completing procedure: An MPI procedure is called **completing** if return from the procedure indicates that at least one associated operation has finished its completion stage, which implies that the user can rely on the content of the output data buffers and modify the content of input and output data buffers of such operation(s). If a completing procedure is not also a freeing procedure (see below) then the user is not permitted to deallocate the data buffers or to modify the array arguments.

Incomplete procedure: An MPI procedure is called **incomplete** if it is not a completing procedure.

Freeing procedure: An MPI procedure is **freeing** if return from the procedure indicates that at least one associated operation has finished its freeing stage, which implies that the user can reuse all parameters specified when initializing such associated operation(s).

Nonblocking procedure: An MPI procedure is **nonblocking** if it is incomplete and local.

Blocking procedure: An MPI procedure is **blocking** if it is not nonblocking.

Advice to users. Note that for operation-related MPI procedures, in most cases incomplete procedures are local and completing procedures are nonlocal. Exceptions are noted where such procedures are defined. In many cases an additional prefix letter l as an abbreviation of the words **incomplete** and **immediate** marks nonblocking procedures in the procedure name.

Some categorization examples are listed below.

Nonblocking procedures:

- incomplete and local: MPI_ISEND, MPI_IRecv, MPI_IBCAST, MPI_IMPROBE, MPI_SEND_INIT, MPI_RECV_INIT, ...

Blocking procedures:

- completing and nonlocal: MPI_SEND, MPI_RECV, MPI_BCAST, ...
- incomplete and nonlocal: MPI_MPROBE, MPI_BCAST_INIT, ..., MPI_FILE_{READ|WRITE}_{AT_ALL|ALL|ORDERED}_BEGIN.
- completing and local: MPI_BSEND, MPI_RSEND, MPI_MRECV.

MPI procedures that are not MPI operation-related:

- MPI_COMM_RANK, MPI_WTIME, MPI_PROBE, MPI_Iprobe, ...

(End of advice to users.)

Collective procedure: An MPI procedure is **collective** if all processes in a group or groups of MPI processes need to invoke the procedure.

1 Initialization procedures of collective operations over the same process group must be
2 executed in the same order by all members of the process group.

3 An MPI collective procedure is **synchronizing** if it will only return once all pro-
4 cesses in the associated group or groups of MPI processes have called the appropriate
5 matching MPI procedure.

6 The initiation procedures for nonblocking collective operations and the starting pro-
7 cedures for persistent collective operations are local and shall not be synchronizing.

8 All other procedures for collective operations, such as for blocking collective operations
9 and the initialization procedures for persistent collective operations, may or may not
10 be synchronizing.

11
12 *Advice to users.* Calling any synchronizing function is erroneous when there
13 is no possibility of corresponding calls at all other processes in the associated
14 process group.

15
16 Waiting for completion of any collective operation is erroneous when there is no
17 possibility that all other processes in the associated group will be able to start
18 the corresponding operation. (*End of advice to users.*)

19 **Noncollective procedure:** Noncollective procedures are defined as procedures that are
20 not collective.

21
22 The definition of **local** and **nonlocal** MPI procedures can also be applied to a specific
23 procedure invocation or to procedure calls **under certain constraints**. For example, a call
24 to a completing receive procedure that happens after the related send operation was already
25 started may be described as local, even though the completing receive procedure without
26 the constraint is nonlocal. More generally, a call to any completing procedure that happens
27 after the operation was already *enabled* is local, even if the completing procedure without
28 the constraint is nonlocal. Another example, a call to a blocking collective procedure using
29 a process group of size one is local, even if the blocking collective procedure without the
30 constraint is nonlocal.

31 32 2.4.3 MPI Datatypes

33 For datatypes, the following terms are defined:

34
35 **predefined:** A predefined datatype is a datatype with a predefined (constant) name (such
36 as MPI_INT, MPI_FLOAT_INT, or MPI_PACKED) or a datatype constructed with
37 MPI_TYPE_CREATE_F90_INTEGER, MPI_TYPE_CREATE_F90_REAL, or
38 MPI_TYPE_CREATE_F90_COMPLEX. The former are **named** whereas the latter are
39 **unnamed**.

40
41 **derived:** A derived datatype is any datatype that is not predefined.

42
43 **portable:** A datatype is portable if it is a predefined datatype, or it is derived from
44 a portable datatype using only the type constructors MPI_TYPE_CONTIGUOUS,
45 MPI_TYPE_VECTOR, MPI_TYPE_INDEXED,
46 MPI_TYPE_CREATE_INDEXED_BLOCK, MPI_TYPE_CREATE_SUBARRAY,
47 MPI_TYPE_DUP, and MPI_TYPE_CREATE_DARRAY. Such a datatype is portable
48 because all displacements in the datatype are in terms of extents of one predefined

datatype. Therefore, if such a datatype fits a data layout in one memory, it will fit the corresponding data layout in another memory, if the same declarations were used, even if the two systems have different architectures. On the other hand, if a datatype was constructed using `MPI_TYPE_CREATE_HINDEXED`, `MPI_TYPE_CREATE_HINDEXED_BLOCK`, `MPI_TYPE_CREATE_HVECTOR` or `MPI_TYPE_CREATE_STRUCT`, then the datatype contains explicit byte displacements (e.g., providing padding to meet alignment restrictions). These displacements are unlikely to be chosen correctly if they fit data layout on one memory, but are used for data layouts on another process, running on a processor with a different architecture.

equivalent: Two datatypes are equivalent if they appear to have been created with the same sequence of calls (and arguments) and thus have the same typemap. Two equivalent datatypes do not necessarily have the same cached attributes or the same names.

2.5 Datatypes

2.5.1 Opaque Objects

MPI manages **system memory** that is used for buffering messages and for storing internal representations of various MPI objects such as groups, communicators, datatypes, etc. This memory is not directly accessible to the user, and objects stored there are **opaque**: their size and shape is not visible to the user. Opaque objects are accessed via **handles**, which exist in user space. MPI procedures that operate on opaque objects are passed handle arguments to access these objects. In addition to their use by MPI calls for object access, handles can participate in assignments and comparisons.

In Fortran with `USE mpi` or (deprecated) `INCLUDE 'mpif.h'`, all handles have type `INTEGER`. In Fortran with `USE mpi_f08`, and in C, a different handle type is defined for each category of objects. With Fortran `USE mpi_f08`, the handles are defined as Fortran `BIND(C)` derived types that consist of only one element `INTEGER :: MPI_VAL`. The internal handle value is identical to the Fortran `INTEGER` value used in the `mpi` module and (deprecated) `mpif.h`. The operators `.EQ.`, `.NE.`, `=` and `/=` are overloaded to allow the comparison of these handles. The type names are identical to the names in C, except that they are not case sensitive. For example:

```
TYPE, BIND(C) :: MPI_Comm
  INTEGER      :: MPI_VAL
END TYPE MPI_Comm
```

The C types must support the use of the assignment and equality operators.

Advice to implementors. In Fortran, the handle can be an index into a table of opaque objects in a system table; in C it can be such an index or a pointer to the object. (*End of advice to implementors.*)

Rationale. Since the Fortran integer values are equivalent, applications can easily convert MPI handles between all three supported Fortran methods. For example, an integer communicator handle `COMM` can be converted directly into an exactly equivalent `mpi_f08` communicator handle named `comm_f08` by `comm_f08%MPI_VAL=COMM`, and vice versa. The use of the `INTEGER` defined handles and the `BIND(C)` derived type handles

1 is different: Fortran 2003 (and later) define that BIND(C) derived types can be used
2 within user defined common blocks, but it is up to the rules of the companion C
3 compiler how many numerical storage units are used for these BIND(C) derived type
4 handles. Most compilers use one unit for both, the INTEGER handles and the handles
5 defined as BIND(C) derived types. (*End of rationale.*)

6
7 *Advice to users.* If a user wants to substitute the `mpi` module or the (deprecated)
8 `mpif.h` by the `mpi_f08` module and the application program stores a handle in a
9 Fortran common block then it is necessary to change the Fortran support method in
10 all application routines that use this common block, because the number of numerical
11 storage units of such a handle can be different in the two modules. (*End of advice to*
12 *users.*)

13
14 Opaque objects are allocated and deallocated by calls that are specific to each object
15 type. These are listed in the sections where the objects are described. The calls accept
16 a handle argument of matching type. In an allocate call this is an OUT argument that
17 returns a valid reference to the object. In a call to deallocate this is an INOUT argument
18 that returns with an “invalid handle” value. MPI provides an “invalid handle” constant for
19 each object type. Comparisons to this constant are used to test for validity of the handle.

20 A call to a deallocate routine invalidates the handle and marks the object for deal-
21 location. The object is not accessible to the user after the call. However, MPI need not
22 deallocate the object immediately. Any operation *pending* (at the time of the deallocate)
23 and *decoupled MPI activity* (see Section 2.9) that involves this object will complete normally;
24 the object will be deallocated afterwards.

25 An opaque object and its handle are significant only at the process where the object
26 was created and cannot be transferred to another process.

27 MPI provides certain predefined opaque objects and predefined, static handles to these
28 objects. The user must not free such objects.

29
30 *Rationale.* This design hides the internal representation used for MPI data structures,
31 thus allowing similar calls in C and Fortran. It also avoids conflicts with the typing
32 rules in these languages, and easily allows future extensions of functionality. The
33 mechanism for opaque objects used here loosely follows the POSIX Fortran binding
34 standard.

35 The explicit separation of handles in user space and objects in system space allows
36 space-reclaiming and deallocation calls to be made at appropriate points in the user
37 program. If the opaque objects were in user space, one would have to be very careful
38 not to go out of scope before any pending operation requiring that object completed.
39 The specified design allows an object to be marked for deallocation, the user program
40 can then go out of scope, and the object itself still persists until any pending operations
41 are complete.

42 The requirement that handles support assignment/comparison is made since such op-
43 erations are common. This restricts the domain of possible implementations. The
44 alternative in C would have been to allow handles to have been an arbitrary, opaque
45 type. This would force the introduction of routines to do assignment and compar-
46 ison, adding complexity, and was therefore ruled out. In Fortran, the handles are
47 defined such that assignment and comparison are available through the operators of
48 the language or overloaded versions of these operators. (*End of rationale.*)

Advice to users. A user may accidentally create a dangling reference by assigning to a handle the value of another handle, and then deallocating the object associated with these handles. Conversely, if a handle variable is deallocated before the associated object is freed, then the object becomes inaccessible (this may occur, for example, if the handle is a local variable within a subroutine, and the subroutine is exited before the associated object is deallocated). It is the user's responsibility to avoid adding or deleting references to opaque objects, except as a result of MPI calls that allocate or deallocate such objects. (*End of advice to users.*)

Advice to implementors. The intended semantics of opaque objects is that opaque objects are separate from one another; each call to allocate such an object copies all the information required for the object. Implementations may avoid excessive copying by substituting referencing for copying. For example, a derived datatype may contain references to its components, rather than copies of its components; a call to `MPI_COMM_GROUP` may return a reference to the group associated with the communicator, rather than a copy of this group. In such cases, the implementation must maintain reference counts, and allocate and deallocate objects in such a way that the visible effect is as if the objects were copied. (*End of advice to implementors.*)

2.5.2 Array Arguments

An MPI call may need an argument that is an array of opaque objects, or an array of handles. The array-of-handles is a regular array with entries that are handles to objects of the same type in consecutive locations in the array. Whenever such an array is used, an additional `len` argument is required to indicate the number of valid entries (unless this number can be derived otherwise). The valid entries are at the beginning of the array; `len` indicates how many of them there are, and need not be the size of the entire array. The same approach is followed for other array arguments. In some cases NULL handles are considered valid entries. When a NULL argument is desired for an array of statuses, one uses `MPI_STATUSES_IGNORE`.

2.5.3 State

MPI procedures use at various places arguments with *state* types. The values of such a datatype are all identified by names, and no operation is defined on them. For example, the `MPI_TYPE_CREATE_SUBARRAY` routine has a state argument `order` with values `MPI_ORDER_C` and `MPI_ORDER_FORTRAN`.

2.5.4 Named Constants

MPI procedures sometimes assign a special meaning to a special value of a basic type argument; e.g., `tag` is an integer-valued argument of point-to-point communication operations, with a special wild-card value, `MPI_ANY_TAG`. Such arguments will have a range of regular values, which is a proper subrange of the range of values of the corresponding basic type; special values (such as `MPI_ANY_TAG`) will be outside the regular range. The range of regular values, such as `tag`, can be queried using environmental inquiry functions, see Chapter 9. The range of other values, such as `source`, depends on values given by other MPI routines (in the case of `source` it is the communicator size).

MPI also provides predefined named constant handles, such as `MPI_COMM_WORLD`.

1 All named MPI constants, with the exceptions noted below for Fortran, can be used in
 2 initialization expressions or assignments. Opaque objects accessed by constant handles are
 3 defined and do not change value between MPI initialization (e.g., with MPI_INIT) and MPI
 4 finalization (e.g., with MPI_FINALIZE). The handles themselves are constants and can be
 5 also used in initialization expressions or assignments.

6 In C, all named MPI constants that are described as “integer constant expression” in
 7 Section A.1.1 must be implemented as *C integer constant expressions* of the specified integer
 8 type. All other MPI constants in C are not required to be *C integer constant expressions* but
 9 must be usable in initialization expressions and assignments. Thus, they are not guaranteed
 10 to be usable in array declarations or as case-labels in `switch` statements.

11 In Fortran, all named MPI constants (with the exceptions below) must be declared with
 12 the `PARAMETER` attribute.

13 The constants that cannot be used in initialization expressions or assignments in Fortran
 14 are as follows:

```
15
16 MPI_BOTTOM
17 MPI_BUFFER_AUTOMATIC
18 MPI_STATUS_IGNORE
19 MPI_STATUSES_IGNORE
20 MPI_ERRCODES_IGNORE
21 MPI_IN_PLACE
22 MPI_ARGV_NULL
23 MPI_ARGVS_NULL
24 MPI_UNWEIGHTED
25 MPI_WEIGHTS_EMPTY
```

26 *Advice to implementors.* In Fortran the implementation of these special constants
 27 may require the use of language constructs that are outside the Fortran standard. Us-
 28 ing special values for the constants (e.g., by defining them through `PARAMETER` state-
 29 ments) is not possible because an implementation cannot distinguish these values from
 30 valid data. Typically, these constants are implemented as predefined static variables
 31 (e.g., a variable in an MPI-declared `COMMON` block), relying on the fact that the target
 32 compiler passes data by address. Inside the subroutine, this address can be extracted
 33 by some mechanism outside the Fortran standard (e.g., by Fortran extensions or by
 34 implementing the function in C). (*End of advice to implementors.*)

36 2.5.5 Choice

37 MPI functions sometimes use arguments with a *choice* (or union) data type. Distinct calls to
 38 the same routine may pass by reference actual arguments of different types. The mechanism
 39 for providing such arguments will differ from language to language. For Fortran with the
 40 (deprecated) include file `mpif.h` or the `mpi` module, the document uses `<type>` to represent
 41 a choice variable; with the Fortran `mpi_f08` module, such arguments are declared with the
 42 Fortran 2018 syntax `TYPE(*)`, `DIMENSION(..)`; for C, we use `void*`.

44 *Advice to implementors.* Implementors can freely choose how to implement choice
 45 arguments in the `mpi` module, e.g., with a nonstandard compiler-dependent method
 46 that has the quality of the call mechanism in the implicit Fortran interfaces, or with
 47 the method defined for the `mpi_f08` module. See details in Section 19.1.1. (*End of*
 48 *advice to implementors.*)

2.5.6 Absolute Addresses and Relative Address Displacements

Some MPI procedures use *address* arguments that represent an *absolute address* in the calling program, or *relative displacement* arguments that represent differences of two absolute addresses. The datatype of such arguments is MPI_Aint in C and INTEGER(KIND=MPI_ADDRESS_KIND) in Fortran. These types must have the same width and encode address values in the same manner such that address values in one language may be passed directly to another language without conversion. There is the MPI constant MPI_BOTTOM to indicate the start of the address range. For retrieving absolute addresses or any calculation with absolute addresses, one should use the routines and functions provided in Section 5.1.5. Section 5.1.12 provides additional rules for the correct use of absolute addresses. For expressions with relative displacements or other usage without absolute addresses, intrinsic operators (e.g., +, -, *) can be used.

Rationale. Byte displacement values need to be large enough to encode any value used for expressing absolute or relative memory addresses. Prior to MPI-4.0, some MPI routines used int in C and INTEGER in Fortran as the type for *byte displacement* arguments. To avoid breaking backward compatibility, this version of the standard continues to support int in C as well as INTEGER in Fortran in such routines. In addition, this version of the standard supports using MPI_Aint in C (via separate “_c” suffixed procedures) as well as INTEGER(KIND=MPI_ADDRESS_KIND) in Fortran (via polymorphic interfaces in newer MPI Fortran bindings (USE mpi_f08)) in such routines. See Section 19.2 for a full explanation. (*End of rationale.*)

2.5.7 File Offsets

For I/O there is a need to give the size, displacement, and offset into a file. These quantities can easily be larger than 32 bits, which can be the default size of a Fortran integer. To overcome this, these quantities are declared to be INTEGER(KIND=MPI_OFFSET_KIND) in Fortran. In C one uses MPI_Offset. These types must have the same width and encode address values in the same manner such that offset values in one language may be passed directly to another language without conversion.

2.5.8 Counts

As described above, MPI defines types (e.g., MPI_Aint) to address locations within memory and other types (e.g., MPI_Offset) to address locations within files. In addition, some MPI procedures use *count* arguments that represent a number of MPI datatypes on which to operate. Furthermore, *timestamps* in the context of the MPI Tool Information Interface are a count of clock ticks elapsed since some time in the past. At times, one needs a single type that can be used to address locations within either memory or files as well as express *count* values, and that type is MPI_Count in C and INTEGER(KIND=MPI_COUNT_KIND) in Fortran. These types must have the same width and encode values in the same manner such that count values in one language may be passed directly to another language without conversion. The size of the MPI_Count type is determined by the MPI implementation with the restriction that it must be minimally capable of encoding any value that may be stored in a variable of type int, MPI_Aint, or MPI_Offset in C and of type INTEGER, INTEGER(KIND=MPI_ADDRESS_KIND), or INTEGER(KIND=MPI_OFFSET_KIND) in Fortran. Even

1 though the MPI_Count type is large enough to encode address locations, the MPI_Count type
2 shall not be used to represent an *absolute address*.

3
4 *Rationale.* Count values need to be large enough to encode any value used for
5 expressing element counts, strides, offsets, indexes, displacements, typemaps in mem-
6 ory, typemaps in file views, etc. Prior to MPI-4.0, many MPI routines used `int` in C
7 and `INTEGER` in Fortran as the type for *count* arguments. To avoid breaking back-
8 ward compatibility, this version of the standard continues to support `int` in C as
9 well as `INTEGER` in Fortran in such routines. In addition, this version of the stan-
10 dard supports using `MPI_Count` in C (via separate “_c” suffixed procedures) as well
11 as `INTEGER(KIND=MPI_COUNT_KIND)` in Fortran (via polymorphic interfaces in newer
12 MPI Fortran bindings (USE `mpi_f08`)) in such routines. See Section 19.2 for a full
13 explanation. (*End of rationale.*)

14 The phrase **large count** refers to the use of `MPI_Count` and
15 `INTEGER(KIND=MPI_COUNT_KIND)` parameter types.

16 There are cases where `MPI_UNDEFINED` can be returned in a **large count** OUT pa-
17 rameter. Per Table A.1.1 (page 851), the `MPI_UNDEFINED` constant is defined to be a C
18 `int` (or unnamed enum) and a Fortran `INTEGER`. Implementations shall therefore choose the
19 underlying types for `MPI_Count` and `INTEGER(KIND=MPI_COUNT_KIND)` such that they can be
20 compared to `MPI_UNDEFINED`.
21

22 *Advice to implementors.* The comparison of `MPI_UNDEFINED` to an `MPI_Count` or
23 `INTEGER(KIND=MPI_COUNT_KIND)` may need to be via a casting operation. (*End of*
24 *advice to implementors.*)
25

26 2.6 Language Binding

27 This section defines the rules for MPI language binding in general and for Fortran, and ISO
28 C, in particular. (Note that ANSI C has been replaced by ISO C.) Defined here are various
29 object representations, as well as the naming conventions used for expressing this standard.
30 The actual calling sequences are defined elsewhere.
31

32 MPI bindings are for Fortran 90 or later, though they were originally designed to be
33 usable in Fortran 77 environments. With the `mpi_f08` module, two new Fortran features,
34 *assumed type* (i.e., `TYPE(*)`) and *assumed rank* (i.e., `DIMENSION(..)`), are also required, see
35 Section 2.5.5.
36

37 Since the word `PARAMETER` is a keyword in the Fortran language, we use the word
38 “argument” to denote the arguments to a subroutine. These are normally referred to
39 as parameters in C, however, we expect that C programmers will understand the word
40 “argument” (which has no specific meaning in C), thus allowing us to avoid unnecessary
41 confusion for Fortran programmers.

42 Since Fortran is case insensitive, linkers may use either lower case or upper case when
43 resolving Fortran names. Users of case sensitive languages should avoid any prefix of the
44 form “MPI_” and “PMPI_”, where any of the letters are either upper or lower case.

45 2.6.1 Deprecated and Removed Interfaces

46
47 A number of chapters refer to deprecated or replaced MPI constructs. These are constructs
48 that continue to be part of the MPI standard, as documented in Chapter 16, but that users

are recommended not to continue using, since better solutions were provided with newer versions of MPI. For example, the Fortran binding for MPI-1 functions that have address arguments uses `INTEGER`. This is not consistent with the C binding, and causes problems on machines with 32 bit `INTEGER`s and 64 bit addresses. In MPI-2, these functions were given new names with new bindings for the address arguments. The use of the old functions was declared as deprecated. For consistency, here and in a few other cases, new C functions are also provided, even though the new functions are equivalent to the old functions. The old names are deprecated.

Some of the previously deprecated constructs are now removed, as documented in Chapter 17. They may still be provided by an implementation for backwards compatibility, but are not required.

Table 2.1 shows a list of all of the deprecated and removed constructs. Note that some C typedefs and Fortran subroutine names are included in this list; they are the types of callback functions.

2.6.2 Fortran Binding Issues

Originally, MPI-1.1 provided bindings for Fortran 77. These bindings are retained, but they are now interpreted in the context of the Fortran 90 standard. MPI can still be used with most Fortran 77 compilers, as noted below. When the term “Fortran” is used it means Fortran 90 or later; it means Fortran 2008 with TS 29113, which is now an integral part of Fortran 2018 and later if the `mpi_f08` module is used.

All Fortran MPI names have an `MPI_` prefix. Although Fortran is not case sensitive, if the `mpi_f08` module is used, the first character after the `MPI_` prefix is capital and all others are lower case. If the `mpi_f08` module is not used, all characters are capitals. Programs must not declare names, e.g., for variables, subroutines, functions, parameters, derived types, abstract interfaces, or modules, beginning with the prefix `MPI_`. To avoid conflicting with the profiling interface, programs must also avoid subroutines and functions with the prefix `PMPI_`. This is mandated to avoid possible name collisions.

All MPI Fortran subroutines have an error code in the last argument. With `USE mpi_f08`, this last argument is declared as `OPTIONAL`, except for user-defined callback functions (e.g., `COMM_COPY_ATTR_FUNCTION`) and their predefined callbacks (e.g., `MPI_COMM_NULL_COPY_FN`). A few MPI operations that are functions do not have the error code argument. The error code value for successful completion is `MPI_SUCCESS`. Other error codes are implementation dependent; see the error codes in Chapter 9 and Annex A.

Constants representing the maximum length of a string are one smaller in Fortran than in C as discussed in Section 19.3.9.

Handles are represented in Fortran as `INTEGER`s, or as a `BIND(C)` derived type with the `mpi_f08` module; see Section 2.5.1. Binary-valued variables are of type `LOGICAL`.

Array arguments are indexed from one.

The older MPI Fortran bindings—use `mpi` and (deprecated) `mpif.h`—are inconsistent with the Fortran standard in several respects. These inconsistencies, such as register optimization problems, have implications for user codes that are discussed in detail in Section 19.1.16.

The support for large count and displacement in Fortran is only available when using newer MPI Fortran bindings (`USE mpi_f08`). For better readability, all Fortran large count procedure declarations are marked with a comment “!(`_c`)”.

Table 2.1: Deprecated and removed constructs

Deprecated or removed construct	Deprecated since	Removed since	Replacement
MPI_ADDRESS	MPI-2.0	MPI-3.0	MPI_GET_ADDRESS
MPI_TYPE_HINDEXED	MPI-2.0	MPI-3.0	MPI_TYPE_CREATE_HINDEXED
MPI_TYPE_HVECTOR	MPI-2.0	MPI-3.0	MPI_TYPE_CREATE_HVECTOR
MPI_TYPE_STRUCT	MPI-2.0	MPI-3.0	MPI_TYPE_CREATE_STRUCT
MPI_TYPE_EXTENT	MPI-2.0	MPI-3.0	MPI_TYPE_GET_EXTENT
MPI_TYPE_UB	MPI-2.0	MPI-3.0	MPI_TYPE_GET_EXTENT
MPI_TYPE_LB	MPI-2.0	MPI-3.0	MPI_TYPE_GET_EXTENT
MPI_LB ¹	MPI-2.0	MPI-3.0	MPI_TYPE_CREATE_RESIZED
MPI_UB ¹	MPI-2.0	MPI-3.0	MPI_TYPE_CREATE_RESIZED
MPI_ERRHANDLER_CREATE	MPI-2.0	MPI-3.0	MPI_COMM_CREATE_ERRHANDLER
MPI_ERRHANDLER_GET	MPI-2.0	MPI-3.0	MPI_COMM_GET_ERRHANDLER
MPI_ERRHANDLER_SET	MPI-2.0	MPI-3.0	MPI_COMM_SET_ERRHANDLER
MPI_Handler_function ²	MPI-2.0	MPI-3.0	MPI_Comm_errhandler_function ²
MPI_KEYVAL_CREATE	MPI-2.0		MPI_COMM_CREATE_KEYVAL
MPI_KEYVAL_FREE	MPI-2.0		MPI_COMM_FREE_KEYVAL
MPI_DUP_FN ³	MPI-2.0		MPI_COMM_DUP_FN ³
MPI_NULL_COPY_FN ³	MPI-2.0		MPI_COMM_NULL_COPY_FN ³
MPI_NULL_DELETE_FN ³	MPI-2.0		MPI_COMM_NULL_DELETE_FN ³
MPI_Copy_function ²	MPI-2.0		MPI_Comm_copy_attr_function ²
COPY_FUNCTION ²	MPI-2.0		COMM_COPY_ATTR_FUNCTION ²
MPI_Delete_function ²	MPI-2.0		MPI_Comm_delete_attr_function ²
DELETE_FUNCTION ²	MPI-2.0		COMM_DELETE_ATTR_FUNCTION ²
MPI_ATTR_DELETE	MPI-2.0		MPI_COMM_DELETE_ATTR
MPI_ATTR_GET	MPI-2.0		MPI_COMM_GET_ATTR
MPI_ATTR_PUT	MPI-2.0		MPI_COMM_SET_ATTR
MPI_COMBINER_HVECTOR_INTEGER ⁴	-	MPI-3.0	MPI_COMBINER_HVECTOR ⁴
MPI_COMBINER_HINDEXED_INTEGER ⁴	-	MPI-3.0	MPI_COMBINER_HINDEXED ⁴
MPI_COMBINER_STRUCT_INTEGER ⁴	-	MPI-3.0	MPI_COMBINER_STRUCT ⁴
MPI::...	MPI-2.2	MPI-3.0	C language binding
MPI_CANCEL for send requests	MPI-4.0		no direct replacement
MPI_INFO_GET	MPI-4.0		MPI_INFO_GET_STRING
MPI_INFO_GET_VALUELEN	MPI-4.0		MPI_INFO_GET_STRING
MPI_T_ERR_INVALID_ITEM	MPI-4.0		MPI_T_ERR_INVALID_INDEX
MPI_SIZEOF	MPI-4.0		storage_size() ⁵ or c_sizeof()
mpif.h	MPI-4.1		mpi module and mpi_f08 module
MPI_TYPE_SIZE_X	MPI-4.1		MPI_Type_size_c / !(_c) / - ⁶
MPI_TYPE_GET_EXTENT_X	MPI-4.1		MPI_Type_get_extent_c / !(_c) / - ⁶
MPI_TYPE_GET_TRUE_EXTENT_X	MPI-4.1		MPI_Type_get_true_extent_c / !(_c) / - ⁶
MPI_GET_ELEMENTS_X	MPI-4.1		MPI_Get_elements_c / !(_c) / - ⁶
MPI_STATUS_SET_ELEMENTS_X	MPI-4.1		MPI_Status_set_elements_c / !(_c) / - ⁶
MPI_HOST	MPI-4.1		no direct replacement

¹ Predefined datatype.
² Callback prototype definition.
³ Predefined callback routine.
⁴ Constant.
⁵ Fortran intrinsic. storage_size() returns the size in bits instead of bytes; see Section 16.3.
⁶ in C / Fortran with the mpi_f08 module. No substitute for the mpi module and mpif.h.
Other entries are regular MPI routines.

2.6.3 C Binding Issues

We use the ISO C declaration format. All MPI names have an MPI_ prefix, defined constants are in all capital letters, and defined types and functions have one capital letter after the prefix. Programs must not declare names (identifiers), e.g., for variables, functions, constants, types, or macros, beginning with any prefix of the form MPI_, where any of the letters are either upper or lower case. To support the profiling interface, programs must not declare functions with names beginning with any prefix of the form PMPI_, where any

of the letters are either upper or lower case. 1

The definition of named constants, function prototypes, and type definitions must be supplied in an include file `mpi.h`. 2

Almost all C functions return an error code. The successful return value will be `MPI_SUCCESS`, but error codes raised after a failure are implementation dependent. 3

Type declarations are provided for handles to each category of opaque objects. 4

Array arguments are indexed from zero. 5

Logical flags are integers with value 0 meaning “false” and a nonzero value meaning “true.” 6

Choice arguments are pointers of type `void*`. 7

2.6.4 Functions and Macros 8

An implementation is allowed to implement `MPI_AINT_ADD`, `PMPI_AINT_ADD`, `MPI_AINT_DIFF`, and `PMPI_AINT_DIFF`, and no others, as macros in C. 9

Advice to implementors. Implementors should document which routines are implemented as macros. (*End of advice to implementors.*) 10

Advice to users. If these routines are implemented as macros, they will not work with the MPI profiling interface. (*End of advice to users.*) 11

2.7 Processes 12

An MPI program consists of autonomous processes, executing their own code, in an MIMD style. The codes executed by each process need not be identical. The processes communicate via calls to MPI communication primitives. Typically, each process executes in its own address space, although shared-memory implementations of MPI are possible. 13

This document specifies the behavior of a parallel program assuming that only MPI calls are used. The interaction of an MPI program with other possible means of communication, I/O, and process management is not specified. Unless otherwise stated in the specification of the standard, MPI places no requirements on the result of its interaction with external mechanisms that provide similar or equivalent functionality. This includes, but is not limited to, interactions with external mechanisms for process control, shared and remote memory access, file system access and control, interprocess communication, process signaling, and terminal I/O. High quality implementations should strive to make the results of such interactions intuitive to users, and attempt to document restrictions where deemed necessary. 14

Advice to implementors. Implementations that support such additional mechanisms for functionality supported within MPI are expected to document how these interact with MPI. (*End of advice to implementors.*) 15

The interaction of MPI and threads is defined in Section 11.6. 16

MPI processes reside in the same **shared memory domain** if it is possible to share a segment of memory between them, i.e., to make a segment of memory (**shared memory segment**) concurrently accessible from all of those MPI processes through load/store accesses. For a group of processes belonging to more than one *shared memory domain* the creation of a subgroup of processes belonging to the same *shared memory domain* is defined in Section 7.4.2. 17

2.8 Error Handling

MPI provides the user with reliable message transmission. A message sent is always received correctly, and the user does not need to check for transmission errors, time-outs, or other error conditions. In other words, MPI does not provide mechanisms for dealing with **transmission failures** in the communication system. If the MPI implementation is built on an unreliable underlying mechanism, then it is the job of the implementor of the MPI subsystem to insulate the user from this unreliability, and to reflect only unrecoverable transmission failures. Whenever possible, such failures will be reflected as errors in the relevant communication call.

Similarly, MPI itself provides no mechanisms for handling **MPI process failures**, that is, when an MPI process unexpectedly and permanently stops communicating (e.g., a software or hardware crash results in an MPI process terminating unexpectedly).

Of course, MPI programs may still be erroneous. A **program error** can occur when an MPI call is made with an incorrect argument (nonexisting destination in a send operation, buffer too small in a receive operation, etc.). This type of error would occur in any implementation. In addition, a **resource error** may occur when a program exceeds the amount of available system resources (number of pending messages, system buffers, etc.). The occurrence of this type of error depends on the amount of available resources in the system and the resource allocation mechanism used; this may differ from system to system. A high-quality implementation will provide generous limits on the important resources so as to alleviate the portability problem this represents.

In C and Fortran, almost all MPI calls return a code that indicates successful completion of the operation. Whenever possible, MPI calls return an error code if an error occurred during the call. By default, an error detected during the execution of the MPI library causes the parallel computation to abort, except for file operations. However, MPI provides mechanisms for users to change this default and to handle recoverable errors. The user may specify that no error is fatal, and handle error codes returned by MPI calls by themselves. Also, the user may provide user-defined error-handling routines, which will be invoked whenever an MPI call returns abnormally. The MPI error handling facilities are described in Section 9.3.

Several factors limit the ability of MPI calls to return with meaningful error codes when an error occurs. MPI may not be able to detect some errors; other errors may be too expensive to detect in normal execution mode; some faults (e.g., memory faults) may corrupt the state of the MPI library and its outputs; finally some errors may be “catastrophic” and may prevent MPI from returning control to the caller.

In addition, some errors may be detected in operations that do not refer to an MPI object from which the associated error handler can be obtained. Error handler associations are further described in Section 9.3. In such cases, these errors will be raised on the communicator `MPI_COMM_SELF` when using the World Model (see Section 11.2). When `MPI_COMM_SELF` is not initialized (i.e., before `MPI_INIT` / `MPI_INIT_THREAD`, after `MPI_FINALIZE`, or when using the Sessions Model exclusively) the error raises the **initial error handler** (set during the launch operation, see 11.8.4). The Sessions Model is described in Section 11.3.

Lastly, some errors may be detected after the associated operation has completed locally. An example of such a case arises because of the nature of asynchronous communications: MPI calls may initiate operations that continue asynchronously after the call returned. Thus, the operation may return with a code indicating successful completion,

yet later cause an error to be raised. If there is a subsequent call that relates to the same operation (e.g., a call that verifies that an asynchronous operation has completed) then the error argument associated with this call will be used to indicate the nature of the error. In a few cases, the error may occur after all calls that relate to the operation have returned, so that no error value can be used to indicate the nature of the error (e.g., an erroneous program on the receiver in a send with the ready mode).

This document does not specify the state of a computation after an erroneous MPI call has occurred. The desired behavior is that a relevant error code be returned, and the effect of the error be localized to the greatest possible extent. E.g., it is highly desirable that an erroneous receive call will not cause any part of the receiver's memory to be overwritten, beyond the area specified for receiving the message.

Implementations may go beyond this document in supporting in a meaningful manner MPI calls that are defined here to be erroneous. For example, MPI specifies strict type matching rules between matching send and receive operations: it is erroneous to send a floating point variable and receive an integer. Implementations may go beyond these type matching rules, and provide automatic type conversion in such situations. It will be helpful to generate warnings for such nonconforming behavior.

MPI defines a way for users to create new error codes as defined in Section 9.5.

2.9 Progress

MPI communication operations or parallel I/O patterns typically comprise several related operations executed in one or multiple MPI processes. Examples are the point-to-point communications with one MPI process executing a send operation and another (or the same) MPI process executing a receive operation, or all MPI processes of a group executing a collective operation.

Within each MPI process parts of the communication or parallel I/O pattern are executed within the MPI procedure calls that belong to the operation in that MPI process, whereas other parts are **decoupled MPI activities**, i.e., they may be executed within an additional progress thread, offloaded to the network interface controller (NIC), or executed within other MPI procedure calls that are not semantically related to the given communication or parallel I/O pattern.

An MPI procedure invocation is **blocked** if it delays its return until some specific activity or state-change has occurred in another MPI process. An MPI procedure call that is *blocked* can be

- a *nonlocal* MPI procedure call that delays its return until a specific semantically-related MPI call on another MPI process, or
- a *local* MPI procedure call that delays its return until some unspecific MPI call in another MPI process causes a specific state-change in that other MPI process, or
- an MPI finalization procedure (MPI_FINALIZE or MPI_SESSION_FINALIZE) that delays its return or exit because this MPI finalization must guarantee that all decoupled MPI activities that are related to that MPI finalization call in the calling MPI process will be executed before this MPI finalization is finished. Note that an MPI finalization procedure may execute attribute deletion callback functions prior to the finalization (see Section 11.2.4); these callback functions may generate additional decoupled MPI activities.

1 Some examples of a *nonlocal* blocked MPI procedure call:

- 2 • MPI_SSEND delays its return until the matching receive operation is *started* at the
- 3 destination MPI process (for example, by a call to MPI_RECV or to MPI_IRecv).
- 4
- 5 • MPI_RECV delays its return until the matching send operation is *started* at the source
- 6 MPI process (for example, by a call to MPI_SEND or to MPI_ISEND).
- 7

8 Some examples of a *local* blocked MPI procedure call:

- 9 • MPI_RSEND, if the message data cannot be entirely buffered, delays its return until
- 10 the destination MPI process has received the portion of message data that cannot
- 11 be buffered, which may require one or more unspecific MPI procedure call(s) at the
- 12 destination MPI process.
- 13
- 14 • MPI_RECV, in case the message was buffered at the sending MPI process (e.g. with
- 15 MPI_BSEND), delays its return until the message is received, which may require one
- 16 or more unspecific MPI procedure calls at the sending MPI process to send the buffered
- 17 data.

18 All MPI processes are required to **guarantee progress**, i.e., all decoupled MPI activi-
 19 ties will eventually be executed. This guarantee is required to be provided during

- 20 • blocked MPI procedures, and
- 21
- 22 • repeatedly called MPI test procedures (see below) that return `flag=false`.
- 23

24 The *progress* must be provided independently of whether a decoupled MPI activity
 25 belongs to a specific session or to the World Model(see Sections 11.2 and 11.3). Other ways
 26 of fulfilling this guarantee are possible and permitted (for example, a dedicated progress
 27 thread or off-loading to a network interface controller (NIC)).

28 MPI test procedures are MPI_TEST, MPI_TESTANY, MPI_TESTALL,
 29 MPI_TESTSOME, MPI_IPROBE, MPI_IMPROBE, MPI_REQUEST_GET_STATUS,
 30 MPI_WIN_TEST, and MPI_PARRIVED.

31 **Strong progress** is provided by an MPI implementation if all *local* procedures return
 32 independently of MPI procedure calls in other MPI processes (operation-related or not). An
 33 MPI implementation provides **weak progress** if it does not provide *strong progress*.

34
 35 *Advice to users.* The type of *progress* may influence the performance of MPI opera-
 36 tions. A correct MPI application must be written under the assumption that only
 37 *weak progress* is provided. Every MPI application that is correct under *weak progress*
 38 will be correctly executed if *strong progress* is provided. In addition, the MPI stan-
 39 dard is designed such that correctness under the assumption of *strong progress* should
 40 imply also correctness if only *weak progress* is provided by the implementation. (*End*
 41 *of advice to users.*)

42
 43 *Rationale.* MPI does not guarantee progress when using synchronization methods
 44 that are not based on MPI procedures. Without guaranteed *strong progress* in MPI
 45 this may lead to a *deadlock*, see for example Section 2.7 and Example 12.13 in Sec-
 46 tion 12.7.3. (*End of rationale.*)

47 For further rules, see in Section 2.4.2 the definition of *local* MPI procedures, and all
 48 references to *progress* in the [general index](#).

2.10 Implementation Issues

There are a number of areas where an MPI implementation may interact with the operating environment and system. While MPI does not mandate that any services (such as signal handling) be provided, it does strongly suggest the behavior to be provided if those services are available. This is an important point in achieving portability across platforms that provide the same set of services.

2.10.1 Independence of Basic Runtime Routines

MPI programs require that library routines that are part of the basic language environment (such as `write` in Fortran and `printf` and `malloc` in ISO C) and are executed after `MPI_INIT` and before `MPI_FINALIZE` operate independently and that their *completion* is independent of the action of other processes in an MPI program.

Note that this in no way prevents the creation of library routines that provide parallel services whose operation is collective. However, the following program is expected to complete in an ISO C environment regardless of the size of `MPI_COMM_WORLD` (assuming that `printf` is available at the executing MPI processes).

```
int commworld_rank;
MPI_Init((void *)0, (void *)0);
MPI_Comm_rank(MPI_COMM_WORLD, &commworld_rank);
if (commworld_rank == 0) printf("Starting program\n");
MPI_Finalize();
```

The corresponding Fortran programs are also expected to complete.

An example of what is *not* required is any particular ordering of the action of these routines when called by several MPI processes. For example, MPI makes neither requirements nor recommendations for the output from the following program (again assuming that I/O is available at the executing MPI processes).

```
MPI_Comm_rank(MPI_COMM_WORLD, &commworld_rank);
printf("Output from MPI process where commworld_rank=%d\n",
      commworld_rank);
```

In addition, calls that fail because of resource exhaustion or other error are not considered a violation of the requirements here (however, they are required to complete, just not to complete successfully).

2.10.2 Interaction with Signals

MPI does not specify the interaction of processes with signals and does not require that MPI be signal safe. The implementation may reserve some signals for its own use. It is required that the implementation document which signals it uses, and it is strongly recommended that it not use `SIGALRM`, `SIGFPE`, or `SIGIO`. Implementations may also prohibit the use of MPI calls from within signal handlers.

In multithreaded environments, users can avoid conflicts between signals and the MPI library by catching signals only on threads that do not execute MPI calls. High quality single-threaded implementations will be signal safe: an MPI call suspended by a signal will resume and complete normally after the signal is handled.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

2.11 Examples

The examples in this document are for illustration purposes only. They are not intended to specify the standard. Many of the examples have been compiled by tools that extract the examples from the source files for the MPI standard. However, the examples have not been carefully checked or verified.

DRAFT

Chapter 3

Point-to-Point Communication

3.1 Introduction

Sending and receiving of *messages* by MPI processes is the basic MPI communication mechanism. The basic point-to-point communication operations are *send* and *receive*. Their use is illustrated in Example 3.1.

Example 3.1. A simple ‘hello world’ example usage of point-to-point communication.

```
#include "mpi.h"
int main(int argc, char *argv[])
{
    char message[20];
    int myrank;
    MPI_Status status;
    MPI_Init(&argc, &argv);
    MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
    if (myrank == 0) /* code for process zero */
    {
        strcpy(message, "Hello, there");
        MPI_Send(message, strlen(message)+1, MPI_CHAR, 1, 99,
                 MPI_COMM_WORLD);
    }
    else if (myrank == 1) /* code for process one */
    {
        MPI_Recv(message, 20, MPI_CHAR, 0, 99, MPI_COMM_WORLD, &status);
        printf("received :%s:\n", message);
    }
    MPI_Finalize();
    return 0;
}
```

In Example 3.1, process zero (`myrank = 0`, strictly ‘the MPI process with rank 0 in communicator `MPI_COMM_WORLD`’) sends a *message* to process one using the *send* operation `MPI_SEND`. The operation specifies a *send buffer* in the sender memory from which the *message data* is taken. In the example above, the send buffer consists of the storage containing the variable `message` in the memory of process zero. The location, size and type of the send buffer are specified by the first three parameters of the send operation. The message sent will contain the 13 characters of this variable. In addition, the send operation associates an *envelope* with the message. This *envelope* specifies the message destination and contains distinguishing information that can be used by the *receive* operation to select a particular message. The last three parameters of the send operation, along with the rank

of the sender, specify the *envelope* for the message sent.

Process one (`myrank = 1`, strictly ‘the MPI process with rank 1 in communicator `MPI_COMM_WORLD`’) receives this message with the *receive* operation `MPI_RECV`. The message to be received is selected according to the value of its *envelope*, and the *message data* is stored into the *receive buffer*. In the example above, the receive buffer consists of the storage containing the string `message` in the memory of process one. The first three parameters of the receive operation specify the location, size and type of the receive buffer. The next three parameters are used for selecting the incoming message. The last parameter is used to return information on the message just received.

Advice to users. Colloquial usage commonly permits references to “rank 0” or “process 0”, which are strictly ambiguous and ideally should be qualified by including the relevant context, for example, the MPI communicator in the case above. (*End of advice to users.*)

The next sections describe the blocking send and receive operations. We discuss send, receive, blocking communication semantics, type matching requirements, type conversion in heterogeneous environments, and more general communication modes. Nonblocking communication is addressed next, followed by probing and cancelling a message, channel-like constructs and send-receive operations, ending with a description of the “dummy” MPI process, `MPI_PROC_NULL`.

3.2 Blocking Send and Receive Operations

3.2.1 Blocking Send

The syntax of the **blocking send** procedure is given below.

```
MPI_SEND(buf, count, datatype, dest, tag, comm)
```

IN	<code>buf</code>	initial address of send buffer (choice)
IN	<code>count</code>	number of elements in send buffer (non-negative integer)
IN	<code>datatype</code>	datatype of each send buffer element (handle)
IN	<code>dest</code>	rank of destination (integer)
IN	<code>tag</code>	message tag (integer)
IN	<code>comm</code>	communicator (handle)

C binding

```
int MPI_Send(const void *buf, int count, MPI_Datatype datatype, int dest,
             int tag, MPI_Comm comm)
```

```
int MPI_Send_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
               int dest, int tag, MPI_Comm comm)
```

Fortran 2008 binding

```
MPI_Send(buf, count, datatype, dest, tag, comm, ierror)
```

Table 3.1: Predefined MPI datatypes corresponding to Fortran datatypes

MPI datatype	Fortran datatype
MPI_INTEGER	INTEGER
MPI_REAL	REAL
MPI_DOUBLE_PRECISION	DOUBLE PRECISION
MPI_COMPLEX	COMPLEX
MPI_LOGICAL	LOGICAL
MPI_CHARACTER	CHARACTER(1)
MPI_BYTE	
MPI_PACKED	

```

TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER, INTENT(IN) :: count, dest, tag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Send(buf, count, datatype, dest, tag, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: dest, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_SEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
<type> BUF(*)
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR

```

The blocking semantics of this call are described in Section 3.4.

3.2.2 Message Data

The send buffer specified by the `MPI_SEND` procedure consists of `count` successive entries of the type indicated by `datatype`, starting with the entry at address `buf`. Note that we specify the message length in terms of number of *elements*, not number of *bytes*. The former is machine independent and closer to the application level.

The data part of the message consists of a sequence of `count` values, each of the type indicated by `datatype`. `count` may be zero, in which case the data part of the message is empty. The **basic datatypes** that can be specified for message data values correspond to the basic datatypes of the host language. Possible values of this argument for Fortran and the corresponding Fortran types are listed in Table 3.1. Possible values for this argument for C and the corresponding C types are listed in Table 3.2.

The datatypes `MPI_BYTE` and `MPI_PACKED` do not correspond to a Fortran or C datatype. A value of type `MPI_BYTE` consists of a byte (8 binary digits). A byte is uninterpreted and is different from a character. Different machines may have different representations for characters, or may use more than one byte to represent characters. On the other

Table 3.2: Predefined MPI datatypes corresponding to C datatypes

MPI datatype	C datatype
MPI_CHAR	char (treated as printable character)
MPI_SHORT	signed short int
MPI_INT	signed int
MPI_LONG	signed long int
MPI_LONG_LONG_INT	signed long long int
MPI_LONG_LONG (as a synonym)	signed long long int
MPI_SIGNED_CHAR	signed char (treated as integral value)
MPI_UNSIGNED_CHAR	unsigned char (treated as integral value)
MPI_UNSIGNED_SHORT	unsigned short int
MPI_UNSIGNED	unsigned int
MPI_UNSIGNED_LONG	unsigned long int
MPI_UNSIGNED_LONG_LONG	unsigned long long int
MPI_FLOAT	float
MPI_DOUBLE	double
MPI_LONG_DOUBLE	long double
MPI_WCHAR	wchar_t (defined in <stddef.h>) (treated as printable character)
MPI_C_BOOL	_Bool
MPI_INT8_T	int8_t
MPI_INT16_T	int16_t
MPI_INT32_T	int32_t
MPI_INT64_T	int64_t
MPI_UINT8_T	uint8_t
MPI_UINT16_T	uint16_t
MPI_UINT32_T	uint32_t
MPI_UINT64_T	uint64_t
MPI_C_COMPLEX	float _Complex
MPI_C_FLOAT_COMPLEX (as a synonym)	float _Complex
MPI_C_DOUBLE_COMPLEX	double _Complex
MPI_C_LONG_DOUBLE_COMPLEX	long double _Complex
MPI_BYTE	
MPI_PACKED	

hand, a byte has the same binary value on all machines. The use of the type MPI_PACKED is explained in Section 5.2.

MPI requires support of these datatypes, which match the basic datatypes of Fortran and ISO C. Additional MPI datatypes should be provided if the host language has additional datatypes¹: MPI_DOUBLE_COMPLEX for double precision complex in Fortran de-

¹These types, such as DOUBLE_COMPLEX and INTEGER*4, are not specified by any Fortran standard but are

Table 3.3: Predefined MPI datatypes corresponding to both C and Fortran datatypes

MPI datatype	C datatype	Fortran datatype
MPI_AINT	MPI_Aint	INTEGER(KIND=MPI_ADDRESS_KIND)
MPI_OFFSET	MPI_Offset	INTEGER(KIND=MPI_OFFSET_KIND)
MPI_COUNT	MPI_Count	INTEGER(KIND=MPI_COUNT_KIND)

Table 3.4: Predefined MPI datatypes corresponding to C++ datatypes

MPI datatype	C++ datatype
MPI_CXX_BOOL	bool
MPI_CXX_FLOAT_COMPLEX	std::complex<float>
MPI_CXX_DOUBLE_COMPLEX	std::complex<double>
MPI_CXX_LONG_DOUBLE_COMPLEX	std::complex<long double>

clared to be of type `DOUBLE COMPLEX`; `MPI_REAL2`, `MPI_REAL4`, `MPI_REAL8`, and `MPI_REAL16` for Fortran reals, declared to be of type `REAL*2`, `REAL*4`, `REAL*8`, and `REAL*16`, respectively; `MPI_INTEGER1`, `MPI_INTEGER2`, `MPI_INTEGER4`, and `MPI_INTEGER8` for Fortran integers, declared to be of type `INTEGER*1`, `INTEGER*2`, `INTEGER*4`, and `INTEGER*8`, respectively; `MPI_COMPLEX4`, `MPI_COMPLEX8`, `MPI_COMPLEX16`, and `MPI_COMPLEX32` for complex numbers in Fortran declared to be of type `COMPLEX*4`, `COMPLEX*8`, `COMPLEX*16`, and `COMPLEX*32`, respectively; etc.

Rationale. One goal of the design is to allow for MPI to be implemented as a library, with no need for additional preprocessing or compilation. Thus, one cannot assume that a communication call has information on the datatype of variables in the communication buffer; this information must be supplied by an explicit argument. The need for such datatype information will become clear in Section 3.3.2. (*End of rationale.*)

The datatypes `MPI_AINT`, `MPI_OFFSET`, and `MPI_COUNT` correspond to the MPI-defined C types `MPI_Aint`, `MPI_Offset`, and `MPI_Count` and their Fortran equivalents `INTEGER(KIND=MPI_ADDRESS_KIND)`, `INTEGER(KIND=MPI_OFFSET_KIND)`, and `INTEGER(KIND=MPI_COUNT_KIND)`. This is described in Table 3.3. All predefined datatype handles are available in all language bindings. See Sections 19.3.6 and 19.3.10 on page 840 and 847 for information on interlanguage communication with these types.

If there is an accompanying C++ compiler then the datatypes in Table 3.4 are also supported in C and Fortran.

3.2.3 Message Envelope

In addition to the data part, messages carry information that can be used to distinguish messages and selectively receive them. This information consists of a fixed number of fields, which we collectively call the **message envelope**. These fields are

source

extensions commonly accepted by Fortran compilers.

1 **destination**
 2 **tag**
 3 **communicator**

4 The *message source* is implicitly determined by the identity of the message sender. The
 5 other fields are specified by arguments in the send procedure.

6 The *message destination* is specified by the **dest** argument.

7 The integer-valued *message tag* is specified by the **tag** argument. This integer can be
 8 used by the program to distinguish different types of messages. The range of valid tag
 9 values is $0, \dots, \text{UB}$, where the value of **UB** is implementation dependent. It can be found by
 10 querying the value of the attribute `MPI_TAG_UB`, as described in Chapter 9. MPI requires
 11 that **UB** be no less than 32767.

12 The **comm** argument specifies the *communicator* that is used for the send operation.
 13 Communicators are explained in Chapter 7; below is a brief summary of their usage.

14 A communicator specifies the communication context for a communication operation.
 15 Each communication context provides a separate “communication universe”: messages are
 16 always received within the context they were sent, and messages sent in different contexts
 17 do not interfere.

18 The communicator also specifies the group of MPI processes that share this commu-
 19 nication context. This MPI *process group* is ordered and MPI processes are identified by
 20 their rank within this group. Thus, the range of valid values for **dest** is $0, \dots, n - 1 \cup$
 21 $\{\text{MPI_PROC_NULL}\}$, where n is the number of MPI processes in the group. (If the communi-
 22 cator is an inter-communicator, then destinations are identified by their rank in the remote
 23 group. See Chapter 7.)

24 An MPI process may have a different rank in each group in which it is a member.

25 When using the World Model (see Section 11.2), a predefined communicator
 26 `MPI_COMM_WORLD` is provided by MPI. It allows communication with all MPI processes
 27 that are accessible after MPI initialization and MPI processes are identified by their rank in
 28 the group of `MPI_COMM_WORLD`.

29
 30 *Advice to users.* Users that are comfortable with the notion of a flat name space
 31 for MPI processes, and a single communication context, as offered by most existing
 32 communication libraries, need only use the World Model for MPI initialization, and
 33 the predefined variable `MPI_COMM_WORLD` as the **comm** argument. This will allow
 34 communication with all the MPI processes available at initialization time.

35 Users may define new communicators, as explained in Chapter 7. Communicators
 36 provide an important encapsulation mechanism for libraries and modules. They allow
 37 modules to have their own disjoint communication universe and their own MPI process
 38 numbering scheme. (*End of advice to users.*)

39
 40 *Advice to implementors.* The *message envelope* would normally be encoded by a
 41 fixed-length message header. However, the actual encoding is implementation de-
 42 pendent. Some of the information (e.g., source or destination) may be implicit, and
 43 need not be explicitly carried by messages. Also, MPI processes may be identified by
 44 relative ranks, or absolute ids, etc. (*End of advice to implementors.*)

46 3.2.4 Blocking Receive

47
 48 The syntax of the **blocking receive** procedure is given below.

MPI_RECV(buf, count, datatype, source, tag, comm, status)			1
OUT	buf	initial address of receive buffer (choice)	2
			3
IN	count	number of elements in receive buffer (non-negative integer)	4
			5
IN	datatype	datatype of each receive buffer element (handle)	6
			7
IN	source	rank of source or MPI_ANY_SOURCE (integer)	8
			9
IN	tag	message tag or MPI_ANY_TAG (integer)	10
			11
IN	comm	communicator (handle)	12
			13
OUT	status	status object (status)	14
			15

C binding

```
int MPI_Recv(void *buf, int count, MPI_Datatype datatype, int source, int tag,
             MPI_Comm comm, MPI_Status *status)
```

```
int MPI_Recv_c(void *buf, MPI_Count count, MPI_Datatype datatype, int source,
               int tag, MPI_Comm comm, MPI_Status *status)
```

Fortran 2008 binding

```
MPI_Recv(buf, count, datatype, source, tag, comm, status, ierror)
```

```
TYPE(*), DIMENSION(..) :: buf
INTEGER, INTENT(IN) :: count, source, tag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Recv(buf, count, datatype, source, tag, comm, status, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: source, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_RECV(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, STATUS, IERROR)
```

```
<type> BUF(*)
INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
```

The blocking semantics of this call are described in Section 3.4.

The receive buffer consists of the storage containing `count` consecutive elements of the type specified by `datatype`, starting at address `buf`. The length of the received message must be less than or equal to the length of the receive buffer. An overflow error occurs if all incoming data does not fit, without truncation, into the receive buffer.

If a message that is shorter than the receive buffer arrives, then only those locations corresponding to the (shorter) message are modified.

1 *Advice to users.* The MPI_PROBE function described in Section 3.8 can be used to
2 receive messages of unknown length. (*End of advice to users.*)

3
4 *Advice to implementors.* Even though no specific behavior is mandated by MPI for
5 *erroneous programs*, the recommended handling of overflow situations is to return in
6 status information about the source and tag of the incoming message. The receive
7 procedure will return an error code. High-quality implementations will also ensure
8 that no memory that is outside the receive buffer will ever be overwritten.

9
10 In the case of a message shorter than the receive buffer, MPI is quite strict in that it
11 allows no modification of the other locations. A more lenient statement would allow
12 for some optimizations but this is not allowed. The implementation must be ready to
13 end a copy into the receiver memory exactly at the end of the receive buffer, even if
14 it is an odd address. (*End of advice to implementors.*)

15 The selection of a message by a receive operation is governed by the value of the
16 *message envelope*. A message can be received by a receive operation if its *envelope* matches
17 the source, tag and comm values specified by the receive operation. The receiver may specify
18 a **wildcard** MPI_ANY_SOURCE value for source, and/or a wildcard MPI_ANY_TAG value for
19 tag, indicating that any source and/or tag are acceptable. It cannot specify a wildcard value
20 for comm. Thus, a message can be received by a receive operation only if it is addressed to
21 the receiving MPI process, has a matching communicator, has matching source unless source
22 = MPI_ANY_SOURCE in the pattern, and has a matching tag unless tag = MPI_ANY_TAG in
23 the pattern.

24 The message tag is specified by the tag argument of the receive operation. The
25 argument source, if different from MPI_ANY_SOURCE, is specified as a rank within the
26 MPI process group associated with that same communicator (remote MPI process group,
27 for inter-communicators). Thus, the range of valid values for the source argument is
28 $\{0, \dots, n - 1\} \cup \{\text{MPI_ANY_SOURCE}\} \cup \{\text{MPI_PROC_NULL}\}$, where n is the number of MPI
29 processes in this group.

30 Note the asymmetry between send and receive operations: A receive operation may
31 accept messages from an arbitrary sender, on the other hand, a send operation must specify
32 a unique receiver. This matches a “push” communication mechanism, where data transfer
33 is effected by the sender (rather than a “pull” mechanism, where data transfer is effected
34 by the receiver).

35 Source = destination is allowed, that is, an MPI process can send a message to itself.
36 However, it is unsafe to do so with the blocking send and receive operations described above,
37 since this may lead to deadlock. See Section 3.5.

38
39 *Advice to implementors.* Message context and other communicator information can
40 be implemented as an additional tag field. It differs from the regular message tag
41 in that wild card matching is not allowed on this field, and that value setting for
42 this field is controlled by communicator manipulation functions. (*End of advice to*
43 *implementors.*)

44
45 The use of dest = MPI_PROC_NULL or source = MPI_PROC_NULL to define a “dummy”
46 destination or source in any send or receive call is described in Section 3.10.

3.2.5 Return Status

The source or tag of a received message may not be known if wildcard values were used in the receive operation. Also, if multiple requests are completed by a single MPI function (see Section 3.7.5), a distinct error code may need to be returned for each request. The information is returned by the `status` argument of `MPI_RECV`. The type of `status` is MPI-defined. Status variables need to be explicitly allocated by the user, that is, they are not system objects.

In C, `status` is a structure that contains three fields named `MPI_SOURCE`, `MPI_TAG`, and `MPI_ERROR`; the structure may contain additional fields. Thus, `status.MPI_SOURCE`, `status.MPI_TAG`, and `status.MPI_ERROR` contain the source, tag, and error code, respectively, of the received message.

In Fortran with `USE mpi` or (deprecated) `INCLUDE 'mpif.h'`, `status` is an array of `INTEGER`s of size `MPI_STATUS_SIZE`. The constants `MPI_SOURCE`, `MPI_TAG`, and `MPI_ERROR` are the indices of the entries that store the source, tag, and error fields. Thus, `status(MPI_SOURCE)`, `status(MPI_TAG)`, and `status(MPI_ERROR)` contain, respectively, the source, tag, and error code of the received message.

With Fortran `USE mpi_f08`, `status` is defined as the Fortran `BIND(C)` derived type `TYPE(MPI_Status)` containing three public `INTEGER` fields named `MPI_SOURCE`, `MPI_TAG`, and `MPI_ERROR`. `TYPE(MPI_Status)` may contain additional, implementation-specific fields. Thus, `status%MPI_SOURCE`, `status%MPI_TAG`, and `status%MPI_ERROR` contain the source, tag, and error code of a received message respectively. Additionally, within both the `mpi` and the `mpi_f08` modules, the constants `MPI_STATUS_SIZE`, `MPI_SOURCE`, `MPI_TAG`, `MPI_ERROR`, and `TYPE(MPI_Status)` are defined to allow conversion between both status representations. Conversion routines are provided in Section 19.3.5.

Rationale. The Fortran `TYPE(MPI_Status)` is defined as a `BIND(C)` derived type so that it can be used at any location where the status integer array representation can be used, e.g., in user defined common blocks. (*End of rationale.*)

Rationale. It is allowed to have the same name (e.g., `MPI_SOURCE`) defined as a constant (e.g., Fortran parameter) and as a field of a derived type. (*End of rationale.*)

In general, message-passing calls do not modify the value of the error code field of status variables. This field may be updated only by the functions in Section 3.7.5 that return multiple statuses. The field is updated if and only if such function returns with an error code of `MPI_ERR_IN_STATUS`.

Rationale. The error field in `status` is not needed for calls that return only one status, such as `MPI_WAIT`, since that would only duplicate the information returned by the function itself. The current design avoids the additional overhead of setting it, in such cases. The field is needed for calls that return multiple statuses, since each request may have had a different failure. (*End of rationale.*)

The `status` argument also returns information on the length of the message received. However, this information is not directly available as a field of the status variable and a call to `MPI_GET_COUNT` is required to “decode” this information.

1 MPI_GET_COUNT(status, datatype, count)

2 IN status return status of receive operation (status)
 3
 4 IN datatype datatype of each receive buffer entry (handle)
 5
 6 OUT count number of received entries (integer)

7 C binding

8 int MPI_Get_count(const MPI_Status *status, MPI_Datatype datatype, int *count)

9
 10 int MPI_Get_count_c(const MPI_Status *status, MPI_Datatype datatype,
 11 MPI_Count *count)

12 Fortran 2008 binding

13 MPI_Get_count(status, datatype, count, ierror)

14 TYPE(MPI_Status), INTENT(IN) :: status
 15 TYPE(MPI_Datatype), INTENT(IN) :: datatype
 16 INTEGER, INTENT(OUT) :: count
 17 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

18 MPI_Get_count(status, datatype, count, ierror) !(_c)

19 TYPE(MPI_Status), INTENT(IN) :: status
 20 TYPE(MPI_Datatype), INTENT(IN) :: datatype
 21 INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: count
 22 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

24 Fortran binding

25 MPI_GET_COUNT(STATUS, DATATYPE, COUNT, IERROR)
 26 INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR

27 Returns the number of entries received. (Again, we count *entries*, each of type *datatype*,
 28 not *bytes*.) The *datatype* argument should match the argument provided by the receive call
 29 that set the *status* variable. If the number of entries received exceeds the limits of the *count*
 30 parameter, then `MPI_GET_COUNT` sets the value of *count* to `MPI_UNDEFINED`. There are
 31 other situations where the value of *count* can be set to `MPI_UNDEFINED`; see Section 5.1.11.

32
 33 *Rationale.* Some message-passing libraries use `INOUT count`, `tag` and `source` argu-
 34 ments, thus using them both to specify the selection criteria for incoming messages
 35 and return the actual *envelope* values of the received message. The use of a separate
 36 status argument prevents errors that are often attached with `INOUT` argument (e.g.,
 37 using the `MPI_ANY_TAG` constant as the tag in a receive). Some libraries use calls that
 38 refer implicitly to the “last message received.” This is not thread safe.

39 The *datatype* argument is passed to `MPI_GET_COUNT` so as to improve performance.
 40 A message might be received without counting the number of elements it contains,
 41 and the *count* value is often not needed. Also, this allows the same function to be
 42 used after a call to `MPI_PROBE` or `MPI_IPROBE`. With a *status* from `MPI_PROBE`
 43 or `MPI_IPROBE`, the same datatypes are allowed as in a call to `MPI_RECV` to receive
 44 this message. (*End of rationale.*)

45
 46 The value returned as the *count* argument of `MPI_GET_COUNT` for a *datatype* of length
 47 zero where zero bytes have been transferred is zero. If the number of bytes transferred is
 48 greater than zero, `MPI_UNDEFINED` is returned.

Rationale. Zero-length datatypes may be created in a number of cases. An important case is `MPI_TYPE_CREATE_DARRAY`, where the definition of the particular darray results in an empty block on some MPI process. Programs written in an SPMD style will not check for this special case and may want to use `MPI_GET_COUNT` to check the status. (*End of rationale.*)

Advice to users. The buffer size required for the receive can be affected by data conversions and by the stride of the receive datatype. In most cases, the safest approach is to use the same datatype with `MPI_GET_COUNT` and the receive. (*End of advice to users.*)

All send and receive operations use the `buf`, `count`, `datatype`, `source`, `dest`, `tag`, `comm`, and `status` arguments in the same way as the blocking `MPI_SEND` and `MPI_RECV` procedures described in this section.

While the `MPI_SOURCE`, `MPI_TAG`, and `MPI_ERROR` status values are directly accessible by the user, for convenience in some contexts, users can also access them via procedure calls, as described below.

`MPI_STATUS_GET_SOURCE(status, source)`

IN	status	status from which to retrieve source rank (status)
OUT	source	rank set in the <code>MPI_SOURCE</code> field (integer)

C binding

```
int MPI_Status_get_source(MPI_Status *status, int *source)
```

Fortran 2008 binding

```
MPI_Status_get_source(status, source, ierror)
  TYPE(MPI_Status), INTENT(IN) :: status
  INTEGER, INTENT(OUT) :: source
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_STATUS_GET_SOURCE(STATUS, SOURCE, IERROR)
  INTEGER STATUS(MPI_STATUS_SIZE), SOURCE, IERROR
```

Returns in `source` the value of the `MPI_SOURCE` field in the `status` object.

`MPI_STATUS_GET_TAG(status, tag)`

IN	status	status from which to retrieve tag (status)
OUT	tag	tag set in the <code>MPI_TAG</code> field (integer)

C binding

```
int MPI_Status_get_tag(MPI_Status *status, int *tag)
```

Fortran 2008 binding

```
MPI_Status_get_tag(status, tag, ierror)
  TYPE(MPI_Status), INTENT(IN) :: status
```

```

1     INTEGER, INTENT(OUT) :: tag
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

4 MPI_STATUS_GET_TAG(STATUS, TAG, IERROR)
5     INTEGER STATUS(MPI_STATUS_SIZE), TAG, IERROR

```

7 Returns in tag the value in the MPI_TAG field of the status object.

```

9 MPI_STATUS_GET_ERROR(status, err)
10

```

```

11     IN          status          status from which to retrieve error (status)

```

```

12     OUT         err             error set in the MPI_ERROR field (integer)

```

C binding

```

15 int MPI_Status_get_error(MPI_Status *status, int *err)

```

Fortran 2008 binding

```

18 MPI_Status_get_error(status, err, ierror)
19     TYPE(MPI_Status), INTENT(IN) :: status
20     INTEGER, INTENT(OUT) :: err
21     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

23 MPI_STATUS_GET_ERROR(STATUS, ERR, IERROR)
24     INTEGER STATUS(MPI_STATUS_SIZE), ERR, IERROR

```

25 Returns in err the value in the MPI_ERROR field of the status object.

26 Procedures for setting these fields in a status object are defined in Section 13.3.

3.2.6 Passing MPI_STATUS_IGNORE for Status

30 Every call to MPI_RECV includes a status argument, wherein the system can return details
31 about the message received. There are also a number of other MPI calls where status
32 is returned. An object of type MPI_Status is not an MPI opaque object; its structure is
33 declared in mpi.h and (deprecated) mpif.h, and it exists in the user's program. In many
34 cases, application programs are constructed so that it is unnecessary for them to examine
35 the status fields. In these cases, it is a waste for the user to allocate a status object, and it
36 is particularly wasteful for the MPI implementation to fill in fields in this object.

37 To cope with this problem, there are two predefined constants, MPI_STATUS_IGNORE
38 and MPI_STATUSES_IGNORE, which when passed to a receive, probe, wait, or test function,
39 inform the implementation that the status fields are not to be filled in. Note that
40 MPI_STATUS_IGNORE is not a special type of MPI_Status object; rather, it is a special value for
41 the argument. In C one would expect it to be NULL, not the address of a special MPI_Status.

42 MPI_STATUS_IGNORE, and the array version MPI_STATUSES_IGNORE, can be used every-
43 where a status argument is passed to a receive, wait, or test function. MPI_STATUS_IGNORE
44 cannot be used when status is an IN argument. Note that in Fortran MPI_STATUS_IGNORE
45 and MPI_STATUSES_IGNORE are objects like MPI_BOTTOM (not usable for initialization or
46 assignment), see Section 2.5.4.

47 In general, this optimization can apply to all functions for which status or an array of
48 statuses is an OUT argument. Note that this converts status into an INOUT argument. The

functions that can be passed `MPI_STATUS_IGNORE` are all the various forms of `MPI_RECV`, `MPI_PROBE`, `MPI_TEST`, and `MPI_WAIT`, as well as `MPI_REQUEST_GET_STATUS`. When an array is passed, as in the `MPI_{TEST|WAIT}{ALL|SOME}` functions, a separate constant, `MPI_STATUSES_IGNORE`, is passed for the array argument. It is possible for an MPI function to return `MPI_ERR_IN_STATUS` even when `MPI_STATUS_IGNORE` or `MPI_STATUSES_IGNORE` has been passed to that function.

`MPI_STATUS_IGNORE` and `MPI_STATUSES_IGNORE` are not required to have the same values in C and Fortran.

It is not allowed to have some of the statuses in an array of statuses for `MPI_{TEST|WAIT}{ALL|SOME}` functions set to `MPI_STATUS_IGNORE`; one either specifies ignoring *all* of the statuses in such a call with `MPI_STATUSES_IGNORE`, or *none* of them by passing normal statuses in all positions in the array of statuses.

3.2.7 Blocking Send-Receive

The **send-recv** operations combine in one operation the sending of a message to one destination and the receiving of another message, from another MPI process. The two (source and destination) are possibly the same. A send-recv operation is very useful for executing a shift operation across a chain of MPI processes. If blocking sends and receives are used for such a shift, then one needs to order the sends and receives correctly (for example, MPI processes with even rank in the communicator send, then receive, MPI processes with odd rank in the communicator receive first, then send) so as to prevent cyclic dependencies that may lead to *deadlock*. When a send-recv operation is used, the communication subsystem takes care of these issues. The send-recv operation can be used in conjunction with the procedures described in Chapter 8 in order to perform shifts on various logical topologies. Also, a send-recv operation is useful for implementing remote procedure calls.

A message sent by a send-recv operation can be received by a regular receive operation or probed by a probe operation; a send-recv operation can receive a message sent by a regular send operation.

`MPI_SENDRCV`(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount, recvtype, source, recvtag, comm, status)

IN	sendbuf	initial address of send buffer (choice)
IN	sendcount	number of elements in send buffer (non-negative integer)
IN	sendtype	type of elements in send buffer (handle)
IN	dest	rank of destination (integer)
IN	sendtag	send tag (integer)
OUT	recvbuf	initial address of receive buffer (choice)
IN	recvcount	number of elements in receive buffer (non-negative integer)
IN	recvtype	type of elements receive buffer element (handle)
IN	source	rank of source or <code>MPI_ANY_SOURCE</code> (integer)

```

1      IN      recvtag      receive tag or MPI_ANY_TAG (integer)
2      IN      comm        communicator (handle)
3
4      OUT     status       status object (status)
5

```

C binding

```

7      int MPI_Sendrecv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
8                      int dest, int sendtag, void *recvbuf, int recvcount,
9                      MPI_Datatype recvtype, int source, int recvtag, MPI_Comm comm,
10                     MPI_Status *status)

```

```

11     int MPI_Sendrecv_c(const void *sendbuf, MPI_Count sendcount,
12                       MPI_Datatype sendtype, int dest, int sendtag, void *recvbuf,
13                       MPI_Count recvcount, MPI_Datatype recvtype, int source,
14                       int recvtag, MPI_Comm comm, MPI_Status *status)
15

```

Fortran 2008 binding

```

16     MPI_Sendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
17                 recvtype, source, recvtag, comm, status, ierror)
18     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
19     INTEGER, INTENT(IN) :: sendcount, dest, sendtag, recvcount, source, recvtag
20     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
21     TYPE(*), DIMENSION(..) :: recvbuf
22     TYPE(MPI_Comm), INTENT(IN) :: comm
23     TYPE(MPI_Status) :: status
24     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
25
26     MPI_Sendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
27                 recvtype, source, recvtag, comm, status, ierror) !(_c)
28     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
29     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
30     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
31     INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag
32     TYPE(*), DIMENSION(..) :: recvbuf
33     TYPE(MPI_Comm), INTENT(IN) :: comm
34     TYPE(MPI_Status) :: status
35     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36

```

Fortran binding

```

37     MPI_SENDRCV(SENDBUF, SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVBUF, RECVCOUNT,
38               RECVTYPE, SOURCE, RECVTAG, COMM, STATUS, IERROR)
39     <type> SENDBUF(*), RECVBUF(*)
40     INTEGER SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVCOUNT, RECVTYPE, SOURCE,
41               RECVTAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
42

```

43 Execute a blocking send-receive operation. Both send and receive use the same com-
44 municator, but possibly different tags. The send buffer and receive buffers must be disjoint,
45 and may have different lengths and datatypes.

46 The semantics of a send-receive operation is what would be obtained if the caller forked
47 two concurrent threads, one to execute the send, and one to execute the receive, followed
48 by a join of these two threads.

MPI_SENDRECV_REPLACE(buf, count, datatype, dest, sendtag, source, recvtag, comm,		1	
status)		2	
INOUT	buf	initial address of send and receive buffer (choice)	3
IN	count	number of elements in send and receive buffer	4
		(non-negative integer)	5
IN	datatype	type of elements in send and receive buffer (handle)	6
IN	dest	rank of destination (integer)	7
IN	sendtag	send message tag (integer)	8
IN	source	rank of source or MPI_ANY_SOURCE (integer)	9
IN	recvtag	receive message tag or MPI_ANY_TAG (integer)	10
IN	comm	communicator (handle)	11
OUT	status	status object (status)	12

C binding

```
int MPI_Sendrecv_replace(void *buf, int count, MPI_Datatype datatype, int dest,
                        int sendtag, int source, int recvtag, MPI_Comm comm,
                        MPI_Status *status)
```

```
int MPI_Sendrecv_replace_c(void *buf, MPI_Count count, MPI_Datatype datatype,
                          int dest, int sendtag, int source, int recvtag, MPI_Comm comm,
                          MPI_Status *status)
```

Fortran 2008 binding

```
MPI_Sendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
                    comm, status, ierror)
```

```
TYPE(*), DIMENSION(..) :: buf
```

```
INTEGER, INTENT(IN) :: count, dest, sendtag, source, recvtag
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Status) :: status
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Sendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
                    comm, status, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..) :: buf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Status) :: status
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SENDRECV_REPLACE(BUF, COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG,
                    COMM, STATUS, IERROR)
```

```
<type> BUF(*)
```

1 INTEGER COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG, COMM,
2 STATUS(MPI_STATUS_SIZE), IERROR

3 Execute a blocking send and receive. The same buffer is used both for the send and
4 for the receive, so that the message sent is replaced by the message received.
5

6 *Advice to implementors.* Additional intermediate buffering is needed for the “replace”
7 variant. (*End of advice to implementors.*)
8

10 3.3 Datatype Matching and Data Conversion

11 3.3.1 Type Matching Rules

12 One can think of message transfer as consisting of the following three phases.
13

- 14 1. Data is pulled out of the send buffer and a message is assembled.
- 15 2. A message is transferred from sender to receiver.
- 16 3. Data is pulled from the incoming message and disassembled into the receive buffer.
17

18 **Type matching** has to be observed at each of these three phases: The type of each
19 variable in the sender buffer has to match the type specified for that entry by the send
20 operation; the type specified by the send operation has to match the type specified by the
21 receive operation; and the type of each variable in the receive buffer has to match the type
22 specified for that entry by the receive operation. A program that fails to observe these three
23 rules is *erroneous*.
24

25 To define type matching more precisely, we need to deal with two issues: matching of
26 types of the host language with types specified in communication operations; and matching
27 of types at sender and receiver.
28

29 The types of a send and receive match (phase two) if both operations use identical
30 names. That is, MPI_INTEGER matches MPI_INTEGER, MPI_REAL matches MPI_REAL, and
31 so on. There is one exception to this rule, discussed in Section 5.2: the type MPI_PACKED
32 can match any other type.

33 The type of a variable in a host program matches the type specified in the commu-
34 nication operation if the datatype name used by that operation corresponds to the basic
35 type of the host program variable. For example, an entry with type name MPI_INTEGER
36 matches a Fortran variable of type INTEGER. A table giving this correspondence for Fortran
37 and C appears in Section 3.2.2. There are two exceptions to this last rule: an entry with
38 type name MPI_BYTE or MPI_PACKED can be used to match any byte of storage (on a byte-
39 addressable machine), irrespective of the datatype of the variable that contains this byte.
40 The type MPI_PACKED is used to send data that has been explicitly packed, or receive data
41 that will be explicitly unpacked, see Section 5.2. The type MPI_BYTE allows one to transfer
42 the binary value of a byte in memory unchanged.

43 To summarize, the type matching rules fall into the three categories below.

- 44 • Communication of typed values (e.g., with datatype different from MPI_BYTE), where
45 the datatypes of the corresponding entries in the sender program, in the send call, in
46 the receive call and in the receiver program must all match.
47

- Communication of untyped values (e.g., of datatype `MPI_BYTE`), where both sender and receiver use the datatype `MPI_BYTE`. In this case, there are no requirements on the types of the corresponding entries in the sender and the receiver programs, nor is it required that they be the same.
- Communication involving packed data, where `MPI_PACKED` is used.

The following examples illustrate the first two cases.

Example 3.2. Sender and receiver specify matching types.

```
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SEND(a(1), 10, MPI_REAL, 1, tag, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(b(1), 15, MPI_REAL, 0, tag, comm, status, ierr)
END IF
```

This code is correct if both `a` and `b` are real arrays of size ≥ 10 . (In Fortran, it might be correct to use this code even if `a` or `b` have size < 10 : e.g., when `a(1)` can be equivalenced to an array with ten reals.)

Example 3.3. Sender and receiver do not specify matching types.

```
! ----- THIS EXAMPLE IS ERRONEOUS -----
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SEND(a(1), 10, MPI_REAL, 1, tag, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(b(1), 40, MPI_BYTE, 0, tag, comm, status, ierr)
END IF
```

This code is *erroneous*, since sender and receiver do not provide matching datatype arguments.

Example 3.4. Sender and receiver specify communication of untyped values.

```
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SEND(a(1), 40, MPI_BYTE, 1, tag, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(b(1), 60, MPI_BYTE, 0, tag, comm, status, ierr)
END IF
```

This code is correct, irrespective of the type and size of `a` and `b` (unless this results in an out of bounds memory access).

Advice to users. If a buffer of type `MPI_BYTE` is passed as an argument to `MPI_SEND`, then MPI will send the data stored at contiguous locations, starting from the address indicated by the `buf` argument. This may have unexpected results when the data layout is not as a casual user would expect it to be. For example, some Fortran compilers implement variables of type `CHARACTER` as a structure that contains the character length and a pointer to the actual string. In such an environment, sending and receiving a Fortran `CHARACTER` variable using the `MPI_BYTE` type will not have

the anticipated result of transferring the character string. For this reason, the user is advised to use typed communication operations whenever possible. (*End of advice to users.*)

Type MPI_CHARACTER

The type MPI_CHARACTER matches one character of a Fortran variable of type CHARACTER, rather than the entire character string stored in the variable. Fortran variables of type CHARACTER or substrings are transferred as if they were arrays of characters. This is illustrated in the example below.

Example 3.5. Transfer of Fortran CHARACTERS.

```

CHARACTER*10 a
CHARACTER*10 b

CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SEND(a, 5, MPI_CHARACTER, 1, tag, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(b(6:10), 5, MPI_CHARACTER, 0, tag, comm, status, ierr)
END IF

```

The last five characters of string b at the MPI process with rank = 1 are replaced by the first five characters of string a at the MPI process with rank = 0.

Rationale. The alternative choice would be for MPI_CHARACTER to match a character of arbitrary length. This runs into problems.

A Fortran character variable is a constant length string, with no special termination symbol. There is no fixed convention on how to represent characters, and how to store their length. Some compilers pass a character argument to a routine as a pair of arguments, one holding the address of the string and the other holding the length of string. Consider the case of an MPI communication call that is passed a communication buffer with type defined by a derived datatype (Section 5.1). If this communicator buffer contains variables of type CHARACTER then the information on their length will not be passed to the MPI routine.

This problem forces us to provide explicit information on character length with the MPI call. One could add a length parameter to the type MPI_CHARACTER, but this does not add much convenience and the same functionality can be achieved by defining a suitable derived datatype. (*End of rationale.*)

Advice to implementors. Some compilers pass Fortran CHARACTER arguments as a structure with a length and a pointer to the actual string. In such an environment, the MPI call needs to dereference the pointer in order to reach the string. (*End of advice to implementors.*)

3.3.2 Data Conversion

One of the goals of MPI is to support parallel computations across heterogeneous environments. Communication in a heterogeneous environment may require data conversions. We use the following terminology.

type conversion changes the datatype of a value, e.g., by rounding a REAL to an INTEGER. 1

representation conversion changes the binary representation of a value, e.g., from Hex floating point to IEEE floating point. 2
3
4

The type matching rules imply that MPI communication never entails type conversion. On the other hand, MPI requires that a representation conversion be performed when a typed value is transferred across environments that use different representations for the datatype of this value. MPI does not specify rules for representation conversion. Such conversion is expected to preserve integer, logical and character values, and to convert a floating point value to the nearest value that can be represented on the target system. 5
6
7
8
9
10

Overflow and underflow exceptions may occur during floating point conversions. Conversion of integers or characters may also lead to exceptions when a value that can be represented in one system cannot be represented in the other system. An exception occurring during representation conversion results in a failure of the communication. An error occurs either in the send operation, or the receive operation, or both. 11
12
13
14
15

If a value sent in a message is untyped (i.e., of type MPI_BYTE), then the binary representation of the byte stored at the receiver is identical to the binary representation of the byte loaded at the sender. This holds true, whether sender and receiver run in the same or in distinct environments. No representation conversion is required. (Note that representation conversion may occur when values of type MPI_CHARACTER or MPI_CHAR are transferred, for example, from an EBCDIC encoding to an ASCII encoding.) 16
17
18
19
20
21

No conversion need occur when an MPI program executes in a homogeneous system, where all MPI processes run in the same environment. 22
23

Consider the three examples, 3.2–3.4. The first program is correct, assuming that *a* and *b* are REAL arrays of size ≥ 10 . If the sender and receiver execute in different environments, then the ten real values that are fetched from the send buffer will be converted to the representation for reals on the receiver site before they are stored in the receive buffer. While the number of real elements fetched from the send buffer equal the number of real elements stored in the receive buffer, the number of bytes stored need not equal the number of bytes loaded. For example, the sender may use a four byte representation and the receiver an eight byte representation for reals. 24
25
26
27
28
29
30
31

The second program is *erroneous*, and its behavior is undefined. 32

The third program is correct. The exact same sequence of forty bytes that were loaded from the send buffer will be stored in the receive buffer, even if sender and receiver run in a different environment. The message sent has exactly the same length (in bytes) and the same binary representation as the message received. If *a* and *b* are of different types, or if they are of the same type but different data representations are used, then the bits stored in the receive buffer may encode values that are different from the values they encoded in the send buffer. 33
34
35
36
37
38
39

Data representation conversion also applies to the *envelope* of a message: source, destination and tag are all integers that may need to be converted. 40
41

Advice to implementors. The current definition does not require messages to carry data type information. Both sender and receiver provide complete data type information. In a heterogeneous environment, one can either use a machine independent encoding such as XDR, or have the receiver convert from the sender representation to its own, or even have the sender do the conversion. 42
43
44
45
46
47

1 Additional type information might be added to messages in order to allow the sys-
2 tem to detect mismatches between datatype at sender and receiver. This might be
3 particularly useful in a slower but safer debug mode. (*End of advice to implementors.*)
4

5 MPI requires support for inter-language communication, e.g., if messages are sent using
6 an MPI procedure from the MPI C language interface and received using an MPI procedure
7 from one of the MPI Fortran language interfaces. The behavior is defined in Section 19.3.
8

9 3.4 Communication Modes

11 The send call described in Section 3.2.1 is *blocking*: it does not return until the *message*
12 *data* and *envelope* have been safely stored away so that the sender is free to modify the
13 send buffer. The message might be copied directly into the matching receive buffer, or it
14 might be copied into a temporary system buffer.

15 Message buffering decouples the send and receive operations. A blocking send can com-
16 plete as soon as the message was buffered, even if no matching receive has been executed by
17 the receiver. On the other hand, message buffering can be expensive, as it entails additional
18 memory-to-memory copying, and it requires the allocation of memory for buffering. MPI
19 offers the choice of several **communication modes** that allow one to control the choice of
20 the communication protocol.

21 The send call described in Section 3.2.1 uses the **standard** communication mode. In
22 this mode, it is up to MPI to decide whether outgoing messages will be buffered. MPI may
23 buffer outgoing messages. In such a case, the send call may complete before a matching
24 receive is invoked. On the other hand, buffer space may be unavailable, or MPI may choose
25 not to buffer outgoing messages, for performance reasons. In this case, the send call will
26 not complete until a matching receive has been *started*, and the data has been moved to
27 the receiver.

28 Thus, a *standard mode send* can be *started* whether or not a matching receive has been
29 *started*. It may *complete* before a matching receive is *started*. The standard mode send is
30 *nonlocal*: successful completion of the send operation may depend on the occurrence of a
31 matching receive.
32

33 *Rationale.* The reluctance of MPI to mandate whether standard sends are buffering
34 or not stems from the desire to achieve portable programs. Since any system will run
35 out of buffer resources as message sizes are increased, and some implementations may
36 want to provide little buffering, MPI takes the position that correct (and therefore,
37 portable) programs do not rely on system buffering in standard mode. Buffering may
38 improve the performance of a correct program, but it doesn't affect the result of the
39 program. If the user wishes to guarantee a certain amount of buffering, the user-
40 provided buffer system of Section 3.6 should be used, along with the buffered-mode
41 send. (*End of rationale.*)
42

43 There are three additional communication modes.

44 A **buffered** mode send operation can be started whether or not a matching receive
45 has been *started*. It may complete before a matching receive is *started*. However, unlike the
46 standard send, this operation is *local*, and its completion does not depend on the occurrence
47 of a matching receive. Thus, if a send is executed and no matching receive is *started*, then
48 MPI must buffer the outgoing message, so as to allow the send call to complete. An error will

occur if there is insufficient buffer space. The amount of available buffer space is controlled by the user—see Section 3.6. Buffer allocation by the user may be required for the buffered mode to be effective.

A send that uses the **synchronous** mode can be started whether or not a matching receive was *started*. However, the send will complete successfully only if a matching receive is *started*, and the receive operation has started to receive the message sent by the synchronous send. Thus, the completion of a synchronous send not only indicates that the send buffer can be reused, but it also indicates that the receiver has reached a certain point in its execution, namely that it has started executing the matching receive. If both sends and receives are blocking operations then the use of the synchronous mode provides synchronous communication semantics: a communication does not complete at either end before both MPI processes rendezvous at the communication. A send executed in this mode is *nonlocal*.

A send that uses the **ready** communication mode may be started *only* if the matching receive is already *started*. Otherwise, the operation is *erroneous* and its outcome is undefined. On some systems, this allows the removal of a hand-shake protocol that is otherwise required and results in improved performance. The completion of the send operation does not depend on the status of a matching receive, and merely indicates that the send buffer can be reused. A send operation that uses the ready mode has the same semantics as a standard send operation, or a synchronous send operation; it is merely that the sender provides additional information to the system (namely that a matching receive is already *started*), that can save some overhead. In a correct program, therefore, a ready send could be replaced by a standard send with no effect on the behavior of the program other than performance.

Three additional send functions are provided for the three additional communication modes. The communication mode is indicated by a one letter prefix: B for buffered, S for synchronous, and R for ready.

`MPI_BSEND(buf, count, datatype, dest, tag, comm)`

IN	<code>buf</code>	initial address of send buffer (choice)
IN	<code>count</code>	number of elements in send buffer (non-negative integer)
IN	<code>datatype</code>	datatype of each send buffer element (handle)
IN	<code>dest</code>	rank of destination (integer)
IN	<code>tag</code>	message tag (integer)
IN	<code>comm</code>	communicator (handle)

C binding

```
int MPI_Bsend(const void *buf, int count, MPI_Datatype datatype, int dest,
             int tag, MPI_Comm comm)
```

```
int MPI_Bsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
               int dest, int tag, MPI_Comm comm)
```

Fortran 2008 binding

```
MPI_Bsend(buf, count, datatype, dest, tag, comm, ierror)
      TYPE(*), DIMENSION(..), INTENT(IN) :: buf
```

```

1     INTEGER, INTENT(IN) :: count, dest, tag
2     TYPE(MPI_Datatype), INTENT(IN) :: datatype
3     TYPE(MPI_Comm), INTENT(IN) :: comm
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6 MPI_Bsend(buf, count, datatype, dest, tag, comm, ierror) !(_c)
7     TYPE(*), DIMENSION(..), INTENT(IN) :: buf
8     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
9     TYPE(MPI_Datatype), INTENT(IN) :: datatype
10    INTEGER, INTENT(IN) :: dest, tag
11    TYPE(MPI_Comm), INTENT(IN) :: comm
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

14 MPI_BSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
15 <type> BUF(*)
16 INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR

```

Send in buffered mode.

According to the definitions in Section 2.4.2, MPI_BSEND is a completing procedure and the user can re-use all resources given as arguments, including the *message data buffer*. It is also a local procedure because it returns immediately without depending on the execution of any MPI procedure in any other MPI process.

Advice to users. This is one of the exceptions in which a completing and therefore blocking operation-related procedure is local. (*End of advice to users.*)

```

27 MPI_SSEND(buf, count, datatype, dest, tag, comm)

```

29	IN	buf	initial address of send buffer (choice)
30	IN	count	number of elements in send buffer (non-negative integer)
32	IN	datatype	datatype of each send buffer element (handle)
34	IN	dest	rank of destination (integer)
35	IN	tag	message tag (integer)
36	IN	comm	communicator (handle)

C binding

```

40 int MPI_Ssend(const void *buf, int count, MPI_Datatype datatype, int dest,
41             int tag, MPI_Comm comm)

```

```

42 int MPI_Ssend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
43               int dest, int tag, MPI_Comm comm)

```

Fortran 2008 binding

```

45 MPI_Ssend(buf, count, datatype, dest, tag, comm, ierror)
46     TYPE(*), DIMENSION(..), INTENT(IN) :: buf
47     INTEGER, INTENT(IN) :: count, dest, tag

```

TYPE(MPI_Datatype), INTENT(IN) :: datatype	1
TYPE(MPI_Comm), INTENT(IN) :: comm	2
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	3
MPI_Ssend(buf, count, datatype, dest, tag, comm, ierror) !(_c)	4
TYPE(*), DIMENSION(..), INTENT(IN) :: buf	5
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count	6
TYPE(MPI_Datatype), INTENT(IN) :: datatype	7
INTEGER, INTENT(IN) :: dest, tag	8
TYPE(MPI_Comm), INTENT(IN) :: comm	9
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	10
Fortran binding	11
MPI_SSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)	12
<type> BUF(*)	13
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR	14
Send in synchronous mode.	15
MPI_RSEND(buf, count, datatype, dest, tag, comm)	16
IN buf initial address of send buffer (choice)	17
IN count number of elements in send buffer (non-negative integer)	18
IN datatype datatype of each send buffer element (handle)	19
IN dest rank of destination (integer)	20
IN tag message tag (integer)	21
IN comm communicator (handle)	22
C binding	23
int MPI_Rsend(const void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm)	24
int MPI_Rsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm)	25
Fortran 2008 binding	26
MPI_Rsend(buf, count, datatype, dest, tag, comm, ierror)	27
TYPE(*), DIMENSION(..), INTENT(IN) :: buf	28
INTEGER, INTENT(IN) :: count, dest, tag	29
TYPE(MPI_Datatype), INTENT(IN) :: datatype	30
TYPE(MPI_Comm), INTENT(IN) :: comm	31
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	32
MPI_Rsend(buf, count, datatype, dest, tag, comm, ierror) !(_c)	33
TYPE(*), DIMENSION(..), INTENT(IN) :: buf	34
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count	35
TYPE(MPI_Datatype), INTENT(IN) :: datatype	36
INTEGER, INTENT(IN) :: dest, tag	37

```

1     TYPE(MPI_Comm), INTENT(IN) :: comm
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

4 MPI_RSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
5     <type> BUF(*)
6     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR

```

8 Send in ready mode.

9 There is only one receive operation, but it matches any of the send modes. The receive
10 procedure described in the last section is *blocking*: it returns only after the receive buffer
11 contains the newly received message. A receive can complete before the matching send has
12 completed (of course, it can complete only after the matching send has started).

13 In a multithreaded implementation of MPI, the system may de-schedule a thread that
14 is blocked on a send or receive operation, and schedule another thread for execution in
15 the same address space. In such a case it is the user's responsibility not to modify a
16 communication buffer until the communication completes. Otherwise, the outcome of the
17 computation is undefined.

18
19 *Advice to implementors.* Since a synchronous send cannot complete before a matching
20 receive is *started*, one will not normally buffer messages sent by such an operation.

21 It is recommended to choose buffering over blocking the sender, whenever possible,
22 for standard sends. The programmer can signal a preference for blocking the sender
23 until a matching receive occurs by using the synchronous send mode.

24 A possible communication protocol for the various communication modes is outlined
25 below.

26
27 **ready send:** The message is sent as soon as possible.

28 **synchronous send:** The sender sends a request-to-send message. The receiver stores
29 this request. When a matching receive is *started*, the receiver sends back a
30 permission-to-send message, and the sender now sends the message.

31
32 **standard send:** First protocol may be used for short messages, and second protocol
33 for long messages.

34 **buffered send:** The sender copies the message into a buffer and then sends it with
35 a nonblocking send (using the same protocol as for standard send).

36 Additional control messages might be needed for flow control and error recovery. Of
37 course, there are many other possible protocols.

38
39 Ready send can be implemented as a standard send. In this case there will be no
40 performance advantage (or disadvantage) for the use of ready send.

41 A standard send can be implemented as a synchronous send. In such a case, no data
42 buffering is needed. However, users may expect some buffering.

43 In a multithreaded environment, the execution of a blocking communication should
44 block only the executing thread, allowing the thread scheduler to de-schedule this
45 thread and schedule another thread for execution. (*End of advice to implementors.*)
46
47
48

3.5 Semantics of Point-to-Point Communication

A valid MPI implementation guarantees certain general properties of point-to-point communication, which are described in this section.

Order. Messages are **nonovertaking**: If a sender sends two messages in succession to the same destination, and both match the same receive, then this operation cannot receive the second message if the first one is still pending. If a receiver posts two receives in succession, and both match the same message, then the second receive operation cannot be satisfied by this message, if the first one is still pending. This requirement facilitates matching of sends to receives. It guarantees that message-passing code is deterministic, if MPI processes are single-threaded and the wildcard `MPI_ANY_SOURCE` is not used in receives. (Some of the calls described later, such as `MPI_CANCEL` or `MPI_WAITANY`, are additional sources of nondeterminism.)

If an MPI process has a single thread of execution, then any two communication operations executed by this MPI process are **ordered**.

Advice to users. The MPI Forum believes the following paragraph is ambiguous and may clarify the meaning in a future version of the MPI Standard. (*End of advice to users.*)

On the other hand, if the MPI process is multithreaded, then the semantics of thread execution may not define a relative order between two send operations executed by two distinct threads. The operations are **logically concurrent**, even if one physically precedes the other. In such a case, the two messages sent can be received in any order. Similarly, if two receive operations that are **logically concurrent** receive two successively sent messages, then the two messages can match the two receives in either order.

Advice to implementors. The MPI Forum believes the previous paragraph is ambiguous and may clarify the meaning in a future version of the MPI Standard. (*End of advice to implementors.*)

Example 3.6. An example of nonovertaking messages.

```
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_BSEND(buf1, count, MPI_REAL, 1, tag, comm, ierr)
  CALL MPI_BSEND(buf2, count, MPI_REAL, 1, tag, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(buf1, count, MPI_REAL, 0, MPI_ANY_TAG, comm, status, &
               ierr)
  CALL MPI_RECV(buf2, count, MPI_REAL, 0, tag, comm, status, ierr)
END IF
```

The message sent by the first send must be received by the first receive, and the message sent by the second send must be received by the second receive.

Progress. If a pair of matching send and receive operations have been initiated, then at least one of these two operations will complete, independently of other actions in the system: the send operation will complete, unless the receive is satisfied by another message,

and completes; the receive operation will complete, unless the message sent is consumed by another matching receive that was *started* at the same destination MPI process.

Example 3.7. An example of two, intertwined matching pairs.

```

CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_BSEND(buf1, count, MPI_REAL, 1, tag1, comm, ierr)
  CALL MPI_SSEND(buf2, count, MPI_REAL, 1, tag2, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(buf1, count, MPI_REAL, 0, tag2, comm, status, ierr)
  CALL MPI_RECV(buf2, count, MPI_REAL, 0, tag1, comm, status, ierr)
END IF

```

Both MPI processes invoke their first communication call. Since the first send at the MPI process with `rank = 0` uses the buffered mode, it must complete, irrespective of the state of the other MPI process(es). Since no matching receive is *started*, the message will be copied into buffer space. (If insufficient buffer space is available, then the program will fail.) The second send is then invoked. At that point, a matching pair of send and receive operation is enabled, and both operations must complete. Next, the second receive call is invoked, which will be satisfied by the buffered message. Note that the MPI process with `rank = 1` received the messages in the reverse order they were sent.

Fairness. MPI makes no guarantee of **fairness** in the handling of communication. Suppose that a send is *started*. Then it is possible that the destination MPI process repeatedly posts a receive that matches this send, yet the message is never received, because it is each time overtaken by another message, sent from another source. Similarly, suppose that a receive was *started* by a multithreaded MPI process. Then it is possible that messages that match this receive are repeatedly received, yet the receive is never satisfied, because it is overtaken by other receives *started* at this MPI process (by other executing threads). It is the programmer's responsibility to prevent starvation in such situations.

Resource limitations. Any *pending* communication operation and *decoupled MPI activity* consumes system resources that are limited. Errors may occur when lack of resources prevent the execution of an MPI call. High-quality implementations will use a (small) fixed amount of resources for each *pending* send in the ready or synchronous mode and for each *pending* receive. However, buffer space may be consumed to store messages sent in standard mode, and must be consumed to store messages sent in buffered mode, when no matching receive is available. The amount of space available for buffering will be much smaller than program data memory on many systems. Then, it will be easy to write programs that overrun available buffer space.

MPI allows the user to provide buffer memory for messages sent in the buffered mode. Furthermore, MPI specifies a detailed operational model for the use of this buffer. An MPI implementation is required to do no worse than implied by this model. This allows users to avoid buffer overflows when they use buffered sends. Buffer allocation and use is described in Section 3.6.

A buffered send operation that cannot complete because of a lack of buffer space is *erroneous*. When such a situation is detected, an error is signaled that may cause the program to terminate abnormally. On the other hand, a standard send operation that

cannot complete because of lack of buffer space will merely block, waiting for buffer space to become available or for a matching receive to be *started*. This behavior is preferable in many situations. Consider a situation where a producer repeatedly produces new values and sends them to a consumer. Assume that the producer produces new values faster than the consumer can consume them. If buffered sends are used, then a buffer overflow will result. Additional synchronization has to be added to the program so as to prevent this from occurring. If standard sends are used, then the producer will be automatically throttled, as its send operations will block when buffer space is unavailable.

In some situations, a lack of buffer space leads to deadlock situations. This is illustrated by the examples below.

Example 3.8. An exchange of messages.

```
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SEND(sendbuf, count, MPI_REAL, 1, tag, comm, ierr)
  CALL MPI_RECV(recvbuf, count, MPI_REAL, 1, tag, comm, status, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(recvbuf, count, MPI_REAL, 0, tag, comm, status, ierr)
  CALL MPI_SEND(sendbuf, count, MPI_REAL, 0, tag, comm, ierr)
END IF
```

This program will succeed even if no buffer space for data is available. The standard send operation can be replaced, in this example, with a synchronous send.

Example 3.9. An errant attempt to exchange messages.

```
! ----- THIS EXAMPLE IS ERRONEOUS -----
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_RECV(recvbuf, count, MPI_REAL, 1, tag, comm, status, ierr)
  CALL MPI_SEND(sendbuf, count, MPI_REAL, 1, tag, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(recvbuf, count, MPI_REAL, 0, tag, comm, status, ierr)
  CALL MPI_SEND(sendbuf, count, MPI_REAL, 0, tag, comm, ierr)
END IF
```

The receive operation of the MPI process with rank = 0 must complete before its send, and can complete only if the matching send of the MPI process with rank = 1 is executed. The receive operation of the MPI process with rank = 1 must complete before its send and can complete only if the matching send of the MPI process with rank = 0 is executed. This program will always deadlock. The same holds for any other send mode.

Example 3.10. An unsafe exchange that relies on MPI to provide sufficient buffering.

```
! ----- THIS EXAMPLE IS ERRONEOUS -----
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SEND(sendbuf, count, MPI_REAL, 1, tag, comm, ierr)
  CALL MPI_RECV(recvbuf, count, MPI_REAL, 1, tag, comm, status, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_SEND(sendbuf, count, MPI_REAL, 0, tag, comm, ierr)
```

```

1  CALL MPI_RECV(recvbuf, count, MPI_REAL, 0, tag, comm, status, ierr)
2  END IF

```

The message sent by each MPI process has to be copied out before the send operation completes and the receive operation starts. For the program to complete, it is necessary that at least one of the two messages sent be buffered. Thus, this program can succeed only if the communication system can buffer at least `count` words of data.

Advice to users. If standard mode send operations are used as in Example 3.10, then a deadlock situation may occur where both MPI processes are blocked because sufficient buffer space is not available. The same will certainly happen, if the synchronous mode is used. If the buffered mode is used, and not enough buffer space is available, then the program will not complete either. However, rather than a deadlock situation, we shall have a buffer overflow error.

A portable program using standard mode send operations should not rely on message buffering for the program to complete without *deadlock*. All sends in such a portable program can be replaced with synchronous mode sends and the program will still run correctly. The buffered send mode can be used for programs that require buffering.

Nonblocking message-passing operations, as described in Section 3.7, can be used to avoid the need for buffering outgoing messages. This can prevent unintentional *serialization* or *deadlock* due to lack of buffer space, and improves performance, by allowing *overlap* of communication with other communication or with computation, and avoiding the overheads of allocating buffers and copying messages into buffers. (*End of advice to users.*)

3.6 Buffer Allocation and Usage

A user may specify up to one buffer per communicator, up to one buffer per session, and up to one buffer per MPI process to be used for buffering messages sent in buffered mode. Buffering is done by the sender.

```
MPI_COMM_ATTACH_BUFFER(comm, buffer, size)
```

IN	comm	communicator (handle)
IN	buffer	initial buffer address (choice)
IN	size	buffer size, in bytes (non-negative integer)

C binding

```
int MPI_Comm_attach_buffer(MPI_Comm comm, void *buffer, int size)
```

```
int MPI_Comm_attach_buffer_c(MPI_Comm comm, void *buffer, MPI_Count size)
```

Fortran 2008 binding

```
MPI_Comm_attach_buffer(comm, buffer, size, ierror)
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
```

```
INTEGER, INTENT(IN) :: size
```

```

    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_attach_buffer(comm, buffer, size, ierror) !(_c)
    TYPE(MPI_Comm), INTENT(IN) :: comm
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: size
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_COMM_ATTACH_BUFFER(COMM, BUFFER, SIZE, IERROR)
    INTEGER COMM, SIZE, IERROR
    <type> BUFFER(*)

```

Provides to MPI a communicator-specific buffer in memory. This is to be used for buffering outgoing messages sent when a buffered mode send is started that uses the communicator `comm`.

If `MPI_BUFFER_AUTOMATIC` is passed as the argument `buffer`, no explicit buffer is attached; rather, automatic buffering is enabled for all buffered mode communication associated with the communicator `comm` (see Section 3.6). Further, if `MPI_BUFFER_AUTOMATIC` is passed as the argument `buffer`, the value of `size` is irrelevant. Note that in Fortran `MPI_BUFFER_AUTOMATIC` is an object like `MPI_BOTTOM` (not usable for initialization or assignment), see Section 2.5.4.

```

MPI_SESSION_ATTACH_BUFFER(session, buffer, size)
    IN      session      session (handle)
    IN      buffer       initial buffer address (choice)
    IN      size         buffer size, in bytes (non-negative integer)

```

C binding

```

int MPI_Session_attach_buffer(MPI_Session session, void *buffer, int size)
int MPI_Session_attach_buffer_c(MPI_Session session, void *buffer,
    MPI_Count size)

```

Fortran 2008 binding

```

MPI_Session_attach_buffer(session, buffer, size, ierror)
    TYPE(MPI_Session), INTENT(IN) :: session
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
    INTEGER, INTENT(IN) :: size
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Session_attach_buffer(session, buffer, size, ierror) !(_c)
    TYPE(MPI_Session), INTENT(IN) :: session
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: size
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_SESSION_ATTACH_BUFFER(SESSION, BUFFER, SIZE, IERROR)
    INTEGER SESSION, SIZE, IERROR

```

1 <type> BUFFER(*)

2 Provides to MPI a session-specific buffer in memory. This buffer is to be used for
3 buffering outgoing messages sent when using a communicator that is created from a group
4 that is derived from the session `session`. However, if there is a communicator-specific buffer
5 attached to the particular communicator at the time of the buffered mode send is started,
6 that buffer is used.

7 If `MPI_BUFFER_AUTOMATIC` is passed as the argument `buffer`, no explicit buffer is at-
8 tached; rather, automatic buffering is enabled for all buffered mode communication associ-
9 ated with the session `session` that is not explicitly covered by a buffer provided at communi-
10 cator level (see Section 3.6). Further, if `MPI_BUFFER_AUTOMATIC` is passed as the argument
11 `buffer`, the value of `size` is irrelevant. Note that in Fortran `MPI_BUFFER_AUTOMATIC` is an
12 object like `MPI_BOTTOM` (not usable for initialization or assignment), see Section 2.5.4.
13

14
15 `MPI_BUFFER_ATTACH(buffer, size)`

16 IN `buffer` initial buffer address (choice)
17
18 IN `size` buffer size, in bytes (non-negative integer)

20 C binding

21 `int MPI_Buffer_attach(void *buffer, int size)`

22 `int MPI_Buffer_attach_c(void *buffer, MPI_Count size)`

24 Fortran 2008 binding

25 `MPI_Buffer_attach(buffer, size, ierror)`
26 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: `buffer`
27 INTEGER, INTENT(IN) :: `size`
28 INTEGER, OPTIONAL, INTENT(OUT) :: `ierror`

29 `MPI_Buffer_attach(buffer, size, ierror) !(_c)`
30 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: `buffer`
31 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: `size`
32 INTEGER, OPTIONAL, INTENT(OUT) :: `ierror`
33

34 Fortran binding

35 `MPI_BUFFER_ATTACH(BUFFER, SIZE, IERROR)`
36 <type> BUFFER(*)
37 INTEGER SIZE, IERROR

38 Provides to MPI an MPI process-specific buffer in memory. This buffer is to be used for
39 buffering outgoing messages sent when using a communicator to which no communicator-
40 specific buffer is attached or which is derived from a session to which no session-specific
41 buffer is attached at the time the buffered mode send is started.

42 If `MPI_BUFFER_AUTOMATIC` is passed as the argument `buffer`, no explicit buffer is at-
43 tached; rather, automatic buffering is enabled for all buffered mode communication not
44 explicitly covered by a buffer provided at session or communicator level (see Section 3.6).
45 Further, if `MPI_BUFFER_AUTOMATIC` is passed as the argument `buffer`, the value of `size` is
46 irrelevant. Note that in Fortran `MPI_BUFFER_AUTOMATIC` is an object like `MPI_BOTTOM`
47 (not usable for initialization or assignment), see Section 2.5.4.
48

Advice to users. The use of a global shared buffer can be problematic when used for communication in different libraries, as the buffer represents a shared resource used for all buffered mode communication. Further, with the introduction of the Sessions Model, the use of a single shared buffer violates the concept of resource isolation that is intended with MPI Sessions. It is therefore recommended, especially for libraries and programs using the Sessions Model, to use only communicator-specific or session-specific buffers. (*End of advice to users.*)

Any of these buffers are used only for messages sent in buffered mode. Only one MPI process-specific buffer can be attached to an MPI process at a time, only one session-specific buffer can be attached to a session at a time and only one communicator-specific buffer can be attached to a communicator at a time.

If automatic buffering is enabled at any level, no other buffer can be attached at that level.

A particular memory region can only be used in one buffer; reusing buffer space for multiple sessions, communicators and/or the global buffer is erroneous. Further, only one buffer is used for any one communication following the rules above; buffer space is not combined, even if two buffers are directly or indirectly provided to a communicator to be used for buffered sends.

In C, `buffer` is the starting address of a memory region. In Fortran, one can pass the first element of a memory region or a whole array, which must be ‘simply contiguous’ (for ‘simply contiguous,’ see also Section 19.1.12).

`MPI_COMM_DETACH_BUFFER(comm, buffer_addr, size)`

IN	<code>comm</code>	communicator (handle)
OUT	<code>buffer_addr</code>	initial buffer address (choice)
OUT	<code>size</code>	buffer size, in bytes (integer)

C binding

```
int MPI_Comm_detach_buffer(MPI_Comm comm, void *buffer_addr, int *size)
```

```
int MPI_Comm_detach_buffer_c(MPI_Comm comm, void *buffer_addr, MPI_Count *size)
```

Fortran 2008 binding

```
MPI_Comm_detach_buffer(comm, buffer_addr, size, ierror)
```

```
USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(C_PTR), INTENT(OUT) :: buffer_addr
```

```
INTEGER, INTENT(OUT) :: size
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Comm_detach_buffer(comm, buffer_addr, size, ierror) !(_c)
```

```
USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(C_PTR), INTENT(OUT) :: buffer_addr
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```

1 MPI_COMM_DETACH_BUFFER(COMM, BUFFER_ADDR, SIZE, IERROR)
2   INTEGER COMM, SIZE, IERROR
3   <type> BUFFER_ADDR(*)

```

Detach the communicator-specific buffer currently attached to the communicator.

```

8 MPI_SESSION_DETACH_BUFFER(session, buffer_addr, size)

```

```

9   IN      session          session (handle)
10
11  OUT     buffer_addr      initial buffer address (choice)
12
13  OUT     size              buffer size, in bytes (integer)

```

C binding

```

15 int MPI_Session_detach_buffer(MPI_Session session, void *buffer_addr,
16                               int *size)
17
18 int MPI_Session_detach_buffer_c(MPI_Session session, void *buffer_addr,
19                               MPI_Count *size)

```

Fortran 2008 binding

```

20 MPI_Session_detach_buffer(session, buffer_addr, size, ierror)
21   USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
22   TYPE(MPI_Session), INTENT(IN) :: session
23   TYPE(C_PTR), INTENT(OUT) :: buffer_addr
24   INTEGER, INTENT(OUT) :: size
25   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27 MPI_Session_detach_buffer(session, buffer_addr, size, ierror) !(_c)
28   USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
29   TYPE(MPI_Session), INTENT(IN) :: session
30   TYPE(C_PTR), INTENT(OUT) :: buffer_addr
31   INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
32   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

34 MPI_SESSION_DETACH_BUFFER(SESSION, BUFFER_ADDR, SIZE, IERROR)
35   INTEGER SESSION, SIZE, IERROR
36   <type> BUFFER_ADDR(*)

```

Detach the session-specific buffer currently attached to the session.

```

41 MPI_BUFFER_DETACH(buffer_addr, size)

```

```

42  OUT     buffer_addr      initial buffer address (choice)
43
44  OUT     size              buffer size, in bytes (integer)

```

C binding

```

46 int MPI_Buffer_detach(void *buffer_addr, int *size)

```



```
int MPI_Buffer_detach_c(void *buffer_addr, MPI_Count *size) 1
```

Fortran 2008 binding 2

```
MPI_Buffer_detach(buffer_addr, size, ierror) 3
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR 4
    TYPE(C_PTR), INTENT(OUT) :: buffer_addr 5
    INTEGER, INTENT(OUT) :: size 6
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror 7
```

```
MPI_Buffer_detach(buffer_addr, size, ierror) !(_c) 8
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR 9
    TYPE(C_PTR), INTENT(OUT) :: buffer_addr 10
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size 11
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror 12
```

Fortran binding 13

```
MPI_BUFFER_DETACH(BUFFER_ADDR, SIZE, IERROR) 14
    <type> BUFFER_ADDR(*) 15
    INTEGER SIZE, IERROR 16
```

Detach the MPI process-specific buffer `buffer` currently attached to MPI. 17

The procedure calls return the address and the size of the detached buffer. If 18
`MPI_BUFFER_AUTOMATIC` was used in the corresponding attach procedure, then 19
`MPI_BUFFER_AUTOMATIC` is also returned in the detach procedure and the value returned 20
in argument `size` is undefined. In this case, automatic buffering is disabled upon return 21
from the detach procedure. When using Fortran `mpi_f08`, the returned value is identical 22
to `c_loc(MPI_BUFFER_AUTOMATIC)`. Note that `c_loc()` is an intrinsic in the Fortran 23
`ISO_C_BINDING` module. 24

Advice to implementors. In Fortran, the implementation of `MPI_BUFFER_AUTOMATIC` 25
must allow the intrinsic `c_loc` to be applied to it. (*End of advice to implementors.*) 26

These procedures will delay their return until all messages currently in the (explicit or 27
automatic) buffer have been transmitted. Upon return of these procedures, the user may 28
reuse or deallocate the space taken by the buffer. 29

If the size of the detached buffer cannot be represented in `size`, it is set to 30
`MPI_UNDEFINED`. 31

The following `MPI_XXX_FLUSH_BUFFER` procedures will not return until all messages 32
currently in the buffer have been transmitted without detaching the buffer. 33

Rationale. These flush procedures provide the same functionality as an atomic 34
combination of first detaching the buffer and then attaching it again (but without 35
having to actually execute the detaching and the re-attaching of the buffer), but they 36
may be implemented with less internal overhead. (*End of rationale.*) 37

```
MPI_COMM_FLUSH_BUFFER(comm) 38
```

```
    IN      comm      communicator (handle) 39
```

C binding 40

```
int MPI_Comm_flush_buffer(MPI_Comm comm) 41
```

Fortran 2008 binding

```

1 MPI_Comm_flush_buffer(comm, ierror)
2     TYPE(MPI_Comm), INTENT(IN) :: comm
3     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

5 MPI_COMM_FLUSH_BUFFER(COMM, IERROR)
6     INTEGER COMM, IERROR

```

9 MPI_COMM_FLUSH_BUFFER will not return until all messages currently in the communicator-specific buffer of the calling MPI process have been transmitted.

```

12 MPI_SESSION_FLUSH_BUFFER(session)

```

```

13     IN          session          session (handle)

```

C binding

```

16 int MPI_Session_flush_buffer(MPI_Session session)

```

Fortran 2008 binding

```

17 MPI_Session_flush_buffer(session, ierror)
18     TYPE(MPI_Session), INTENT(IN) :: session
19     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

20 MPI_SESSION_FLUSH_BUFFER(SESSION, IERROR)
21     INTEGER SESSION, IERROR

```

22 MPI_SESSION_FLUSH_BUFFER will not return until all messages currently in the session-specific buffer of the calling MPI process have been transmitted.

```

23 MPI_BUFFER_FLUSH()

```

C binding

```

24 int MPI_Buffer_flush(void)

```

Fortran 2008 binding

```

25 MPI_Buffer_flush(ierror)
26     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

27 MPI_BUFFER_FLUSH(IERROR)
28     INTEGER IERROR

```

29 MPI_BUFFER_FLUSH will not return until all messages currently in the MPI process-specific buffer of the calling MPI process have been transmitted.

30 For all MPI_XXX_FLUSH_BUFFER procedures, there also exist the following nonblocking variants, which start the respective flush operation. These operations will not complete until all messages currently in the respective buffer of the calling MPI process have been transmitted.

MPI_COMM_IFLUSH_BUFFER(comm, request)	1
IN comm communicator (handle)	2
OUT request communication request (handle)	3
	4
	5
C binding	6
int MPI_Comm_iflush_buffer(MPI_Comm comm, MPI_Request *request)	7
	8
Fortran 2008 binding	9
MPI_Comm_iflush_buffer(comm, request, ierror)	10
TYPE(MPI_Comm), INTENT(IN) :: comm	11
TYPE(MPI_Request), INTENT(OUT) :: request	12
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	13
	14
Fortran binding	15
MPI_COMM_IFLUSH_BUFFER(COMM, REQUEST, IERROR)	16
INTEGER COMM, REQUEST, IERROR	17
	18
MPI_SESSION_IFLUSH_BUFFER(session, request)	19
IN session session (handle)	20
OUT request communication request (handle)	21
	22
	23
C binding	24
int MPI_Session_iflush_buffer(MPI_Session session, MPI_Request *request)	25
	26
Fortran 2008 binding	27
MPI_Session_iflush_buffer(session, request, ierror)	28
TYPE(MPI_Session), INTENT(IN) :: session	29
TYPE(MPI_Request), INTENT(OUT) :: request	30
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	31
	32
Fortran binding	33
MPI_SESSION_IFLUSH_BUFFER(SESSION, REQUEST, IERROR)	34
INTEGER SESSION, REQUEST, IERROR	35
	36
MPI_BUFFER_IFLUSH(request)	37
OUT request communication request (handle)	38
	39
	40
C binding	41
int MPI_Buffer_iflush(MPI_Request *request)	42
	43
Fortran 2008 binding	44
MPI_Buffer_iflush(request, ierror)	45
TYPE(MPI_Request), INTENT(OUT) :: request	46
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	47
	48
Fortran binding	
MPI_BUFFER_IFLUSH(REQUEST, IERROR)	

1 INTEGER REQUEST, IERROR
2

3 **Example 3.11.** Calls to attach and detach buffers.

```
4 #define BUFFSIZE 10000+MPI_BSEND_OVERHEAD
5 int size;
6 char *buff;
7 MPI_Buffer_attach(malloc(BUFFSIZE), BUFFSIZE);
8 /* a buffer of 10000 bytes can now be used by MPI_Bsend */
9 /* on all communicators, assuming only one message at a time is sent */
10 MPI_Buffer_detach(&buff, &size);
11 /* Buffer size reduced to zero */
12 MPI_Buffer_attach(buff, size);
13 /* Buffer of 10000 bytes available again */
```

14
15 **Example 3.12.** Calls to attach and detach communicator-specific buffers.

```
16 #define BUFFSIZE1 10000+MPI_BSEND_OVERHEAD
17 #define BUFFSIZE2 20000+MPI_BSEND_OVERHEAD
18 int size;
19 char *buff1, *buff2;
20 MPI_Comm world_dup;
21 MPI_Comm_dup(MPI_COMM_WORLD, &world_dup);
22 MPI_Comm_attach_buffer(MPI_COMM_WORLD, malloc(BUFFSIZE2), BUFFSIZE2);
23 MPI_Buffer_attach(malloc(BUFFSIZE1), BUFFSIZE1);
24 /* a buffer of 20000 bytes can now be used by MPI_Bsend for */
25 /* communication using MPI_COMM_WORLD, assuming only one message */
26 /* at a time is sent a buffer of 10000 bytes can now be used by */
27 /* MPI_Bsend for communication using any other communicator, */
28 /* including world_dup assuming only one message at a time is sent */
29 MPI_Comm_detach_buffer(MPI_COMM_WORLD, &buff1, &size);
30 MPI_Buffer_detach(&buff2, &size);
31 /* Both buffers are detached and no specific or MPI process-specific */
32 /* buffer can be used for further MPI_Bsend */
```

33 *Advice to users.* Even though the C procedures `MPI_Buffer_attach`,
34 `MPI_Session_attach_buffer`, `MPI_Comm_attach_buffer`, `MPI_Buffer_detach`,
35 `MPI_Session_detach_buffer` and `MPI_Comm_detach_buffer` have an argument of type
36 `void*`, these arguments are used differently: a pointer to the buffer is passed to
37 `MPI_Buffer_attach`, `MPI_Session_attach_buffer` and `MPI_Comm_attach_buffer`; the ad-
38 dress of the pointer is passed to `MPI_Buffer_detach`, `MPI_Session_detach_buffer` and
39 `MPI_Comm_detach_buffer`, so that this call can return the pointer value. In Fortran
40 with the `mpi` module or (deprecated) `mpif.h`, the type of the `buffer_addr` argument is
41 wrongly defined and the argument is therefore unused. In Fortran with the `mpi_f08`
42 module, the address of the buffer is returned as `TYPE(C_PTR)`, see also Example 9.1
43 about the use of `C_PTR` pointers. (*End of advice to users.*)

44
45 *Rationale.* In all cases, arguments are defined to be of type `void*` (rather than
46 `void*` and `void**`, respectively), so as to avoid complex type casts. E.g., in the
47 last two examples, `&buff`, which is of type `char**`, can be passed as argument to
48 `MPI_Buffer_detach`, `MPI_Session_detach_buffer` and `MPI_Comm_detach_buffer` without

type casting. If the formal parameter had type `void**` then we would need a type cast before and after each call. (*End of rationale.*)

General semantics of buffered mode sends. The statements made in this section describe the behavior of MPI for buffered-mode sends.

When no MPI process-specific buffer is currently (explicitly) attached and if no automatic buffering is enabled, MPI behaves as if a zero-sized MPI process-specific buffer is (implicitly) attached.

It is erroneous to detach a communicator-specific, session-specific, or MPI process-specific buffer, if no such buffer had been attached using a corresponding attach procedure. This includes attach procedure calls using `MPI_BUFFER_AUTOMATIC` as the buffer argument. It is erroneous to attach a communicator-specific, session-specific, or MPI process-specific buffer, if such buffer had already been attached using a corresponding attach procedure and not yet been detached again. This includes attach procedure calls using `MPI_BUFFER_AUTOMATIC` as the buffer argument. It is erroneous to flush a communicator-specific, session-specific or MPI process-specific buffer, if there is no buffer attached (including automatic buffering).

`MPI_COMM_ATTACH_BUFFER`, `MPI_SESSION_ATTACH_BUFFER`, and `MPI_BUFFER_ATTACH` are local. `MPI_COMM_DETACH_BUFFER`, `MPI_SESSION_DETACH_BUFFER`, `MPI_BUFFER_DETACH`, `MPI_COMM_FLUSH_BUFFER`, `MPI_SESSION_FLUSH_BUFFER`, and `MPI_BUFFER_FLUSH` are nonlocal; they must not return before all buffered messages in their related buffers are transmitted, and they must eventually return when all corresponding receive operations are started (provided that none are cancelled).

Automatic buffering with buffered mode sends. If the buffer used at the time of buffered mode send is set to the buffer address `MPI_BUFFER_AUTOMATIC`, then a buffer of sufficient size is automatically used by the MPI library.

Advice to users. When using automatic buffering, the user relinquishes control over buffer management, including allocation and deallocation decisions and timing, to the MPI library. If explicit control is needed over when and how much buffer space is allocated, automatic buffering must not be used. (*End of advice to users.*)

Advice to implementors. High-quality implementations of an MPI library should strive to support automatic buffering in a balanced fashion, i.e., providing the right balance between memory allocated for send operations and memory available for the end user. (*End of advice to implementors.*)

The flush operations for the communicator-specific, session-specific, and MPI process-specific buffers can also be used for automatic buffering. The flush procedure will not return until all automatically allocated buffers for the communicator-specific, session-specific, or MPI process-specific buffers, respectively, no longer hold message data and could be deallocated by the MPI library, if it chooses to do so.

Advice to users. With standard mode send, the limitation of needed buffer space is implemented within the MPI library through switching from internal buffering to internal synchronous mode. If the user wants to limit the automatically allocated

1 buffer space for buffered mode send using automatic buffering, the user may call
2 explicitly the appropriate flush procedure to wait until automatically allocated buffers
3 are deallocated. (*End of advice to users.*)
4

5
6 **Further rules.** In the case of an attached buffer (i.e., not using automatic buffering), the
7 user must provide as much buffering for outgoing messages as would be required if outgoing
8 message data were buffered by the sending MPI process, in the specified buffer space, using
9 a circular, contiguous-space allocation policy. We outline below a model implementation
10 that defines this policy. MPI may provide more buffering, and may use a better buffer
11 allocation algorithm than described below. On the other hand, MPI may signal an error
12 whenever the simple buffering allocator described below would run out of space. MPI must
13 not require more buffer space as described in the model implementation below.

14 MPI does not provide mechanisms for querying or controlling buffering done by standard
15 mode sends. It is expected that vendors will provide such information for their implemen-
16 tations.
17

18 *Rationale.* There is a wide spectrum of possible implementations of buffered commu-
19 nication operations: buffering can be done at sender, at receiver, or both; buffers can
20 be dedicated to one sender-receiver pair, or be shared by all communication opera-
21 tions; buffering can be done in real or in virtual memory; it can use dedicated memory,
22 or memory shared by other MPI processes; buffer space may be allocated statically or
23 be changed dynamically; etc. It does not seem feasible to provide a portable mech-
24 anism for querying or controlling buffering that would be compatible with all these
25 choices, yet provide meaningful information. (*End of rationale.*)
26

27 3.6.1 Model Implementation of Buffered Mode

28 The model implementation uses the packing and unpacking procedures described in Sec-
29 tion 5.2 and the nonblocking communication procedures described in Section 3.7.

30 We assume that a circular queue of pending message entries (PME) is maintained.
31 Each entry contains a communication request handle that identifies a pending nonblocking
32 send, a pointer to the next entry and the packed message data. The entries are stored in
33 successive locations in the buffer. Free space is available between the queue tail and the
34 queue head.
35

36 A buffered send call results in the execution of the following algorithm:

- 37 • Traverse sequentially the PME queue from head towards the tail, deleting all entries
38 for communication operations that have completed, up to the first entry with an
39 uncompleted request; update queue head to point to that entry.
40
- 41 • Compute the number of bytes, n , needed to store an entry for the new message. An
42 upper bound on n can be computed as follows: A call to the function
43 MPI_PACK_SIZE(count, datatype, comm, size), with the count, datatype and comm
44 arguments used in the MPI_BSEND call, returns an upper bound on the amount
45 of space needed to buffer the message data (see Section 5.2). The MPI constant
46 MPI_BSEND_OVERHEAD provides an upper bound on the additional space consumed
47 by the entry (e.g., for pointers or *envelope* information).
48

- Find the next contiguous empty space of n bytes in buffer (space following queue tail, or space at start of buffer if queue tail is too close to end of buffer). If space is not found then raise buffer overflow error.
- Append to end of PME queue in contiguous space the new entry that contains request handle, next pointer and packed message data; `MPI_PACK` is used to pack data.
- Post nonblocking send (standard mode) for packed data.
- Return

3.7 Nonblocking Communication

Nonblocking communication is important both for reasons of correctness and performance. For complex communication patterns, the use of only blocking communication (without buffering) is difficult because the programmer must ensure that each send is matched with a receive in an order that avoids *deadlock*. For communication patterns that are determined only at run time, this is even more difficult. Nonblocking communication can be used to avoid this problem, allowing programmers to express complex and possibly dynamic communication patterns without needing to ensure that all sends and receives are issued in an order that prevents deadlock (see Section 3.5 and the discussion of “safe” programs). Nonblocking communication also allows for the *overlap* of communication with different communication operations, e.g., to prevent the unintentional *serialization* of such operations, and for the *overlap* of communication with computation. Whether an implementation is able to accomplish an effective (from a performance standpoint) overlap of operations depends on the implementation itself and the system on which the implementation is running. Using nonblocking operations *permits* an implementation to overlap communication with computation, but does not require it to do so.

A nonblocking **send start** call *initiates* the send operation, but does not complete it. The send start call can return before the message was copied out of the send buffer. A separate **send complete** call is needed to complete the communication, i.e., to verify that the data has been copied out of the send buffer. With suitable hardware, the transfer of data out of the sender memory may proceed concurrently with computations done at the sender after the send was initiated and before it completed. Similarly, a nonblocking **receive start** call *initiates* the receive operation, but does not complete it. The call can return before a message is stored into the receive buffer. A separate **receive complete** call is needed to complete the receive operation and verify that the data has been received into the receive buffer. With suitable hardware, the transfer of data into the receiver memory may proceed concurrently with computations done after the receive was initiated and before it completed. The use of nonblocking receives may also avoid system buffering and memory-to-memory copying, as information is provided early on the location of the receive buffer.

Nonblocking send start calls can use the same four modes as blocking sends: *standard*, *buffered*, *synchronous*, and *ready*. These carry the same meaning. Sends of all modes, *ready* excepted, can be started whether a matching receive has been started or not; a nonblocking **ready** send can be started only if the matching receive is already started. In all cases, the send start call is *local*: it returns immediately, irrespective of the status of other MPI processes. If the call causes some system resource to be exhausted, then it will fail and return an error code. High-quality implementations of MPI should ensure that this happens only

1 in “pathological” cases. That is, an MPI implementation should be able to support a large
2 number of *pending* nonblocking operations.

3 The send-complete call returns no earlier than when all message data has been copied
4 out of the send buffer. It may carry additional meaning, depending on the send mode.

5 If the send mode is **synchronous**, then the send-complete call is *nonlocal*; the send can
6 complete only if a matching receive has been started and has been matched with the send.
7 Note that a synchronous mode send may complete, if matched by a nonblocking receive,
8 before the receive complete call occurs. (It can complete as soon as the sender “knows” the
9 transfer will complete, but before the receiver “knows” the transfer will complete.)

10 If the send mode is **buffered**, then the send-complete call is *local*; the send must
11 complete irrespective of the status of a matching receive. If there is no *pending* receive
12 operation, then the message must be buffered.

13 If the send mode is **standard**, then the send-complete call can be either *local* or
14 *nonlocal*. If the message is buffered, it is permitted for the send to complete before a
15 matching receive is started. On the other hand, it is permitted for the send not to complete
16 until a matching receive has been started and the message has been copied into the receive
17 buffer.

18 Nonblocking sends can be matched with blocking receives, and vice-versa.

19
20 *Advice to users.* The completion of a send operation may be delayed for standard
21 mode, and must be delayed for synchronous mode, until a matching receive has been
22 started. The use of nonblocking sends in these two cases allows the sender to proceed
23 ahead of the receiver, so that the computation is more tolerant of fluctuations in the
24 speeds of the two MPI processes.

25 Nonblocking sends in the buffered and ready modes have a more limited impact, e.g.,
26 the blocking version of buffered send is capable of completing regardless of when a
27 matching receive call is made. However, separating the start from the completion
28 of these sends still gives some opportunity for optimization within the MPI library.
29 For example, starting a buffered send gives an implementation more flexibility in
30 determining if and how the message is buffered. There are also advantages for both
31 nonblocking buffered and ready modes when data copying can be done concurrently
32 with computation.

33 The message-passing model implies that communication is initiated by the sender.
34 The communication will generally have lower overhead if a receive is already *started*
35 when the sender initiates the communication (data can be moved directly to the
36 receive buffer, and there is no need to queue a pending send request). However, a
37 receive operation can complete only after the matching send has *started*. The use of
38 nonblocking receives allows one to achieve lower communication overheads without
39 blocking the receiver while it waits for the send. (*End of advice to users.*)
40

41 3.7.1 Communication Request Objects

42
43 Nonblocking communication operations use opaque **request** objects to identify commu-
44 nication operations and match the operation that initiates the communication with the
45 operation that terminates it. These are system objects that are accessed via a handle.
46 A request object identifies various properties of a communication operation, such as the
47 send mode, the communication buffer that is associated with it, its context, the tag and
48 destination arguments to be used for a send, or the tag and source arguments to be used

for a receive. In addition, this object stores information about the status of the *pending* communication operation.

3.7.2 Communication Initiation

For the functions defined in this section, we use the same naming conventions as for blocking communication: a prefix of B, S, or R is used for *buffered*, *synchronous*, or *ready* mode. In addition, for these functions a prefix of I (for *immediate* and *incomplete*) indicates that the call is nonblocking.

MPI_ISEND(buf, count, datatype, dest, tag, comm, request)

IN	buf	initial address of send buffer (choice)
IN	count	number of elements in send buffer (non-negative integer)
IN	datatype	datatype of each send buffer element (handle)
IN	dest	rank of destination (integer)
IN	tag	message tag (integer)
IN	comm	communicator (handle)
OUT	request	communication request (handle)

C binding

```
int MPI_Isend(const void *buf, int count, MPI_Datatype datatype, int dest,
             int tag, MPI_Comm comm, MPI_Request *request)
```

```
int MPI_Isend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
               int dest, int tag, MPI_Comm comm, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Isend(buf, count, datatype, dest, tag, comm, request, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
  INTEGER, INTENT(IN) :: count, dest, tag
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Isend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER, INTENT(IN) :: dest, tag
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_ISEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
```

```

1     <type> BUF(*)
2     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
3
4     Start a standard mode nonblocking send.

```

```

6 MPI_IBSEND(buf, count, datatype, dest, tag, comm, request)
7
8     IN        buf                initial address of send buffer (choice)
9     IN        count              number of elements in send buffer (non-negative
10                                integer)
11    IN        datatype            datatype of each send buffer element (handle)
12    IN        dest                rank of destination (integer)
13    IN        tag                 message tag (integer)
14    IN        comm                communicator (handle)
15    IN        comm                communicator (handle)
16    OUT       request             communication request (handle)
17

```

C binding

```

19 int MPI_Ibsend(const void *buf, int count, MPI_Datatype datatype, int dest,
20               int tag, MPI_Comm comm, MPI_Request *request)
21
22 int MPI_Ibsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
23                 int dest, int tag, MPI_Comm comm, MPI_Request *request)
24

```

Fortran 2008 binding

```

25 MPI_Ibsend(buf, count, datatype, dest, tag, comm, request, ierror)
26     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
27     INTEGER, INTENT(IN) :: count, dest, tag
28     TYPE(MPI_Datatype), INTENT(IN) :: datatype
29     TYPE(MPI_Comm), INTENT(IN) :: comm
30     TYPE(MPI_Request), INTENT(OUT) :: request
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Ibsend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
34     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
35     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
36     TYPE(MPI_Datatype), INTENT(IN) :: datatype
37     INTEGER, INTENT(IN) :: dest, tag
38     TYPE(MPI_Comm), INTENT(IN) :: comm
39     TYPE(MPI_Request), INTENT(OUT) :: request
40     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
41

```

Fortran binding

```

42 MPI_IBSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
43     <type> BUF(*)
44     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
45
46     Start a buffered mode nonblocking send.
47
48

```

MPI_ISSEND(buf, count, datatype, dest, tag, comm, request)	1
IN buf	2
	3
IN count	4
	5
IN datatype	6
IN dest	7
	8
IN tag	9
IN comm	10
OUT request	11
	12
	13
C binding	14
int MPI_Issend(const void *buf, int count, MPI_Datatype datatype, int dest,	15
int tag, MPI_Comm comm, MPI_Request *request)	16
int MPI_Issend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,	17
int dest, int tag, MPI_Comm comm, MPI_Request *request)	18
	19
Fortran 2008 binding	20
MPI_Issend(buf, count, datatype, dest, tag, comm, request, ierror)	21
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf	22
INTEGER, INTENT(IN) :: count, dest, tag	23
TYPE(MPI_Datatype), INTENT(IN) :: datatype	24
TYPE(MPI_Comm), INTENT(IN) :: comm	25
TYPE(MPI_Request), INTENT(OUT) :: request	26
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	27
MPI_Issend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)	28
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf	29
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count	30
TYPE(MPI_Datatype), INTENT(IN) :: datatype	31
INTEGER, INTENT(IN) :: dest, tag	32
TYPE(MPI_Comm), INTENT(IN) :: comm	33
TYPE(MPI_Request), INTENT(OUT) :: request	34
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	35
	36
Fortran binding	37
MPI_ISSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	38
<type> BUF(*)	39
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR	40
Start a synchronous mode nonblocking send.	41
	42
	43
	44
	45
	46
	47
	48

```

1 MPI_IRSEND(buf, count, datatype, dest, tag, comm, request)
2   IN      buf                initial address of send buffer (choice)
3
4   IN      count              number of elements in send buffer (non-negative
5                               integer)
6
7   IN      datatype           datatype of each send buffer element (handle)
8
9   IN      dest                rank of destination (integer)
10
11  IN      tag                  message tag (integer)
12
13  IN      comm                 communicator (handle)
14
15  OUT     request              communication request (handle)

```

C binding

```

14 int MPI_Irsend(const void *buf, int count, MPI_Datatype datatype, int dest,
15               int tag, MPI_Comm comm, MPI_Request *request)
16

```

```

17 int MPI_Irsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
18                 int dest, int tag, MPI_Comm comm, MPI_Request *request)
19

```

Fortran 2008 binding

```

20 MPI_Irsend(buf, count, datatype, dest, tag, comm, request, ierror)
21   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
22   INTEGER, INTENT(IN) :: count, dest, tag
23   TYPE(MPI_Datatype), INTENT(IN) :: datatype
24   TYPE(MPI_Comm), INTENT(IN) :: comm
25   TYPE(MPI_Request), INTENT(OUT) :: request
26   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27

```

```

28 MPI_Irsend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
29   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
30   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
31   TYPE(MPI_Datatype), INTENT(IN) :: datatype
32   INTEGER, INTENT(IN) :: dest, tag
33   TYPE(MPI_Comm), INTENT(IN) :: comm
34   TYPE(MPI_Request), INTENT(OUT) :: request
35   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36

```

Fortran binding

```

37 MPI_IRSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
38   <type> BUF(*)
39   INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
40

```

41 Start a ready mode nonblocking send.

MPI_Irecv(buf, count, datatype, source, tag, comm, request)	1
OUT buf	2
	3
IN count	4
	5
IN datatype	6
IN source	7
IN tag	8
IN comm	9
OUT request	10
	11
	12
	13
C binding	14
int MPI_Irecv(void *buf, int count, MPI_Datatype datatype, int source, int tag,	15
MPI_Comm comm, MPI_Request *request)	16
int MPI_Irecv_c(void *buf, MPI_Count count, MPI_Datatype datatype, int source,	17
int tag, MPI_Comm comm, MPI_Request *request)	18
	19
Fortran 2008 binding	20
MPI_Irecv(buf, count, datatype, source, tag, comm, request, ierror)	21
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf	22
INTEGER, INTENT(IN) :: count, source, tag	23
TYPE(MPI_Datatype), INTENT(IN) :: datatype	24
TYPE(MPI_Comm), INTENT(IN) :: comm	25
TYPE(MPI_Request), INTENT(OUT) :: request	26
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	27
MPI_Irecv(buf, count, datatype, source, tag, comm, request, ierror) !(_c)	28
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf	29
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count	30
TYPE(MPI_Datatype), INTENT(IN) :: datatype	31
INTEGER, INTENT(IN) :: source, tag	32
TYPE(MPI_Comm), INTENT(IN) :: comm	33
TYPE(MPI_Request), INTENT(OUT) :: request	34
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	35
	36
Fortran binding	37
MPI_Irecv(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR)	38
<type> BUF(*)	39
INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR	40
Start a nonblocking receive.	41
	42
	43
	44
	45
	46
	47
	48

```

1 MPI_ISENDRECV(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount, recvtype,
2               source, recvtag, comm, request)
3
4   IN      sendbuf      initial address of send buffer (choice)
5   IN      sendcount    number of elements in send buffer (non-negative
6                       integer)
7   IN      sendtype     datatype of each send buffer element (handle)
8   IN      dest         rank of destination (integer)
9   IN      sendtag      send tag (integer)
10  IN      recvbuf      initial address of receive buffer (choice)
11  OUT     recvcount    number of elements in receive buffer (non-negative
12                       integer)
13  IN      recvtag      receive tag or MPI_ANY_TAG (integer)
14  IN      recvttype    datatype of each receive buffer element (handle)
15  IN      source       rank of source or MPI_ANY_SOURCE (integer)
16  IN      recvtag      receive tag or MPI_ANY_TAG (integer)
17  IN      comm         communicator (handle)
18  IN      request      communication request (handle)
19  OUT     request
20

```

C binding

```

21
22
23 int MPI_Isendrecv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
24                 int dest, int sendtag, void *recvbuf, int recvcount,
25                 MPI_Datatype recvtype, int source, int recvtag, MPI_Comm comm,
26                 MPI_Request *request)
27
28 int MPI_Isendrecv_c(const void *sendbuf, MPI_Count sendcount,
29                   MPI_Datatype sendtype, int dest, int sendtag, void *recvbuf,
30                   MPI_Count recvcount, MPI_Datatype recvtype, int source,
31                   int recvtag, MPI_Comm comm, MPI_Request *request)
32

```

Fortran 2008 binding

```

33 MPI_Isendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
34              recvttype, source, recvtag, comm, request, ierror)
35   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
36   INTEGER, INTENT(IN) :: sendcount, dest, sendtag, recvcount, source, recvtag
37   TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvttype
38   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
39   TYPE(MPI_Comm), INTENT(IN) :: comm
40   TYPE(MPI_Request), INTENT(OUT) :: request
41   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Isendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
44              recvttype, source, recvtag, comm, request, ierror) !(_c)
45   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
46   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
47   TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvttype
48   INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag

```

```

TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_ISENDRECV(SENDBUF, SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVBUF, REVCOUNT,
              RECVTYPE, SOURCE, RECVTAG, COMM, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNT, SENDTYPE, DEST, SENDTAG, REVCOUNT, RECVTYPE, SOURCE,
              RECVTAG, COMM, REQUEST, IERROR

```

Initiate a nonblocking communication request for a *send and receive* operation.

```

MPI_ISENDRECV_REPLACE(buf, count, datatype, dest, sendtag, source, recvtag, comm,
                      request)

```

INOUT	buf	initial address of send and receive buffer (choice)
IN	count	number of elements in send and receive buffer (non-negative integer)
IN	datatype	type of elements in send and receive buffer (handle)
IN	dest	rank of destination (integer)
IN	sendtag	send message tag (integer)
IN	source	rank of source or MPI_ANY_SOURCE (integer)
IN	recvtag	receive message tag or MPI_ANY_TAG (integer)
IN	comm	communicator (handle)
OUT	request	communication request (handle)

C binding

```

int MPI_Isendrecv_replace(void *buf, int count, MPI_Datatype datatype,
                          int dest, int sendtag, int source, int recvtag, MPI_Comm comm,
                          MPI_Request *request)
int MPI_Isendrecv_replace_c(void *buf, MPI_Count count, MPI_Datatype datatype,
                            int dest, int sendtag, int source, int recvtag, MPI_Comm comm,
                            MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Isendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
                      comm, request, ierror)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, dest, sendtag, source, recvtag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Isendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
2     comm, request, ierror) !(_c)
3     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
4     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
5     TYPE(MPI_Datatype), INTENT(IN) :: datatype
6     INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag
7     TYPE(MPI_Comm), INTENT(IN) :: comm
8     TYPE(MPI_Request), INTENT(OUT) :: request
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

11 MPI_ISENDRECV_REPLACE(BUF, COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG,
12     COMM, REQUEST, IERROR)
13     <type> BUF(*)
14     INTEGER COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG, COMM, REQUEST,
15     IERROR

```

Initiate a nonblocking communication request for a *send and receive* operation. The same buffer is used both for the send and for the receive, so that the message sent is replaced by the message received.

These calls allocate a communication request object and associate it with the request handle (the argument `request`). The request can be used later to query the status of the communication or wait for its completion.

A nonblocking send call indicates that the system may start copying data out of the send buffer. The sender should not modify any part of the send buffer after a nonblocking send operation is called, until the send completes.

A nonblocking receive call indicates that the system may start writing data into the receive buffer. The receiver should not access any part of the receive buffer after a nonblocking receive operation is called, until the receive completes.

Advice to users. To prevent problems with the argument copying and register optimization done by Fortran compilers, please note the hints in Sections [19.1.10–19.1.20](#). (*End of advice to users.*)

3.7.3 Communication Completion

The functions `MPI_WAIT` and `MPI_TEST` are used to complete a nonblocking communication. The *completion* of a send operation indicates that the sender is now free to update the send buffer (the send operation itself leaves the content of the send buffer unchanged). It does not indicate that the message has been received, rather, it may have been buffered by the communication subsystem. However, if a *synchronous mode send* was used, the *completion* of the send operation indicates that a matching receive was *initiated*, and that the message will eventually be received by this matching receive.

The *completion* of a receive operation indicates that the receive buffer contains the received message, the receiver is now free to access it, and that the status object is set. It does not indicate that the matching send operation has *completed* (but indicates, of course, that the send was *initiated*).

We shall use the following terminology: A **null handle** is a handle with value `MPI_REQUEST_NULL`. A *persistent communication request* and the handle to it are **inactive**

if the request is not associated with any ongoing communication (see Section 3.9). A handle is **active** if it is neither *null* nor *inactive*. An **empty** status is a status that is set to return `tag = MPI_ANY_TAG`, `source = MPI_ANY_SOURCE`, `error = MPI_SUCCESS`, and is also internally configured so that calls to `MPI_GET_COUNT` and `MPI_GET_ELEMENTS` return `count = 0` and `MPI_TEST_CANCELLED` returns false. We set a status variable to *empty* when the value returned by it is not significant. Status is set in this way so as to prevent errors due to accesses of stale information.

The fields in a **status** object returned by a call to `MPI_WAIT`, `MPI_TEST`, or any of the other derived functions (`MPI_{TEST|WAIT}{ALL|SOME|ANY}`), where the **request** corresponds to a send call, are undefined, with two exceptions: The error status field will contain valid information if the wait or test call returned with `MPI_ERR_IN_STATUS`; and the returned status can be queried by the call `MPI_TEST_CANCELLED`.

Error codes belonging to the error class `MPI_ERR_IN_STATUS` should be returned only by the MPI completion functions that take arrays of `MPI_Status`. For the functions that take a single `MPI_Status` argument, the error code is returned by the function, and the value of the `MPI_ERROR` field in the `MPI_Status` argument is undefined (see 3.2.5).

`MPI_WAIT(request, status)`

INOUT	request	request (handle)
OUT	status	status object (status)

C binding

```
int MPI_Wait(MPI_Request *request, MPI_Status *status)
```

Fortran 2008 binding

```
MPI_Wait(request, status, ierror)
  TYPE(MPI_Request), INTENT(INOUT) :: request
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WAIT(REQUEST, STATUS, IERROR)
  INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR
```

A call to `MPI_WAIT` returns when the operation identified by **request** is *complete*. If the request is an *active persistent communication request*, it is marked *inactive*. Any other type of request is deallocated and the request handle is set to `MPI_REQUEST_NULL`. `MPI_WAIT` is in general a *nonlocal* procedure. When the operation represented by the **request** is *enabled* then a call to `MPI_WAIT` is a *local* procedure call.

The call returns, in **status**, information on the completed operation. The content of the status object for a receive operation can be accessed as described in Section 3.2.5. The status object for a send operation may be queried by a call to `MPI_TEST_CANCELLED` (see Section 3.8).

One is allowed to call `MPI_WAIT` with a *null* or *inactive* request argument. In this case the procedure returns immediately with *empty* status.

Advice to users. Successful return of `MPI_WAIT` after a `MPI_IBSEND` implies that the user send buffer can be reused—i.e., data has been sent out or copied into a buffer

1 attached with `MPI_BUFFER_ATTACH`, `MPI_COMM_ATTACH_BUFFER` or
 2 `MPI_SESSION_ATTACH_BUFFER`. Further, at this point, we can no longer *cancel* the
 3 send (see Section 3.8). If a matching receive is never *started*, then the buffer cannot be
 4 freed. This runs somewhat counter to the stated goal of `MPI_CANCEL` (always being
 5 able to free program space that was committed to the communication subsystem).
 6 (*End of advice to users.*)

7
 8 *Advice to implementors.* In a multithreaded environment, a call to `MPI_WAIT` should
 9 block only the calling thread, allowing the thread scheduler to schedule another thread
 10 for execution. (*End of advice to implementors.*)

11
 12
 13 `MPI_TEST(request, flag, status)`

14	INOUT	request	communication request (handle)
15	OUT	flag	true if operation completed (logical)
16	OUT	status	status object (status)
17			
18			

19 C binding

20
 21 `int MPI_Test(MPI_Request *request, int *flag, MPI_Status *status)`

22 Fortran 2008 binding

23 `MPI_Test(request, flag, status, ierror)`
 24 `TYPE(MPI_Request), INTENT(INOUT) :: request`
 25 `LOGICAL, INTENT(OUT) :: flag`
 26 `TYPE(MPI_Status) :: status`
 27 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

28 Fortran binding

29 `MPI_TEST(REQUEST, FLAG, STATUS, IERROR)`
 30 `INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR`
 31 `LOGICAL FLAG`

32
 33 A call to `MPI_TEST` returns `flag = true` if the operation identified by `request` is *complete*.
 34 In such a case, the status object is set to contain information on the completed operation.
 35 If the request is an *active persistent communication request*, it is marked as *inactive*. Any
 36 other type of request is deallocated and the request handle is set to `MPI_REQUEST_NULL`.
 37 The call returns `flag = false` if the operation identified by `request` is not complete. In this
 38 case, the value of the status object is undefined. `MPI_TEST` is a *local* procedure.

39 The return status object for a receive operation carries information that can be accessed
 40 as described in Section 3.2.5. The status object for a send operation carries information
 41 that can be accessed by a call to `MPI_TEST_CANCELLED` (see Section 3.8).

42 One is allowed to call `MPI_TEST` with a *null* or *inactive* request argument. In such a
 43 case the procedure returns with `flag = true` and *empty* status.

44 The procedures `MPI_WAIT` and `MPI_TEST` can be used to complete any request-based
 45 nonblocking or persistent operation.

46
 47 *Advice to users.* The use of the nonblocking `MPI_TEST` call allows the user to
 48 schedule alternative activities within a single thread of execution. An event-driven

thread scheduler can be emulated with periodic calls to `MPI_TEST`. (*End of advice to users.*)

Example 3.13. Simple usage of nonblocking operations and `MPI_WAIT`.

```
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_ISEND(a(1), 10, MPI_REAL, 1, tag, comm, request, ierr)
  ! **** do some computation to mask latency ****
  CALL MPI_WAIT(request, status, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_IRECV(a(1), 15, MPI_REAL, 0, tag, comm, request, ierr)
  ! **** do some computation to mask latency ****
  CALL MPI_WAIT(request, status, ierr)
END IF
```

A request object can be *freed* using the following MPI procedure.

`MPI_REQUEST_FREE(request)`

INOUT request communication request (handle)

C binding

`int MPI_Request_free(MPI_Request *request)`

Fortran 2008 binding

`MPI_Request_free(request, ierror)`
 TYPE(MPI_Request), INTENT(INOUT) :: request
 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

`MPI_REQUEST_FREE(REQUEST, IERROR)`
 INTEGER REQUEST, IERROR

`MPI_REQUEST_FREE` is a *local* procedure. Upon successful return, `MPI_REQUEST_FREE` sets `request` to `MPI_REQUEST_NULL`. For an *inactive* request representing any type of MPI operation, `MPI_REQUEST_FREE` shall do the *freeing stage* of the associated operation during its execution.

For a request representing a *nonblocking* point-to-point or a persistent point-to-point operation, it is permitted (although strongly discouraged) to call `MPI_REQUEST_FREE` when the request is *active*. In this special case, `MPI_REQUEST_FREE` will only mark the request for freeing and MPI will actually do the *freeing stage* of the operation associated with the request later.

The use of this procedure for generalized requests is described in Section 13.2.

Calling `MPI_REQUEST_FREE` with an *active* request representing any other type of MPI operation (e.g., any partitioned operation (see Chapter 4), any collective operation (see Chapter 6), any I/O operation (see Chapter 14), or any request-based RMA operation (see Chapter 12)) is *erroneous*.

Rationale. For point-to-point operations, the `MPI_REQUEST_FREE` mechanism is provided for reasons of performance and convenience on the sending side. (*End of*

1 *rationale.*)

2
3 *Advice to users.* Once a request is freed by a call to `MPI_REQUEST_FREE`, it is not
4 possible to check for the successful completion of the associated communication with
5 calls to `MPI_WAIT` or `MPI_TEST`. Also, if an error occurs subsequently during the
6 communication, an error code cannot be returned to the user—such an error must be
7 treated as fatal. An active receive request should never be freed as the receiver will
8 have no way to verify that the receive has completed and the receive buffer can be
9 reused. (*End of advice to users.*)

11 **Example 3.14.** An example using `MPI_REQUEST_FREE`.

```
12 CALL MPI_COMM_RANK(MPI_COMM_WORLD, rank, ierr)
13 IF (rank .EQ. 0) THEN
14   DO i=1,n
15     CALL MPI_ISEND(outval, 1, MPI_REAL, 1, 0, MPI_COMM_WORLD, req, ierr)
16     CALL MPI_REQUEST_FREE(req, ierr)
17     CALL MPI_Irecv(ival, 1, MPI_REAL, 1, 0, MPI_COMM_WORLD, req, ierr)
18     CALL MPI_WAIT(req, status, ierr)
19   END DO
20 ELSE IF (rank .EQ. 1) THEN
21   CALL MPI_Irecv(ival, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
22   CALL MPI_WAIT(req, status, ierr)
23   DO I=1,n-1
24     CALL MPI_ISEND(outval, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
25     CALL MPI_REQUEST_FREE(req, ierr)
26     CALL MPI_Irecv(ival, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
27     CALL MPI_WAIT(req, status, ierr)
28   END DO
29   CALL MPI_ISEND(outval, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
30   CALL MPI_WAIT(req, status, ierr)
31 END IF
```

3.7.4 Semantics of Nonblocking Communication Operations

32 The semantics of nonblocking communication operations are defined by suitably extending
33 the definitions in Section 3.5.

34
35 **Order.** Nonblocking communication operations are **ordered** according to the execution
36 order of the calls that *initiate* the communication. The **nonovertaking** requirement of
37 Section 3.5 is extended to nonblocking communication, with this definition of order being
38 used.

39
40 **Example 3.15.** Message ordering for nonblocking operations.

```
41 CALL MPI_COMM_RANK(comm, rank, ierr)
42 IF (rank .EQ. 0) THEN
43   CALL MPI_ISEND(a, 1, MPI_REAL, 1, 0, comm, r1, ierr)
44   CALL MPI_ISEND(b, 1, MPI_REAL, 1, 0, comm, r2, ierr)
45 ELSE IF (rank .EQ. 1) THEN
46   CALL MPI_Irecv(a, 1, MPI_REAL, 0, MPI_ANY_TAG, comm, r1, ierr)
47   CALL MPI_Irecv(b, 1, MPI_REAL, 0, 0, comm, r2, ierr)
48 END IF
```

```
CALL MPI_WAIT(r1, status, ierr)
CALL MPI_WAIT(r2, status, ierr)
```

The first send will match the first receive, even if both messages are sent before either receive is executed.

Progress. A call to `MPI_WAIT` that *completes* a receive will eventually terminate and return if a matching send has been *started*, unless the send is satisfied by another receive. In particular, if the matching send is *nonblocking*, then the receive should *complete* even if no call is executed by the sender to *complete* the send. Similarly, a call to `MPI_WAIT` that *completes* a send will eventually return if a matching receive has been *started*, unless the receive is satisfied by another send, and even if no call is executed to *complete* the receive.

Example 3.16. An illustration of progress semantics.

```
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SSEND(a, 1, MPI_REAL, 1, 0, comm, ierr)
  CALL MPI_SEND(b, 1, MPI_REAL, 1, 1, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_IRECV(a, 1, MPI_REAL, 0, 0, comm, r, ierr)
  CALL MPI_RECV(b, 1, MPI_REAL, 0, 1, comm, status, ierr)
  CALL MPI_WAIT(r, status, ierr)
END IF
```

This code should not deadlock in a correct MPI implementation. The first synchronous send must complete once the matching (nonblocking) receive is *started*, even though the completing wait call has not yet been reached. Thus, the sending MPI process will continue and execute the second send procedure, allowing the receiving MPI process to complete execution.

If an `MPI_TEST` that *completes* a receive is repeatedly called with the same arguments, and a matching send has been *started*, then the call will eventually return `flag = true`, unless the send is satisfied by another receive. If an `MPI_TEST` that *completes* a send is repeatedly called with the same arguments, and a matching receive has been *started*, then the call will eventually return `flag = true`, unless the receive is satisfied by another send. See also Section 2.9 on *progress*.

3.7.5 Multiple Completions

It is convenient to be able to wait for the *completion* of any, some, or all the operations in a list, rather than having to wait for a specific message. A call to `MPI_WAITANY` or `MPI_TESTANY` can be used to wait for the *completion* of one out of several operations. A call to `MPI_WAITALL` or `MPI_TESTALL` can be used to wait for all *pending* operations in a list. A call to `MPI_WAIT SOME` or `MPI_TEST SOME` can be used to *complete* all enabled operations in a list.

`MPI_WAITANY(count, array_of_requests, index, status)`

IN	count	list length (non-negative integer)
INOUT	array_of_requests	array of requests (array of handles)

1	OUT	index	index of handle for operation that completed (integer)
2			
3	OUT	status	status object (status)
4			

C binding

```

5
6 int MPI_Waitany(int count, MPI_Request array_of_requests[], int *index,
7               MPI_Status *status)
8

```

Fortran 2008 binding

```

9
10 MPI_Waitany(count, array_of_requests, index, status, ierror)
11     INTEGER, INTENT(IN) :: count
12     TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
13     INTEGER, INTENT(OUT) :: index
14     TYPE(MPI_Status) :: status
15     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
16

```

Fortran binding

```

17 MPI_WAITANY(COUNT, ARRAY_OF_REQUESTS, INDEX, STATUS, IERROR)
18     INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE), IERROR
19

```

20 Blocks until one of the operations associated with the *active* requests in the array has
21 *completed*. If more than one operation is enabled and can terminate, one is arbitrarily
22 chosen. Returns in *index* the index of that request in the array and returns in *status* the
23 status of the completing operation. (The array is indexed from zero in C, and from one in
24 Fortran.) If the request is an *active persistent communication request*, it is marked *inactive*.
25 Any other type of request is deallocated and the request handle is set to
26 MPI_REQUEST_NULL.

27 The *array_of_requests* list may contain *null* or *inactive* handles. If the list contains no
28 *active* handles (list has length zero or all entries are *null* or *inactive*), then the call returns
29 immediately with *index* = MPI_UNDEFINED, and an *empty status*.

30 The execution of MPI_WAITANY with an array containing multiple entries has the
31 same effect as the execution of MPI_WAIT with the array entry indicated by the output
32 value of *index* (unless the output value of *index* is MPI_UNDEFINED). MPI_WAITANY with
33 an array containing one *active* entry is equivalent to MPI_WAIT.

```

34
35 MPI_TESTANY(count, array_of_requests, index, flag, status)
36

```

37	IN	count	list length (non-negative integer)
38	INOUT	array_of_requests	array of requests (array of handles)
39			
40	OUT	index	index of operation that completed or MPI_UNDEFINED if none completed (integer)
41			
42	OUT	flag	true if one of the operations is complete (logical)
43	OUT	status	status object (status)
44			

C binding

```

45
46 int MPI_Testany(int count, MPI_Request array_of_requests[], int *index,
47               int *flag, MPI_Status *status)
48

```

Fortran 2008 binding

```

MPI_Testany(count, array_of_requests, index, flag, status, ierror)
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
  INTEGER, INTENT(OUT) :: index
  LOGICAL, INTENT(OUT) :: flag
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TESTANY(COUNT, ARRAY_OF_REQUESTS, INDEX, FLAG, STATUS, IERROR)
  INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE), IERROR
  LOGICAL FLAG

```

Tests for *completion* of either one or none of the operations associated with *active* handles. In the former case, it returns `flag = true`, returns in `index` the index of this request in the array, and returns in `status` the status of that operation. If the request is an *active persistent communication request*, it is marked as *inactive*. Any other type of request is deallocated and the handle is set to `MPI_REQUEST_NULL`. (The array is indexed from zero in C, and from one in Fortran.) In the latter case (no operation *completed*), it returns `flag = false`, returns a value of `MPI_UNDEFINED` in `index` and `status` is undefined.

The array may contain *null* or inactive handles. If the array contains no *active* handles then the call returns *immediately* with `flag = true`, `index = MPI_UNDEFINED`, and an *empty* status.

If the array of requests contains *active* handles then the execution of `MPI_TESTANY` has the same effect as the execution of `MPI_TEST` with each of the *active* handles in the array in some arbitrary order, until one call returns `flag = true`, or all return `flag = false`. In the former case, `index` is set to indicate which array element returned `flag = true` and in the latter case, it is set to `MPI_UNDEFINED`. `MPI_TESTANY` with an array containing one *active* entry is equivalent to `MPI_TEST`.

```

MPI_WAITALL(count, array_of_requests, array_of_statuses)

```

IN	count	list length (non-negative integer)
INOUT	array_of_requests	array of requests (array of handles)
OUT	array_of_statuses	array of status objects (array of status)

C binding

```

int MPI_Waitall(int count, MPI_Request array_of_requests[],
               MPI_Status array_of_statuses[])

```

Fortran 2008 binding

```

MPI_Waitall(count, array_of_requests, array_of_statuses, ierror)
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
  TYPE(MPI_Status) :: array_of_statuses(*)
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

1 MPI_WAITALL(COUNT, ARRAY_OF_REQUESTS, ARRAY_OF_STATUSES, IERROR)
2   INTEGER COUNT, ARRAY_OF_REQUESTS(*), ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *),
3     IERROR
4
5

```

Blocks until all communication operations associated with *active* handles in the list *complete*, and returns the status of all these operations (this includes the case where no handle in the list is *active*). Both arrays have the same number of valid entries. The *i*-th entry in *array_of_statuses* is set to the return status of the *i*-th operation. *Active persistent requests* are marked *inactive*. Requests of any other type are deallocated and the corresponding handles in the array are set to `MPI_REQUEST_NULL`. The list may contain *null* or *inactive* handles. The call sets to *empty* the status of each such entry.

The error-free execution of `MPI_WAITALL` has the same effect as the execution of `MPI_WAIT` for each of the array elements in some arbitrary order. `MPI_WAITALL` with an array of length one is equivalent to `MPI_WAIT`.

When one or more of the communication operations *completed* by a call to `MPI_WAITALL` fail, it is desirable to return specific information on each communication. The function `MPI_WAITALL` will return in such case the error code `MPI_ERR_IN_STATUS` and will set the error field of each status to a specific error code. This code will be `MPI_SUCCESS`, if the specific communication *completed*; it will be another specific error code, if it failed; or it can be `MPI_ERR_PENDING` if it has neither failed nor *completed*. The function `MPI_WAITALL` will return `MPI_SUCCESS` if no request had an error, or will return another error code if it failed for other reasons (such as invalid arguments). In such cases, it will not update the error fields of the statuses.

Rationale. This design streamlines error handling in the application. The application code need only test the (single) function result to determine if an error has occurred. It needs to check each individual status only when an error occurred. (*End of rationale.*)

```

30 MPI_TESTALL(count, array_of_requests, flag, array_of_statuses)
31
32

```

IN	count	list length (non-negative integer)
INOUT	array_of_requests	array of requests (array of handles)
OUT	flag	true if all of the operations are complete (logical)
OUT	array_of_statuses	array of status objects (array of status)

C binding

```

38
39 int MPI_Testall(int count, MPI_Request array_of_requests[], int *flag,
40               MPI_Status array_of_statuses[])
41

```

Fortran 2008 binding

```

42 MPI_Testall(count, array_of_requests, flag, array_of_statuses, ierror)
43   INTEGER, INTENT(IN) :: count
44   TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
45   LOGICAL, INTENT(OUT) :: flag
46   TYPE(MPI_Status) :: array_of_statuses(*)
47   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48

```


Fortran binding

```

MPI_TESTALL(COUNT, ARRAY_OF_REQUESTS, FLAG, ARRAY_OF_STATUSES, IERROR)
    INTEGER COUNT, ARRAY_OF_REQUESTS(*), ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *),
        IERROR
    LOGICAL FLAG

```

Returns `flag = true` if all communication operations associated with *active* handles in the array have *completed* (this includes the case where no handle in the list is *active*). In this case, each status entry that corresponds to an *active* request is set to the status of the corresponding operation. *Active persistent requests* are marked *inactive*. Requests of any other type are deallocated and the corresponding handles in the array are set to `MPI_REQUEST_NULL`. Each status entry that corresponds to a *null* or *inactive* handle is set to *empty*.

Otherwise, `flag = false` is returned, no request is modified and the values of the status entries are undefined. This is a *local* procedure.

Errors that occurred during the execution of `MPI_TESTALL` are handled in the same manner as errors in `MPI_WAITALL`.

```

MPI_WAITSSOME(incount, array_of_requests, outcount, array_of_indices, array_of_statuses)

```

IN	incount	length of array_of_requests (non-negative integer)
INOUT	array_of_requests	array of requests (array of handles)
OUT	outcount	number of completed requests (integer)
OUT	array_of_indices	array of indices of operations that completed (array of integers)
OUT	array_of_statuses	array of status objects for operations that completed (array of status)

C binding

```

int MPI_Waitssome(int incount, MPI_Request array_of_requests[], int *outcount,
    int array_of_indices[], MPI_Status array_of_statuses[])

```

Fortran 2008 binding

```

MPI_Waitssome(incount, array_of_requests, outcount, array_of_indices,
    array_of_statuses, ierror)

```

```

INTEGER, INTENT(IN) :: incount
TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(incount)
INTEGER, INTENT(OUT) :: outcount, array_of_indices(*)
TYPE(MPI_Status) :: array_of_statuses(*)
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_WAITSSOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT, ARRAY_OF_INDICES,
    ARRAY_OF_STATUSES, IERROR)
    INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),
        ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *), IERROR

```

1 Waits until at least one of the operations associated with *active* handles in the list have
 2 *completed*. Returns in *outcount* the number of requests from the list *array_of_requests* that
 3 have *completed*. Returns in the first *outcount* locations of the array *array_of_indices* the
 4 indices of these operations (index within the array *array_of_requests*; the array is indexed
 5 from zero in C and from one in Fortran). Returns in the first *outcount* locations of the array
 6 *array_of_statuses* the status for these *completed* operations. *Completed active persistent*
 7 *requests* are marked as *inactive*. Any other type or request that *completed* is deallocated,
 8 and the associated handle is set to `MPI_REQUEST_NULL`.

9 If the list contains no *active* handles, then the call returns *immediately* with *outcount*
 10 = `MPI_UNDEFINED`.

11 When one or more of the communication operations *completed* by `MPI_WAITSSOME`
 12 fails, then it is desirable to return specific information on each communication. The argu-
 13 ments *outcount*, *array_of_indices* and *array_of_statuses* will be adjusted to indicate *comple-*
 14 *tion* of all communication operations that have succeeded or failed. The call will return
 15 the error code `MPI_ERR_IN_STATUS` and the error field of each status returned will be set
 16 to indicate success or to indicate the specific error that occurred. The call will return
 17 `MPI_SUCCESS` if no request resulted in an error, and will return another error code if it
 18 failed for other reasons (such as invalid arguments). In such cases, it will not update the
 19 error fields of the statuses.

20
 21
 22 `MPI_TESTSSOME`(*incount*, *array_of_requests*, *outcount*, *array_of_indices*, *array_of_statuses*)

23			
24	IN	<i>incount</i>	length of <i>array_of_requests</i> (non-negative integer)
25	INOUT	<i>array_of_requests</i>	array of requests (array of handles)
26			
27	OUT	<i>outcount</i>	number of completed requests (integer)
28	OUT	<i>array_of_indices</i>	array of indices of operations that completed (array 29 of integers)
30			
31	OUT	<i>array_of_statuses</i>	array of status objects for operations that completed 32 (array of status)
33			

34 C binding

35 `int MPI_Testsome(int incount, MPI_Request array_of_requests[], int *outcount,`
 36 `int array_of_indices[], MPI_Status array_of_statuses[])`

37 Fortran 2008 binding

38 `MPI_Testsome(incount, array_of_requests, outcount, array_of_indices,`
 39 `array_of_statuses, ierror)`

40 `INTEGER, INTENT(IN) :: incount`

41 `TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(incount)`

42 `INTEGER, INTENT(OUT) :: outcount, array_of_indices(*)`

43 `TYPE(MPI_Status) :: array_of_statuses(*)`

44 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

45 Fortran binding

46 `MPI_TESTSSOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT, ARRAY_OF_INDICES,`
 47 `ARRAY_OF_STATUSES, IERROR)`

```
INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),
      ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *), IERROR
```

This procedure behaves like `MPI_WAITSSOME`, except that it returns *immediately*. If no operation has completed it returns `outcount = 0`. If there is no *active* handle in the list it returns `outcount = MPI_UNDEFINED`.

`MPI_TESTSSOME` is a *local* procedure, which returns *immediately*, whereas `MPI_WAITSSOME` will block until a communication *completes*, if it was passed a list that contains at least one *active* handle. Both calls fulfill a **fairness requirement**: If a request for a receive repeatedly appears in a list of requests passed to `MPI_WAITSSOME` or `MPI_TESTSSOME`, and a matching send has been *started*, then the receive will eventually succeed, unless the send is satisfied by another receive; and similarly for send requests.

Errors that occur during the execution of `MPI_TESTSSOME` are handled as for `MPI_WAITSSOME`.

Advice to users. The use of `MPI_TESTSSOME` is likely to be more efficient than the use of `MPI_TESTANY`. The former returns information on all *completed* communication operations, with the latter, a new call is required for each communication that completes.

A server with multiple clients can use `MPI_WAITSSOME` so as not to starve any client. Clients send messages to the server with service requests. The server calls `MPI_WAITSSOME` with one receive request for each client, and then handles all receives that completed. If a call to `MPI_WAITANY` is used instead, then one client could starve while requests from another client always sneak in first. (*End of advice to users.*)

Advice to implementors. `MPI_TESTSSOME` should *complete* as many *pending* communication operations of the `array_of_requests` as possible. (*End of advice to implementors.*)

Example 3.17. Client-server code (starvation can occur).

```
CALL MPI_COMM_SIZE(comm, size, ierr)
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .GT. 0) THEN          ! client code
  DO WHILE(.TRUE.)
    CALL MPI_ISEND(a, n, MPI_REAL, 0, tag, comm, request, ierr)
    CALL MPI_WAIT(request, status, ierr)
  END DO
ELSE                            ! rank=0 -- server code
  DO i=1,size-1
    CALL MPI_IRecv(a(1,i), n, MPI_REAL, i, tag, &
                  comm, request_list(i), ierr)
  END DO
  DO WHILE(.TRUE.)
    CALL MPI_WAITANY(size-1, request_list, index, status, ierr)
    CALL DO_SERVICE(a(1,index)) ! handle one message
    CALL MPI_IRecv(a(1, index), n, MPI_REAL, index, tag, &
                  comm, request_list(index), ierr)
  END DO
END IF
```

Example 3.18. Same code, using MPI_WAITSSOME.

```

1  CALL MPI_COMM_SIZE(comm, size, ierr)
2  CALL MPI_COMM_RANK(comm, rank, ierr)
3  IF (rank .GT. 0) THEN          ! client code
4      DO WHILE(.TRUE.)
5          CALL MPI_ISEND(a, n, MPI_REAL, 0, tag, comm, request, ierr)
6          CALL MPI_WAIT(request, status, ierr)
7      END DO
8  ELSE                          ! rank=0 -- server code
9      DO i=1,size-1
10         CALL MPI_IRECV(a(1,i), n, MPI_REAL, i, tag, &
11             comm, request_list(i), ierr)
12     END DO
13     DO WHILE(.TRUE.)
14         CALL MPI_WAITSSOME(size, request_list, numdone, &
15             indices, statuses, ierr)
16         DO i=1,numdone
17             CALL DO_SERVICE(a(1, indices(i)))
18             CALL MPI_IRECV(a(1, indices(i)), n, MPI_REAL, 0, tag, &
19                 comm, request_list(indices(i)), ierr)
20         END DO
21     END DO
22 END IF

```

3.7.6 Non-Destructive Test of status

These procedures are useful for accessing the information associated with a request, without *freeing* the request (in case the user is expected to access it later). It allows one to layer libraries more conveniently, since multiple layers of software may access the same *completed* request and extract from it the status information.

MPI_REQUEST_GET_STATUS(request, flag, status)

IN	request	request (handle)
OUT	flag	boolean flag, same as from MPI_TEST (logical)
OUT	status	status object if flag is true (status)

C binding

```
int MPI_Request_get_status(MPI_Request request, int *flag, MPI_Status *status)
```

Fortran 2008 binding

```

MPI_Request_get_status(request, flag, status, ierror)
    TYPE(MPI_Request), INTENT(IN) :: request
    LOGICAL, INTENT(OUT) :: flag
    TYPE(MPI_Status) :: status
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```
MPI_REQUEST_GET_STATUS(REQUEST, FLAG, STATUS, IERROR)
```

INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR
 LOGICAL FLAG

Sets `flag = true` if the operation is *complete*, and, if so, returns in `status` the request status. However, unlike `test` or `wait`, it does not deallocate or *inactivate* the request; a subsequent call to `test`, `wait` or `free` should be executed with that request. It sets `flag = false` if the operation is not *complete*.

One is allowed to call `MPI_REQUEST_GET_STATUS` with a *null* or *inactive* request argument. In such a case the procedure returns with `flag = true` and *empty* status.

The *progress* rule for `MPI_TEST`, as described in Section 3.7.4, also applies to `MPI_REQUEST_GET_STATUS`.

`MPI_REQUEST_GET_STATUS_ANY(count, array_of_requests, index, flag, status)`

IN	count	list length (non-negative integer)
IN	array_of_requests	array of requests (array of handles)
OUT	index	index of operation that completed or MPI_UNDEFINED if none completed (integer)
OUT	flag	true if one of the operations is complete (logical)
OUT	status	status object if flag is true (status)

C binding

```
int MPI_Request_get_status_any(int count,
                              const MPI_Request array_of_requests[], int *index, int *flag,
                              MPI_Status *status)
```

Fortran 2008 binding

```
MPI_Request_get_status_any(count, array_of_requests, index, flag, status,
                           ierror)
INTEGER, INTENT(IN) :: count
TYPE(MPI_Request), INTENT(IN) :: array_of_requests(count)
INTEGER, INTENT(OUT) :: index
LOGICAL, INTENT(OUT) :: flag
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_REQUEST_GET_STATUS_ANY(COUNT, ARRAY_OF_REQUESTS, INDEX, FLAG, STATUS,
                            IERROR)
INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE), IERROR
LOGICAL FLAG
```

Tests for *completion* of either one or none of the operations associated with *active* handles. In the former case, it returns `flag = true`, returns in `index` the index of this request in the array, and returns in `status` the status of that operation. (The array is indexed from zero in C, and from one in Fortran.) In the latter case (no operation *completed*), it returns `flag = false`, returns a value of `MPI_UNDEFINED` in `index` and `status` is undefined.

1 The array may contain *null* or inactive handles. If the array contains no *active* handles
 2 then the call returns *immediately* with `flag = true`, `index = MPI_UNDEFINED`, and an *empty*
 3 *status*.

4 If the array of requests contains active handles then the execution of
 5 `MPI_REQUEST_GET_STATUS_ANY` has the same effect as the execution of
 6 `MPI_REQUEST_GET_STATUS` with each of the active array elements in some arbitrary
 7 order, until one call returns `flag = true`, or all return `flag = false`. In the former case, `index`
 8 is set to indicate which array element returned `flag = true` and in the latter case, it is
 9 set to `MPI_UNDEFINED`. `MPI_REQUEST_GET_STATUS_ANY` with an array containing one
 10 request is equivalent to `MPI_REQUEST_GET_STATUS`.

11
 12
 13 `MPI_REQUEST_GET_STATUS_ALL(count, array_of_requests, flag, array_of_statuses)`

14	IN	count	list length (non-negative integer)
15	IN	array_of_requests	array of requests (array of handles)
16	OUT	flag	true if all of the operations are complete (logical)
17	OUT	array_of_statuses	array of status objects (array of status)

20 C binding

```
21 int MPI_Request_get_status_all(int count,
22                               const MPI_Request array_of_requests[], int *flag,
23                               MPI_Status array_of_statuses[])
24
```

25 Fortran 2008 binding

```
26 MPI_Request_get_status_all(count, array_of_requests, flag, array_of_statuses,
27                             ierror)
28     INTEGER, INTENT(IN) :: count
29     TYPE(MPI_Request), INTENT(IN) :: array_of_requests(count)
30     LOGICAL, INTENT(OUT) :: flag
31     TYPE(MPI_Status) :: array_of_statuses(*)
32     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
33
```

33 Fortran binding

```
34 MPI_REQUEST_GET_STATUS_ALL(COUNT, ARRAY_OF_REQUESTS, FLAG, ARRAY_OF_STATUSES,
35                             IERROR)
36     INTEGER COUNT, ARRAY_OF_REQUESTS(*), ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *),
37     IERROR
38     LOGICAL FLAG
39
```

40 `MPI_REQUEST_GET_STATUS_ALL` returns `flag = true` if all communication operations
 41 associated with *active* handles in the array have *completed* (this includes the case where all
 42 handles in the list are *inactive* or `MPI_REQUEST_NULL`). In this case, each status entry that
 43 corresponds to an *active* request is set to the status of the corresponding operation. Unlike
 44 `test` or `wait`, it does not deallocate or *inactivate* the requests; a subsequent call to `test`, `wait`
 45 or `free` should be executed with each of those requests.

46 Each status entry that corresponds to a *null* or *inactive* handle is set to *empty*.

47 Otherwise, `flag = false` is returned and the values of the status entries are undefined.

The *progress* rule for MPI_TEST, as described in Section 3.7.4, also applies to MPI_REQUEST_GET_STATUS_ALL.

MPI_REQUEST_GET_STATUS_SOME(incount, array_of_requests, outcount,
array_of_indices, array_of_statuses)

IN	incount	length of array_of_requests (non-negative integer)
IN	array_of_requests	array of requests (array of handles)
OUT	outcount	number of completed requests (integer)
OUT	array_of_indices	array of indices of operations that completed (array of integers)
OUT	array_of_statuses	array of status objects for operations that completed (array of status)

C binding

```
int MPI_Request_get_status_some(int incount,
                               const MPI_Request array_of_requests[], int *outcount,
                               int array_of_indices[], MPI_Status array_of_statuses[])
```

Fortran 2008 binding

```
MPI_Request_get_status_some(incount, array_of_requests, outcount,
                             array_of_indices, array_of_statuses, ierror)
    INTEGER, INTENT(IN) :: incount
    TYPE(MPI_Request), INTENT(IN) :: array_of_requests(incount)
    INTEGER, INTENT(OUT) :: outcount, array_of_indices(*)
    TYPE(MPI_Status) :: array_of_statuses(*)
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_REQUEST_GET_STATUS_SOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT,
                             ARRAY_OF_INDICES, ARRAY_OF_STATUSES, IERROR)
    INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),
    ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *), IERROR
```

MPI_REQUEST_GET_STATUS_SOME returns in *outcount* the number of requests from the list *array_of_requests* that have *completed*. Returns in the first *outcount* locations of the array *array_of_indices* the indices of these operations within the array *array_of_requests*; the array is indexed from zero in C and from one in Fortran. Returns in the first *outcount* locations of the array *array_of_statuses* the status for these *completed* operations. However, unlike *test* or *wait*, it does not deallocate or *inactivate* any requests in *array_of_requests*; a subsequent call to *test*, *wait* or *free* should be executed with each completed request. If no operation in *array_of_requests* is complete, it returns *outcount* = 0. If all operations in *array_of_requests* are either MPI_REQUEST_NULL or *inactive*, *outcount* will be set to MPI_UNDEFINED. The *progress* rule for MPI_TEST, as described in Section 3.7.4, also applies to MPI_REQUEST_GET_STATUS_SOME.

Like MPI_WAIT_SOME and MPI_TEST_SOME, MPI_REQUEST_GET_STATUS_SOME fulfills a **fairness requirement**: If a request for a receive repeatedly appears in a list of requests passed to MPI_REQUEST_GET_STATUS_SOME, MPI_WAIT_SOME, or

1 MPI_TEST SOME and a matching send has been *started*, then the receive will eventually
 2 succeed, unless the send is satisfied by another receive; and similarly for send requests.

3 Errors that occur during the execution of MPI_REQUEST_GET_STATUS_SOME are
 4 handled as for MPI_WAIT_SOME.

5 *Advice to implementors.* MPI_REQUEST_GET_STATUS_SOME should *complete* as
 6 many pending communication operations as possible. (*End of advice to implemen-*
 7 *tors.*)

9 *Advice to users.* MPI_REQUEST_GET_STATUS_ANY,
 10 MPI_REQUEST_GET_STATUS_SOME, and MPI_REQUEST_GET_STATUS_ALL offer
 11 tradeoffs between precision and speed, as do the corresponding TEST and WAIT func-
 12 tions. The ANY variants are fast, but imprecise and unfair. The ALL variants will
 13 provide all-or-nothing information and/or completion, which can limit their applica-
 14 bility. The SOME variants, because of their precision and fairness guarantee, will
 15 typically be the slowest on a per-call basis. (*End of advice to users.*)
 16

17 3.8 Probe and Cancel

19 The MPI_PROBE, MPI_IPROBE, MPI_MPROBE, and MPI_IMPROBE procedures allow in-
 20 coming messages to be checked for, without actually receiving them. The user can then
 21 decide how to receive them, based on the information returned by the **probe** (basically,
 22 the information returned by **status**). In particular, the user may allocate memory for the
 23 receive buffer, according to the length of the probed message.
 24

25 The MPI_CANCEL procedure allows *pending* communication operations to be **can-**
 26 **celled**. This is required for cleanup. Posting a send or a receive ties up user resources
 27 (send or receive buffers), and a *cancel* may be needed to free these resources gracefully.

28 *Cancelling* a send request by calling MPI_CANCEL is deprecated. *Cancelling* a send-
 29 recv request by calling MPI_CANCEL is not allowed.

30 3.8.1 Probe

32 MPI_IPROBE(source, tag, comm, flag, status)

35	IN	source	rank of source or MPI_ANY_SOURCE (integer)
36	IN	tag	message tag or MPI_ANY_TAG (integer)
37	IN	comm	communicator (handle)
38	OUT	flag	true if there is a matching message that can be 39 received (logical)
40	OUT	status	status object (status)
41			
42			

43 C binding

44 int MPI_Iprobe(int source, int tag, MPI_Comm comm, int *flag,
 45 MPI_Status *status)

46 Fortran 2008 binding

47 MPI_Iprobe(source, tag, comm, flag, status, ierror)


```

INTEGER, INTENT(IN) :: source, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
LOGICAL, INTENT(OUT) :: flag
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_IPROBE(SOURCE, TAG, COMM, FLAG, STATUS, IERROR)
  INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
  LOGICAL FLAG

```

MPI_IPROBE returns `flag = true` if there is a message that can be received and that matches the pattern specified by the arguments `source`, `tag`, and `comm`. The call matches the same message that would have been received by a call to MPI_RECV with the same argument values for `source`, `tag`, `comm`, and `status` executed at the same point in the program, and returns in `status` the same value that would have been returned by MPI_RECV. Otherwise, the call returns `flag = false`, and leaves `status` undefined.

If MPI_IPROBE returns `flag = true`, then the content of the status object can be subsequently accessed as described in Section 3.2.5 to find the source, tag, and length of the probed message.

MPI_IPROBE is a *local* procedure since its return does not depend on MPI calls in other MPI processes, which is marked with the prefix I (for *immediate*).

A subsequent receive executed with the same communicator, and the source and tag returned in `status` by MPI_IPROBE will receive the message that was matched by the probe, if no other intervening receive occurs after the probe, and the send is not successfully *cancelled* before the receive. If the receiving MPI process is multithreaded, it is the user's responsibility to ensure that the last condition holds.

The `source` argument of MPI_IPROBE can be MPI_ANY_SOURCE, and the `tag` argument can be MPI_ANY_TAG, so that one can *probe* for *messages* from an arbitrary source and/or with an arbitrary tag. However, a specific communication context must be provided with the `comm` argument.

It is not necessary to receive a message immediately after it has been probed for, and the same message may be probed for several times before it is received.

A probe with MPI_PROC_NULL as source returns `flag = true`, and the status object returns `source = MPI_PROC_NULL`, `tag = MPI_ANY_TAG`, and `count = 0`; see Section 3.10.

```

MPI_PROBE(source, tag, comm, status)

```

IN	source	rank of source or MPI_ANY_SOURCE (integer)
IN	tag	message tag or MPI_ANY_TAG (integer)
IN	comm	communicator (handle)
OUT	status	status object (status)

C binding

```

int MPI_Probe(int source, int tag, MPI_Comm comm, MPI_Status *status)

```

Fortran 2008 binding

```

MPI_Probe(source, tag, comm, status, ierror)

```

```

1     INTEGER, INTENT(IN) :: source, tag
2     TYPE(MPI_Comm), INTENT(IN) :: comm
3     TYPE(MPI_Status) :: status
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

6 MPI_PROBE(SOURCE, TAG, COMM, STATUS, IERROR)
7     INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
8

```

MPI_PROBE behaves like MPI_IPROBE except that it is a *nonlocal* call that returns only after a matching message has been found.

The MPI implementation of MPI_PROBE and MPI_IPROBE needs to guarantee *progress*: if a call to MPI_PROBE has been issued by an MPI process, and a send that matches the probe has been *initiated* by some MPI process, then the call to MPI_PROBE will return, unless the message is received by another concurrent receive operation (that is executed by another thread at the probing MPI process).

Similarly, if an MPI process repeatedly calls MPI_IPROBE and a matching message has been issued, then MPI_IPROBE will eventually return `flag = true` unless the message is received by another concurrent receive operation or matched by a concurrent *matching probe*. See also Section 2.9 on *progress*.

Example 3.19. Use probe to wait for an incoming message.

```

22     CALL MPI_COMM_RANK(comm, rank, ierr)
23     IF (rank .EQ. 0) THEN
24         CALL MPI_SEND(i, 1, MPI_INTEGER, 2, 0, comm, ierr)
25     ELSE IF (rank .EQ. 1) THEN
26         CALL MPI_SEND(x, 1, MPI_REAL, 2, 0, comm, ierr)
27     ELSE IF (rank .EQ. 2) THEN
28         DO i=1,2
29             CALL MPI_PROBE(MPI_ANY_SOURCE, 0, comm, status, ierr)
30             IF (status(MPI_SOURCE) .EQ. 0) THEN
31                 100     CALL MPI_RECV(i, 1, MPI_INTEGER, 0, 0, comm, status, ierr)
32             ELSE
33                 200     CALL MPI_RECV(x, 1, MPI_REAL, 1, 0, comm, status, ierr)
34             END IF
35         END DO
36     END IF

```

Each message is received with the right type.

Example 3.20. A similar program to the previous example, but now it has a problem.

```

39     ! ----- THIS EXAMPLE IS ERRONEOUS -----
40     CALL MPI_COMM_RANK(comm, rank, ierr)
41     IF (rank .EQ. 0) THEN
42         CALL MPI_SEND(i, 1, MPI_INTEGER, 2, 0, comm, ierr)
43     ELSE IF (rank .EQ. 1) THEN
44         CALL MPI_SEND(x, 1, MPI_REAL, 2, 0, comm, ierr)
45     ELSE IF (rank .EQ. 2) THEN
46         DO i=1,2
47             CALL MPI_PROBE(MPI_ANY_SOURCE, 0, comm, status, ierr)
48             IF (status(MPI_SOURCE) .EQ. 0) THEN

```

```

100      CALL MPI_RECV(i, 1, MPI_INTEGER, MPI_ANY_SOURCE, &
                                0, comm, status, ierr)
      ELSE
200      CALL MPI_RECV(x, 1, MPI_REAL, MPI_ANY_SOURCE, &
                                0, comm, status, ierr)
      END IF
    END DO
  END IF

```

In Example 3.20, the two receive calls in statements labeled 100 and 200 in Example 3.19 are slightly modified, using `MPI_ANY_SOURCE` as the source argument. The program is now incorrect: the receive operation may receive a message that is distinct from the message probed by the preceding call to `MPI_PROBE`.

Advice to users. In a multithreaded MPI program, `MPI_PROBE` and `MPI_Iprobe` might need special care. If a thread *probes* for a message and then immediately posts a matching receive, the receive may match a message other than that found by the probe since another thread could concurrently receive that original message [34]. `MPI_Mprobe` and `MPI_Improbe` solve this problem by matching the incoming message so that it may only be received with `MPI_Mrecv` or `MPI_Irecv` on the corresponding *message handle*. (*End of advice to users.*)

Advice to implementors. A call to `MPI_PROBE` will match the message that would have been received by a call to `MPI_RECV` with the same argument values for *source*, *tag*, *comm*, and *status* executed at the same point. Suppose that this message has source *s*, tag *t* and communicator *c*. If the tag argument in the probe call has value `MPI_ANY_TAG` then the message probed will be the earliest pending message from source *s* with communicator *c* and any tag; in any case, the message probed will be the earliest pending message from source *s* with tag *t* and communicator *c* (this is the message that would have been received, so as to preserve message order). This message continues as the earliest pending message from source *s* with tag *t* and communicator *c*, until it is received. A receive operation subsequent to the probe that uses the same communicator as the probe and uses the tag and source values returned by the probe, must receive this message, unless it has already been received by another receive operation. (*End of advice to implementors.*)

3.8.2 Matching Probe

The function `MPI_PROBE` checks for incoming *messages* without receiving them. Since the list of incoming *messages* is global among the threads of each MPI process, it can be hard to use this functionality in threaded environments [34, 31].

Like `MPI_PROBE` and `MPI_Iprobe`, the **matching probe** operation (`MPI_Mprobe` and `MPI_Improbe` procedures) allow incoming *messages* to be queried without actually receiving them, except that `MPI_Mprobe` and `MPI_Improbe` provide a mechanism to receive the specific *message* that was matched regardless of other intervening probe or receive operations. This gives the application an opportunity to decide how to receive the message, based on the information returned by the probe. In particular, the user may allocate memory for the receive buffer, according to the length of the probed message.

```

1 MPI_IMPROBE(source, tag, comm, flag, message, status)
2     IN      source      rank of source or MPI_ANY_SOURCE (integer)
3     IN      tag         message tag or MPI_ANY_TAG (integer)
4     IN      comm        communicator (handle)
5     OUT     flag         true if there is a matching message that can be
6                       received (logical)
7     OUT     message     returned message (handle)
8     OUT     status      status object (status)
9
10
11

```

C binding

```

12 int MPI_Improbe(int source, int tag, MPI_Comm comm, int *flag,
13                MPI_Message *message, MPI_Status *status)
14
15

```

Fortran 2008 binding

```

16 MPI_Improbe(source, tag, comm, flag, message, status, ierror)
17     INTEGER, INTENT(IN) :: source, tag
18     TYPE(MPI_Comm), INTENT(IN) :: comm
19     LOGICAL, INTENT(OUT) :: flag
20     TYPE(MPI_Message), INTENT(OUT) :: message
21     TYPE(MPI_Status) :: status
22     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24

```

Fortran binding

```

25 MPI_IMPROBE(SOURCE, TAG, COMM, FLAG, MESSAGE, STATUS, IERROR)
26     INTEGER SOURCE, TAG, COMM, MESSAGE, STATUS(MPI_STATUS_SIZE), IERROR
27     LOGICAL FLAG
28

```

MPI_IMPROBE returns `flag = true` if there is a message that can be received and that matches the pattern specified by the arguments `source`, `tag`, and `comm`. The call matches the same message that would have been received by a call to `MPI_RECV` with the same argument values for `source`, `tag`, `comm`, and `status` executed at the same point in the program and returns in `status` the same value that would have been returned by `MPI_RECV`. In addition, it returns in `message` a **message handle** to the matched message. Otherwise, the call returns `flag = false`, and leaves `status` and `message` undefined.

MPI_IMPROBE is a *local* procedure. According to the definitions in Section 2.4.2 and in contrast to `MPI_IPROBE`, it is a *nonblocking* procedure because it is the *initialization* of a *matched receive* operation.

A *matched receive* (`MPI_MRECV` or `MPI_IMRECV`) executed with the *message handle* will receive the message that was matched by the *matching probe*. Unlike `MPI_IPROBE`, no other probe or receive operation may match the message returned by `MPI_IMPROBE`. Each *message handle* returned by `MPI_IMPROBE` must be received with either `MPI_MRECV` or `MPI_IMRECV`.

The `source` argument of `MPI_IMPROBE` can be `MPI_ANY_SOURCE`, and the `tag` argument can be `MPI_ANY_TAG`, so that one can *probe* for *messages* from an arbitrary source and/or with an arbitrary tag. However, a specific communication context must be provided with the `comm` argument.

A *synchronous mode send* operation that is matched with `MPI_IMPROBE` or `MPI_MPROBE` will *complete* successfully only if both a *matching receive* is started with `MPI_MRECV` or `MPI_IMRECV`, and the *matching receive* operation has started to receive the message sent by the *synchronous mode send*.

There is a special **predefined message handle**: `MPI_MESSAGE_NO_PROC`, which is a message that has `MPI_PROC_NULL` as its source. The predefined constant `MPI_MESSAGE_NULL` is the value used for **invalid message handles**.

A *matching probe* with `source = MPI_PROC_NULL` returns `flag = true`, `message = MPI_MESSAGE_NO_PROC`, and the status object returns `source = MPI_PROC_NULL`, `tag = MPI_ANY_TAG`, and `count = 0`; see Section 3.10. It is not necessary to call `MPI_MRECV` or `MPI_IMRECV` with `MPI_MESSAGE_NO_PROC`, but it is not *erroneous* to do so.

Rationale. `MPI_MESSAGE_NO_PROC` was chosen instead of `MPI_MESSAGE_PROC_NULL` to avoid possible confusion as another null handle constant. (*End of rationale.*)

<code>MPI_MPROBE(source, tag, comm, message, status)</code>			
IN	source	rank of source or <code>MPI_ANY_SOURCE</code> (integer)	
IN	tag	message tag or <code>MPI_ANY_TAG</code> (integer)	
IN	comm	communicator (handle)	
OUT	message	returned message (handle)	
OUT	status	status object (status)	

C binding

```
int MPI_Mprobe(int source, int tag, MPI_Comm comm, MPI_Message *message,
               MPI_Status *status)
```

Fortran 2008 binding

```
MPI_Mprobe(source, tag, comm, message, status, ierror)
  INTEGER, INTENT(IN) :: source, tag
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Message), INTENT(OUT) :: message
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_MPROBE(SOURCE, TAG, COMM, MESSAGE, STATUS, IERROR)
  INTEGER SOURCE, TAG, COMM, MESSAGE, STATUS(MPI_STATUS_SIZE), IERROR
```

`MPI_MPROBE` behaves like `MPI_IMPROBE` except that it is a *blocking* call that returns only after a matching message has been found.

The implementation of `MPI_MPROBE` and `MPI_IMPROBE` needs to guarantee *progress* in the same way as in the case of `MPI_PROBE` and `MPI_IPROBE`. See also Section 2.9 on *progress*.

According to the definitions in Section 2.4.2, `MPI_MPROBE` is *incomplete*. It is also a *nonlocal* procedure.

1 *Advice to users.* This is one of the exceptions in which *incomplete* procedures are
 2 *nonlocal.* (End of advice to users.)

3.8.3 Matched Receives

6 The **matched receive** operation (MPI_MRECV and MPI_IMRECV procedures) receive *mes-*
 7 *sages* that have been previously matched by a *matching probe* operation (Section 3.8.2).

9 MPI_MRECV(buf, count, datatype, message, status)

11	OUT	buf	initial address of receive buffer (choice)
12	IN	count	number of elements in receive buffer (non-negative integer)
13			
14	IN	datatype	datatype of each receive buffer element (handle)
15			
16	INOUT	message	message (handle)
17	OUT	status	status object (status)

19 C binding

20 int MPI_Mrecv(void *buf, int count, MPI_Datatype datatype,
 21 MPI_Message *message, MPI_Status *status)

23 int MPI_Mrecv_c(void *buf, MPI_Count count, MPI_Datatype datatype,
 24 MPI_Message *message, MPI_Status *status)

25 Fortran 2008 binding

26 MPI_Mrecv(buf, count, datatype, message, status, ierror)

27 TYPE(*), DIMENSION(..) :: buf
 28 INTEGER, INTENT(IN) :: count
 29 TYPE(MPI_Datatype), INTENT(IN) :: datatype
 30 TYPE(MPI_Message), INTENT(INOUT) :: message
 31 TYPE(MPI_Status) :: status
 32 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

34 MPI_Mrecv(buf, count, datatype, message, status, ierror) !(_c)

35 TYPE(*), DIMENSION(..) :: buf
 36 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
 37 TYPE(MPI_Datatype), INTENT(IN) :: datatype
 38 TYPE(MPI_Message), INTENT(INOUT) :: message
 39 TYPE(MPI_Status) :: status
 40 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

41 Fortran binding

42 MPI_MRECV(BUF, COUNT, DATATYPE, MESSAGE, STATUS, IERROR)

43 <type> BUF(*)
 44 INTEGER COUNT, DATATYPE, MESSAGE, STATUS(MPI_STATUS_SIZE), IERROR

45 This call receives a message matched by a *matching probe* operation (Section 3.8.2).

46 The *receive buffer* consists of the storage containing **count** consecutive elements of the
 47 type specified by **datatype**, starting at address **buf**. The length of the received message must
 48

be less than or equal to the length of the receive buffer. An overflow error occurs if all incoming data does not fit, without truncation, into the receive buffer.

If the message is shorter than the receive buffer, then only those locations corresponding to the (shorter) message are modified.

On return from this function, the *message handle* is set to MPI_MESSAGE_NULL. All errors that occur during the execution of this operation are handled according to the error handler set for the communicator used in the matching probe call that produced the message handle.

If MPI_MRECV is called with MPI_MESSAGE_NO_PROC as the message argument, the call returns immediately with the status object set to source = MPI_PROC_NULL, tag = MPI_ANY_TAG, and count = 0. This is consistent with the status object produced by a call to MPI_RECV or to MPI_PROBE with source = MPI_PROC_NULL (see Section 3.10). A call to MPI_MRECV with MPI_MESSAGE_NULL is *erroneous*.

MPI_IMRECV(buf, count, datatype, message, request)

OUT	buf	initial address of receive buffer (choice)
IN	count	number of elements in receive buffer (non-negative integer)
IN	datatype	datatype of each receive buffer element (handle)
INOUT	message	message (handle)
OUT	request	communication request (handle)

C binding

```
int MPI_Imrecv(void *buf, int count, MPI_Datatype datatype,
               MPI_Message *message, MPI_Request *request)
```

```
int MPI_Imrecv_c(void *buf, MPI_Count count, MPI_Datatype datatype,
                 MPI_Message *message, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Imrecv(buf, count, datatype, message, request, ierror)
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
```

```
INTEGER, INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Message), INTENT(INOUT) :: message
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Imrecv(buf, count, datatype, message, request, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Message), INTENT(INOUT) :: message
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_IMRECV(BUF, COUNT, DATATYPE, MESSAGE, REQUEST, IERROR)
```

1 <type> BUF(*)
2 INTEGER COUNT, DATATYPE, MESSAGE, REQUEST, IERROR

3
4 MPI_IMRECV is the nonblocking variant of MPI_MRECV and starts a nonblocking
5 receive of a matched message. Completion semantics are similar to MPI_Irecv as described
6 in Section 3.7.2. On return from this function, the *message handle* is set to
7 MPI_MESSAGE_NULL.

8 If MPI_IMRECV is called with MPI_MESSAGE_NO_PROC as the message argument, the
9 call returns immediately with a request object that, when completed, will yield a status
10 object set to source = MPI_PROC_NULL, tag = MPI_ANY_TAG, and count = 0, as if a receive
11 from MPI_PROC_NULL was issued (see Section 3.10). A call to MPI_IMRECV with
12 MPI_MESSAGE_NULL is *erroneous*.

13 *Advice to implementors.* If reception of a matched message is started with
14 MPI_IMRECV, then it is possible to *cancel* the returned request with MPI_CANCEL. If
15 MPI_CANCEL succeeds, the matched message must be found by a subsequent message
16 probe (MPI_PROBE, MPI_Iprobe, MPI_Mprobe, or MPI_Improbe), received by a
17 subsequent receive operation or *cancelled* by the sender. See Section 3.8.4 for details
18 about MPI_CANCEL. The *cancellation* of operations initiated with MPI_IMRECV may
19 fail. (*End of advice to implementors.*)

21 3.8.4 Cancel

22
23
24
25 MPI_CANCEL(request)

26 IN request communication request (handle)

28 C binding

29 int MPI_Cancel(MPI_Request *request)

31 Fortran 2008 binding

32 MPI_Cancel(request, ierror)

33 TYPE(MPI_Request), INTENT(IN) :: request

34 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

35 Fortran binding

36 MPI_CANCEL (REQUEST, IERROR)

37 INTEGER REQUEST, IERROR

38
39 A call to MPI_CANCEL marks for *cancellation* a *pending, nonblocking* communica-
40 tion operation (send or receive). *Cancelling* a send request by calling MPI_CANCEL is
41 deprecated. The *cancel* call is *local*. It returns *immediately*, possibly before the communi-
42 cation is actually *cancelled*. It is still necessary to call MPI_REQUEST_FREE, MPI_WAIT or
43 MPI_TEST (or any of the derived procedures) with the *cancelled* request as argument after
44 the call to MPI_CANCEL. If a communication is marked for *cancellation*, then a MPI_WAIT
45 call for that communication is guaranteed to return, irrespective of the activities of other
46 MPI processes (i.e., MPI_WAIT behaves as a *local* function); similarly if MPI_TEST is re-
47 peatedly called for a *cancelled* communication, then MPI_TEST will eventually return flag

= true.

MPI_CANCEL can be used to *cancel* a communication that uses a *persistent communication request* (see Section 3.9), in the same way as it is described above for nonblocking operations. *Cancelling* a persistent send request by calling MPI_CANCEL is deprecated. A successful *cancellation cancels* the *active* communication, but not the request itself. After the call to MPI_CANCEL and the subsequent call to MPI_WAIT or MPI_TEST, the request becomes *inactive* and can be activated for a new communication.

The successful *cancellation* of a *buffered mode send* frees the buffer space occupied by the pending message. *Cancelling a buffered mode send* request by calling MPI_CANCEL is deprecated.

Either the *cancellation* succeeds, or the communication succeeds, but not both. If a send is marked for *cancellation*, which is deprecated, then it must be the case that either the send *completes* normally, in which case the message sent was received at the destination, or that the send is successfully *cancelled*, in which case no part of the message was received at the destination. Then, any matching receive has to be satisfied by another send. If a receive is marked for *cancellation*, then it must be the case that either the receive *completes* normally, or that the receive is successfully *cancelled*, in which case no part of the receive buffer is altered. Then, any matching send has to be satisfied by another receive.

If the operation has been *cancelled*, then information to that effect will be returned in the status argument of the operation that *completes* the communication.

Rationale. Although the IN request handle parameter should not need to be passed by reference, the C binding has listed the argument type as MPI_Request* since MPI-1.0. This function signature therefore cannot be changed without breaking existing MPI applications. (*End of rationale.*)

MPI_TEST_CANCELLED(status, flag)

IN	status	status object (status)
OUT	flag	true if the operation has been cancelled (logical)

C binding

int MPI_Test_cancelled(const MPI_Status *status, int *flag)

Fortran 2008 binding

```
MPI_Test_cancelled(status, flag, ierror)
  TYPE(MPI_Status), INTENT(IN) :: status
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TEST_CANCELLED(STATUS, FLAG, IERROR)
  INTEGER STATUS(MPI_STATUS_SIZE), IERROR
  LOGICAL FLAG
```

Returns flag = true if the communication associated with the status object was *cancelled* successfully. In such a case, all other fields of status (such as count or tag) are undefined. Returns flag = false, otherwise. If a receive operation might be *cancelled* then one should

1 call `MPI_TEST_CANCELLED` first, to check whether the operation was *cancelled*, before
 2 checking on the other fields of the return status.

3
 4 *Advice to users.* *Cancel* can be an expensive operation that should be used only
 5 exceptionally. (*End of advice to users.*)

6
 7 *Advice to implementors.* If a send operation uses an “eager” protocol (data is
 8 transferred to the receiver before a matching receive is *started*), then the *cancellation*
 9 of this send may require communication with the intended receiver in order to free
 10 allocated buffers. On some systems this may require an interrupt to the intended
 11 receiver. Note that, while communication may be needed to implement `MPI_CANCEL`,
 12 this is still a *local* procedure, since its completion does not depend on the code executed
 13 by other MPI processes. If processing is required on another MPI process, this should
 14 be transparent to the application (hence the need for an interrupt and an interrupt
 15 handler). See also Section 2.9 on *progress*. (*End of advice to implementors.*)

17 3.9 Persistent Communication Requests

18
 19 Often a communication with the same argument list (with the exception of the buffer con-
 20 tents) is repeatedly executed within the inner loop of a parallel computation. In such a
 21 situation, it may be possible to optimize the communication by binding the list of commu-
 22 nication arguments to a *persistent communication request* once and then repeatedly using
 23 the request to *start* and *complete* operations. In the case of point-to-point communication,
 24 the *persistent communication request* thus created can be thought of as a communication
 25 port or a “half-channel.” It does not provide the full functionality of a conventional chan-
 26 nel, since there is no binding of the send port to the receive port. This construct allows
 27 reduction of the overhead for communication between the MPI process and communication
 28 controller, but not of the overhead for communication between one communication con-
 29 troller and another. It is not necessary that messages sent with a persistent point-to-point
 30 request be received by a receive operation using a persistent point-to-point request, or vice
 31 versa.

32 There are also persistent collective communication operations defined in Section 6.13
 33 and Section 8.8. The remainder of this section covers the point-to-point persistent *ini-*
 34 *tialization* operations and the start routines, which are used for persistent point-to-point,
 35 partitioned point-to-point, and persistent collective communication operations.

36 A point-to-point **persistent communication request** is created using one of the five
 37 following calls. These point-to-point persistent *initialization* calls involve no communica-
 38 tion.

39
 40
 41 `MPI_SEND_INIT(buf, count, datatype, dest, tag, comm, request)`

42	IN	<code>buf</code>	initial address of send buffer (choice)
43	IN	<code>count</code>	number of elements sent (non-negative integer)
44	IN	<code>datatype</code>	type of each element (handle)
45	IN	<code>dest</code>	rank of destination (integer)
46	IN	<code>dest</code>	rank of destination (integer)
47	IN	<code>tag</code>	message tag (integer)
48	IN	<code>tag</code>	message tag (integer)

IN	comm	communicator (handle)	1
OUT	request	communication request (handle)	2
			3
			4
C binding			5
int	MPI_Send_init(const void *buf, int count, MPI_Datatype datatype, int dest,		6
	int tag, MPI_Comm comm, MPI_Request *request)		7
int	MPI_Send_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype,		8
	int dest, int tag, MPI_Comm comm, MPI_Request *request)		9
			10
Fortran 2008 binding			11
MPI_Send_init(buf, count, datatype, dest, tag, comm, request, ierror)			12
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf			13
INTEGER, INTENT(IN) :: count, dest, tag			14
TYPE(MPI_Datatype), INTENT(IN) :: datatype			15
TYPE(MPI_Comm), INTENT(IN) :: comm			16
TYPE(MPI_Request), INTENT(OUT) :: request			17
INTEGER, OPTIONAL, INTENT(OUT) :: ierror			18
MPI_Send_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)			19
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf			20
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count			21
TYPE(MPI_Datatype), INTENT(IN) :: datatype			22
INTEGER, INTENT(IN) :: dest, tag			23
TYPE(MPI_Comm), INTENT(IN) :: comm			24
TYPE(MPI_Request), INTENT(OUT) :: request			25
INTEGER, OPTIONAL, INTENT(OUT) :: ierror			26
			27
Fortran binding			28
MPI_SEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)			29
<type> BUF(*)			30
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR			31
Creates a <i>persistent communication request</i> for a <i>standard mode send</i> operation.			32
			33
			34
MPI_BSEND_INIT(buf, count, datatype, dest, tag, comm, request)			35
IN	buf	initial address of send buffer (choice)	36
IN	count	number of elements sent (non-negative integer)	37
IN	datatype	type of each element (handle)	38
IN	dest	rank of destination (integer)	39
IN	tag	message tag (integer)	40
IN	comm	communicator (handle)	41
IN	comm	communicator (handle)	42
OUT	request	communication request (handle)	43
			44
			45
C binding			46
int	MPI_Bsend_init(const void *buf, int count, MPI_Datatype datatype, int dest,		47
	int tag, MPI_Comm comm, MPI_Request *request)		48

```

1 int MPI_Bsend_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
2                     int dest, int tag, MPI_Comm comm, MPI_Request *request)
3

```

Fortran 2008 binding

```

4 MPI_Bsend_init(buf, count, datatype, dest, tag, comm, request, ierror)
5     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
6     INTEGER, INTENT(IN) :: count, dest, tag
7     TYPE(MPI_Datatype), INTENT(IN) :: datatype
8     TYPE(MPI_Comm), INTENT(IN) :: comm
9     TYPE(MPI_Request), INTENT(OUT) :: request
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12 MPI_Bsend_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
13    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
14    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
15    TYPE(MPI_Datatype), INTENT(IN) :: datatype
16    INTEGER, INTENT(IN) :: dest, tag
17    TYPE(MPI_Comm), INTENT(IN) :: comm
18    TYPE(MPI_Request), INTENT(OUT) :: request
19    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20

```

Fortran binding

```

21 MPI_BSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
22     <type> BUF(*)
23     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
24
25     Creates a persistent communication request for a buffered mode send operation.
26

```

```

27 MPI_SSEND_INIT(buf, count, datatype, dest, tag, comm, request)
28
29     IN      buf          initial address of send buffer (choice)
30     IN      count       number of elements sent (non-negative integer)
31     IN      datatype    type of each element (handle)
32     IN      dest        rank of destination (integer)
33     IN      tag         message tag (integer)
34     IN      comm        communicator (handle)
35     IN      request     communication request (handle)
36
37
38

```

C binding

```

39 int MPI_Ssend_init(const void *buf, int count, MPI_Datatype datatype, int dest,
40                  int tag, MPI_Comm comm, MPI_Request *request)
41
42 int MPI_Ssend_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
43                     int dest, int tag, MPI_Comm comm, MPI_Request *request)
44

```

Fortran 2008 binding

```

45 MPI_Ssend_init(buf, count, datatype, dest, tag, comm, request, ierror)
46     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
47     INTEGER, INTENT(IN) :: count, dest, tag
48

```

```

TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ssend_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: dest, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
Fortran binding
MPI_SSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
<type> BUF(*)
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
Creates a persistent communication request for a synchronous mode send operation.
MPI_RSEND_INIT(buf, count, datatype, dest, tag, comm, request)
IN      buf          initial address of send buffer (choice)
IN      count        number of elements sent (non-negative integer)
IN      datatype     type of each element (handle)
IN      dest         rank of destination (integer)
IN      tag          message tag (integer)
IN      comm         communicator (handle)
OUT     request      communication request (handle)
C binding
int MPI_Rsend_init(const void *buf, int count, MPI_Datatype datatype, int dest,
                  int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Rsend_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
                    int dest, int tag, MPI_Comm comm, MPI_Request *request)
Fortran 2008 binding
MPI_Rsend_init(buf, count, datatype, dest, tag, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, dest, tag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Rsend_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf

```

```

1     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
2     TYPE(MPI_Datatype), INTENT(IN) :: datatype
3     INTEGER, INTENT(IN) :: dest, tag
4     TYPE(MPI_Comm), INTENT(IN) :: comm
5     TYPE(MPI_Request), INTENT(OUT) :: request
6     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

8 MPI_RSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
9     <type> BUF(*)
10    INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR

```

Creates a *persistent communication request* for a *ready mode send* operation.

```

14 MPI_RECV_INIT(buf, count, datatype, source, tag, comm, request)

```

16	OUT	buf	initial address of receive buffer (choice)
17	IN	count	number of elements received (non-negative integer)
18	IN	datatype	type of each element (handle)
19	IN	source	rank of source or MPI_ANY_SOURCE (integer)
20	IN	tag	message tag or MPI_ANY_TAG (integer)
21	IN	comm	communicator (handle)
22	OUT	request	communication request (handle)

C binding

```

27 int MPI_Recv_init(void *buf, int count, MPI_Datatype datatype, int source,
28                 int tag, MPI_Comm comm, MPI_Request *request)

```

```

29 int MPI_Recv_init_c(void *buf, MPI_Count count, MPI_Datatype datatype,
30                   int source, int tag, MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```

33 MPI_Recv_init(buf, count, datatype, source, tag, comm, request, ierror)

```

```

34     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
35     INTEGER, INTENT(IN) :: count, source, tag
36     TYPE(MPI_Datatype), INTENT(IN) :: datatype
37     TYPE(MPI_Comm), INTENT(IN) :: comm
38     TYPE(MPI_Request), INTENT(OUT) :: request
39     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

40 MPI_Recv_init(buf, count, datatype, source, tag, comm, request, ierror) !(_c)

```

```

41     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
42     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
43     TYPE(MPI_Datatype), INTENT(IN) :: datatype
44     INTEGER, INTENT(IN) :: source, tag
45     TYPE(MPI_Comm), INTENT(IN) :: comm
46     TYPE(MPI_Request), INTENT(OUT) :: request
47     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48

```

Fortran binding

```
MPI_RECV_INIT(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR)
  <type> BUF(*)
  INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR
```

Creates a *persistent communication request* for a receive operation. The argument `buf` is marked as OUT because the user gives permission to write on the receive buffer by passing the argument to `MPI_RECV_INIT`.

A *persistent communication request* is *inactive* after it was created—no active communication is attached to the request.

A communication that uses a *persistent communication request* is *started* by the function `MPI_START`.

```
MPI_START(request)
```

```
  INOUT  request                communication request (handle)
```

C binding

```
int MPI_Start(MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Start(request, ierror)
  TYPE(MPI_Request), INTENT(INOUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_START(REQUEST, IERROR)
  INTEGER REQUEST, IERROR
```

The argument, `request`, is a handle returned by any of the *initialization* procedures for persistent point-to-point communication (the previous five procedures), or for partitioned point-to-point communication (see Section 4), or for persistent collective communication (see Sections 6.13 and 8.8). The associated request should be *inactive*. The request becomes *active* once the call is made.

If the request is for a *ready mode send* operation, then a matching receive operation should be *started* before the call is made. The communication buffer should not be modified after the call, and until the operation *completes*.

The call is *local*, with similar semantics to the nonblocking communication operations described in Section 3.7. That is, a call to `MPI_START` with a request created by `MPI_SEND_INIT` starts a communication in the same manner as a call to `MPI_ISEND`; a call to `MPI_START` with a request created by `MPI_BSEND_INIT` starts a communication in the same manner as a call to `MPI_IBSEND`; and so on.

```
MPI_STARTALL(count, array_of_requests)
```

```
  IN      count                list length (non-negative integer)
  INOUT  array_of_requests    array of requests (array of handles)
```

C binding

```
int MPI_Startall(int count, MPI_Request array_of_requests[])
```

Fortran 2008 binding

```

MPI_Startall(count, array_of_requests, ierror)
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_STARTALL(COUNT, ARRAY_OF_REQUESTS, IERROR)
  INTEGER COUNT, ARRAY_OF_REQUESTS(*), IERROR

```

The execution of `MPI_STARTALL` has the same effect as the execution of `MPI_START` for each of the array elements in some arbitrary order. `MPI_STARTALL` with an array of length one is equivalent to `MPI_START`.

A communication started with a call to `MPI_START` or `MPI_STARTALL` is completed by a call to `MPI_WAIT`, `MPI_TEST`, or one of the derived functions described in Section 3.7.5. The request becomes *inactive* after successful completion of such call. The request is not deallocated and it can be activated anew by an `MPI_START` or `MPI_STARTALL` call.

A *persistent communication request* is deallocated by a call to `MPI_REQUEST_FREE` (Section 3.7.3). The call to `MPI_REQUEST_FREE` can occur at any point in the program after the persistent request was created. However, the request will be deallocated only after it becomes *inactive*. *Active* receive requests should not be *freed*. Otherwise, it will not be possible to check that the receive has *completed*. *Collective* operation requests (defined in Section 6.12 and Section 8.7 for nonblocking collective operations, and Section 6.13 and Section 8.8 for persistent collective operations) must not be *freed* while *active*. It is preferable, in general, to free requests when they are inactive. If this rule is followed, then the functions described in this section will be invoked in a sequence of the form,

Create (Start Complete)* Free

where *** indicates zero or more repetitions. If the same *persistent communication request* is used in several concurrent threads, it is the user's responsibility to coordinate calls so that the correct sequence is obeyed.

Inactive persistent requests are not automatically *freed* when the associated communicator is disconnected (via `MPI_COMM_DISCONNECT`, see 11.10.4) or the associated World Model or Sessions Model is finalized (via `MPI_FINALIZE`, see 11.2.2, or `MPI_SESSION_FINALIZE`, see 11.3.1). In these situations, any further use of the request handle is erroneous. In particular, freeing associated inactive request handles after such a communicator disconnect or finalization is then impossible.

Advice to users. Persistent request handles may bind internal resources such as MPI buffers in shared memory for providing efficient communication. Therefore, it is highly recommended to explicitly free inactive request handles, using `MPI_REQUEST_FREE`, when they are no longer in use, and in particular before freeing or disconnecting the associated communicator with `MPI_COMM_FREE` or `MPI_COMM_DISCONNECT` or finalizing the associated session with `MPI_SESSION_FINALIZE`. (*End of advice to users.*)

A send operation *started* with `MPI_START` can be *matched* with any receive operation and, likewise, a receive operation *started* with `MPI_START` can receive messages generated by any send operation.

Advice to users. To prevent problems with the argument copying and register optimization done by Fortran compilers, please note the hints in Sections 19.1.10–19.1.20. (*End of advice to users.*)

3.10 Null MPI Processes

In many instances, it is convenient to specify a “dummy” source or destination for communication. This simplifies the code that is needed for dealing with boundaries, for example, in the case of a noncircular shift done with calls to send-receive.

The special value `MPI_PROC_NULL` can be used instead of a rank wherever a source or a destination argument is required in a call. A communication with `MPI_PROC_NULL` has no effect. A send to `MPI_PROC_NULL` succeeds and returns as soon as possible. A receive from `MPI_PROC_NULL` succeeds and returns as soon as possible with no modifications to the receive buffer. When a receive with `source = MPI_PROC_NULL` is executed then the status object returns `source = MPI_PROC_NULL`, `tag = MPI_ANY_TAG` and `count = 0`. A probe or matching probe with `source = MPI_PROC_NULL` succeeds and returns as soon as possible, and the status object returns `source = MPI_PROC_NULL`, `tag = MPI_ANY_TAG` and `count = 0`. A matching probe (cf. Section 3.8.2) with `source = MPI_PROC_NULL` returns `flag = true`, `message = MPI_MESSAGE_NO_PROC`, and the status object returns `source = MPI_PROC_NULL`, `tag = MPI_ANY_TAG`, and `count = 0`.

DRAFT

Chapter 4

Partitioned Point-to-Point Communication

4.1 Introduction

Partitioned communication extends persistent point-to-point communication as defined in Chapter 3. Partitioned communication operations are matched based on the order in which the local initialization calls are performed. Partitioned communication is “partitioned” because it allows for multiple contributions of data to be made, potentially, from multiple actors (e.g., threads or tasks) in an MPI process to a single communication operation.

Advice to users. The techniques of partitioned communication were known as “fine-points” before their adoption into the MPI standard. We refer the interested reader to the original literature describing the design goals, functioning, initial implementation and performance improvements [29, 30]. (*End of advice to users.*)

Partitioned communication operations use a persistent communication style that involves a sequence of start and test or wait operations. For this sequence, partitioned communications use `MPI_START` or `MPI_STARTALL` calls and completion mechanisms (e.g., `MPI_TEST` or `MPI_WAIT`). Partitioned communication is different in three fundamental ways from persistent point-to-point operations in MPI. First, partitioned communication allows additional partitioned test function calls that can expose partial completion of the operation. Second, partitioned communication may perform all of the initialization required to enable data transfer as early as its initialization phase. Third, partitioned communication allows for MPI to be independently notified of multiple contributions from the send-side to a single data buffer of a single MPI message.

Rationale. The rationale behind having different initialization behavior allowed for partitioned communication as opposed to persistent point-to-point communication is to enable flexibility and optimization possibilities in implementations. Buffer setup can occur in the partitioned communication initialization functions (see Section 4.2.1). However, such negotiation can be deferred until data is to be moved between two processes. This means that partitioned communication can lazily negotiate as late as testing for completion of the operation on the first iteration of a sequence of partitioned communication start and test or wait operations. Matching still occurs as if matching happened at the partitioned communication initialization functions as noted in the function descriptions. (*End of rationale.*)

4.2 Semantics of Partitioned Point-to-Point Communication

MPI guarantees certain general properties of partitioned point-to-point communication progress, which are described in this section.

Persistent communications use opaque `MPI_REQUEST` objects as described in Section 3. Partitioned communication uses these same semantics for `MPI_REQUEST` objects.

Partitioned communication provides fine-grained transfers on either or both sides of a send-receive operation described by requests. Persistent communication semantics are ideal for partitioned communication: they provide `MPI_PSEND_INIT` and `MPI_PRECV_INIT` functions that allow partitioned communication setup to occur prior to message transfers. Partitioned communication initialization functions are local. The partitioned communication initialization includes inputs on the number of user-visible partitions on the send-side and receive-side, which may differ. Valid partitioned communication operations must have one or more partitions specified.

Once an `MPI_PSEND_INIT` call has been made, the user may start the operation with a call to a starting procedure and complete the operation with a number of `MPI_PREADY` calls equal to the requested number of send partitions followed by a call to a completing procedure. A call to `MPI_PREADY` notifies the MPI library that a specified portion of the data buffer (a specific partition) is ready to be sent. Notification of partial completion can be done via fine-grained `MPI_PARRIVED` calls at the receiver before a final `MPI_TEST/MPI_WAIT` on the request itself; the latter represents overall operation completion upon success. A full set of methods for starting and completing partitioned communication is given in the following sections.

Advice to users. Having a large number of receiver-side partitions can increase overheads as the completion mechanism may need to work with finer-grained notifications. Using a small number of receiver-side partitions *may* provide higher performance.

A large number of sender-side partitions may be aggregated by an MPI implementation, making performance concerns of a large number of sender-side partitions potentially less impactful than receiver-side granularity. (*End of advice to users.*)

Advice to implementors. It is expected that an MPI implementation will attempt to balance latency and aggregation for data transfers for the requested partition counts on the sender-side and receiver-side to allow optimization for different hardware. A high quality implementation may perform significant optimizations to enhance performance in this way; they may, for example, resize the data transfers of the partitions to combine partitions in fractional partition sizes (e.g., 2.5 partitions in a single data transfer). (*End of advice to implementors.*)

Example 4.1 shows a simple partitioned transfer in which the sender-side and receiver-side partitioning is identical in partition count.

Example 4.1. Simple partitioned communication example.

```
#include <stdlib.h>
#include "mpi.h"
#define PARTITIONS 8
#define COUNT 5
int main(int argc, char *argv[])
{
```

```

1  double message[PARTITIONS*COUNT];
2  MPI_Count partitions = PARTITIONS;
3  int source = 0, dest = 1, tag = 1, flag = 0;
4  int myrank, i;
5  int provided;
6  MPI_Request request;
7  MPI_Init_thread(&argc, &argv, MPI_THREAD_SERIALIZED, &provided);
8  if (provided < MPI_THREAD_SERIALIZED)
9      MPI_Abort(MPI_COMM_WORLD, EXIT_FAILURE);
10 MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
11 if (myrank == 0)
12 {
13     MPI_Psend_init(message, partitions, COUNT, MPI_DOUBLE, dest, tag,
14                   MPI_COMM_WORLD, MPI_INFO_NULL, &request);
15     MPI_Start(&request);
16     for(i = 0; i < partitions; ++i)
17     {
18         /* compute and fill partition #i, then mark ready: */
19         MPI_Pready(i, request);
20     }
21     while(!flag)
22     {
23         /* do useful work #1 */
24         MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
25         /* do useful work #2 */
26     }
27     MPI_Request_free(&request);
28 }
29 else if (myrank == 1)
30 {
31     MPI_Precv_init(message, partitions, COUNT, MPI_DOUBLE, source, tag,
32                   MPI_COMM_WORLD, MPI_INFO_NULL, &request);
33     MPI_Start(&request);
34     while(!flag)
35     {
36         /* do useful work #1 */
37         MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
38         /* do useful work #2 */
39     }
40     MPI_Request_free(&request);
41 }
42 MPI_Finalize();
43 return 0;
44 }

```

Rationale. Partitioned communication is designed to provide opportunities for MPI implementations to optimize data transfers. MPI is free to choose how many transfers to do within a partitioned communication send independent of how many partitions are reported as ready to MPI through MPI_PREADY calls. Aggregation of partitions is permitted but not required. Ordering of partitions is permitted but not required. A naive implementation can simply wait for the entire message buffer to be marked ready before any transfer(s) occur and could wait until the completion function is

1 called on a request before transferring data. However, this modality of communication
 2 gives MPI implementations far more flexibility in data movement than nonpartitioned
 3 communications. (*End of rationale.*)
 4

5 4.2.1 Communication Initialization and Starting with Partitioning

6 Initialization of partitioned communication operations use the initialization calls described
 7 below. Subsequent to initialization, MPI_START/MPI_STARTALL are used as the first
 8 indication to MPI that a message transfer will occur. For send-side operations, neither
 9 initializing nor starting the operation enables transfer of any part of the user buffer. Freeing
 10 or canceling a partitioned communication request that is active (i.e., initialized and started)
 11 and not completed is erroneous. After the partitioned communication operation is started,
 12 individual partitions of a message are indicated as ready to be sent by MPI via the
 13 MPI_PREADY function, described below.
 14

15
 16 MPI_PSEND_INIT(buf, partitions, count, datatype, dest, tag, comm, info, request)

17	IN	buf	initial address of send buffer (choice)
18	IN	partitions	number of partitions (non-negative integer)
19	IN	count	number of elements sent per partition (non-negative integer)
20	IN	datatype	type of each element (handle)
21	IN	dest	rank of destination (integer)
22	IN	tag	message tag (integer)
23	IN	comm	communicator (handle)
24	IN	info	info argument (handle)
25	OUT	request	communication request (handle)

31 C binding

32
 33 int MPI_Psend_init(const void *buf, int partitions, MPI_Count count,
 34 MPI_Datatype datatype, int dest, int tag, MPI_Comm comm,
 35 MPI_Info info, MPI_Request *request)

36 Fortran 2008 binding

37 MPI_Psend_init(buf, partitions, count, datatype, dest, tag, comm, info,
 38 request, ierror)
 39 TYPE(*), DIMENSION(..), INTENT(IN) :: buf
 40 INTEGER, INTENT(IN) :: partitions, dest, tag
 41 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
 42 TYPE(MPI_Datatype), INTENT(IN) :: datatype
 43 TYPE(MPI_Comm), INTENT(IN) :: comm
 44 TYPE(MPI_Info), INTENT(IN) :: info
 45 TYPE(MPI_Request), INTENT(OUT) :: request
 46 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
 47
 48

Fortran binding

```

MPI_PSEND_INIT(BUF, PARTITIONS, COUNT, DATATYPE, DEST, TAG, COMM, INFO,
               REQUEST, IERROR)
<type> BUF(*)
INTEGER PARTITIONS, DATATYPE, DEST, TAG, COMM, INFO, REQUEST, IERROR
INTEGER(KIND=MPI_COUNT_KIND) COUNT

```

MPI_PSEND_INIT creates a partitioned communication request and binds to it all the arguments of a partitioned send operation. Matching follows the same MPI matching rules as for point-to-point communication (see Chapter 3) with communicator, tag, and source dictating message matching. In the event that the communicator, tag, and source do not uniquely identify a message, the order in which partitioned communication *initialization* calls are made is the order in which they will eventually match. This operation can only match with partitioned communication initialization operations, therefore it is required to be matched with a corresponding MPI_PRECV_INIT call. Partitioned communication initialization calls are local. It is erroneous to provide a `partitions` value ≤ 0 . Send-side and receive-side buffers must be identical in size.

Advice to implementors. Unlike MPI_SEND_INIT, MPI_PSEND_INIT can be matched as early as the initialization call. Also, unlike MPI_SEND_INIT, MPI_PSEND_INIT takes an info argument. (*End of advice to implementors.*)

```

MPI_PRECV_INIT(buf, partitions, count, datatype, source, tag, comm, info, request)

```

IN	buf	initial address of recv buffer (choice)
IN	partitions	number of partitions (non-negative integer)
IN	count	number of elements received per partition (non-negative integer)
IN	datatype	type of each element (handle)
IN	source	rank of source (integer)
IN	tag	message tag (integer)
IN	comm	communicator (handle)
IN	info	info argument (handle)
OUT	request	communication request (handle)

C binding

```

int MPI_Precv_init(void *buf, int partitions, MPI_Count count,
                  MPI_Datatype datatype, int source, int tag, MPI_Comm comm,
                  MPI_Info info, MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Precv_init(buf, partitions, count, datatype, source, tag, comm, info,
               request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER, INTENT(IN) :: partitions, source, tag

```

```

1     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
2     TYPE(MPI_Datatype), INTENT(IN) :: datatype
3     TYPE(MPI_Comm), INTENT(IN) :: comm
4     TYPE(MPI_Info), INTENT(IN) :: info
5     TYPE(MPI_Request), INTENT(OUT) :: request
6     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

8     MPI_PRECV_INIT(BUF, PARTITIONS, COUNT, DATATYPE, SOURCE, TAG, COMM, INFO,
9                   REQUEST, IERROR)
10    <type> BUF(*)
11    INTEGER PARTITIONS, DATATYPE, SOURCE, TAG, COMM, INFO, REQUEST, IERROR
12    INTEGER(KIND=MPI_COUNT_KIND) COUNT

```

Rationale. The info argument is provided in order to support per-operation implementation-defined info keys. (*End of rationale.*)

MPI_PRECV_INIT creates a partitioned communication receive request and binds to it all the arguments of a partitioned receive operation. This operation can only match with partitioned communication initialization operations, therefore the MPI library is required to match MPI_PRECV_INIT calls only with a corresponding MPI_PSEND_INIT call. Matching follows the same MPI matching rules as for point-to-point communication (see Chapter 3) with communicator, tag, and source dictating message matching. In the event that the communicator, tag, and source do not uniquely identify a message, the order in which partitioned communication initialization calls are made is the order in which they will eventually match. Partitioned communication initialization calls are local. That is, MPI_PRECV_INIT may return before the operation completes. It is erroneous to provide a partitions value ≤ 0 . Wildcards for source and tag are not allowed.

Advice to implementors. Unlike MPI_RECV_INIT, MPI_PRECV_INIT may communicate. Also unlike MPI_RECV_INIT, MPI_PRECV_INIT takes an info argument. (*End of advice to implementors.*)

```

35    MPI_PREADY(partition, request)
36
37    IN          partition          partition to mark ready for transfer (non-negative
38                                     integer)
39
39    INOUT      request            partitioned communication request (handle)

```

C binding

```

42    int MPI_Pready(int partition, MPI_Request request)

```

Fortran 2008 binding

```

44    MPI_Pready(partition, request, ierror)
45    INTEGER, INTENT(IN) :: partition
46    TYPE(MPI_Request), INTENT(IN) :: request
47    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```


Fortran binding

```
MPI_PREADY(PARTITION, REQUEST, IERROR)
    INTEGER PARTITION, REQUEST, IERROR
```

MPI_PREADY is a send-side call that indicates that a given partition is ready to be transferred. It is erroneous to use MPI_PREADY on any request object that does not correspond to a partitioned send operation. The partitioning is defined by the MPI_PSEND_INIT call. Partition numbering starts at zero and ranges to one less than the number of partitions declared in the MPI_PSEND_INIT call. Specifying a partition number that is equal to or larger than the number of partitions is erroneous. After a call to MPI_START/MPI_STARTALL, all partitions associated with that operation are inactive. A call to MPI_PREADY marks the indicated partition as active. Calling MPI_PREADY on an active partition is erroneous.

```
MPI_PREADY_RANGE(partition_low, partition_high, request)
```

IN	partition_low	lowest partition ready for transfer (non-negative integer)
IN	partition_high	highest partition ready for transfer (non-negative integer)
INOUT	request	partitioned communication request (handle)

C binding

```
int MPI_Pready_range(int partition_low, int partition_high,
    MPI_Request request)
```

Fortran 2008 binding

```
MPI_Pready_range(partition_low, partition_high, request, ierror)
    INTEGER, INTENT(IN) :: partition_low, partition_high
    TYPE(MPI_Request), INTENT(IN) :: request
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_PREADY_RANGE(PARTITION_LOW, PARTITION_HIGH, REQUEST, IERROR)
    INTEGER PARTITION_LOW, PARTITION_HIGH, REQUEST, IERROR
```

A call to MPI_PREADY_RANGE has the same effect as calls to MPI_PREADY, executed for $i = \text{partition_low}, \dots, \text{partition_high}$, in some arbitrary order. Calls to MPI_PREADY_RANGE follow the same rules as those for MPI_PREADY calls.

```
MPI_PREADY_LIST(length, array_of_partitions, request)
```

IN	length	list length (integer)
IN	array_of_partitions	array of partitions (array of non-negative integers)
INOUT	request	partitioned communication request (handle)

C binding

```
int MPI_Pready_list(int length, const int array_of_partitions[],
    MPI_Request request)
```

Fortran 2008 binding

```

MPI_Pready_list(length, array_of_partitions, request, ierror)
  INTEGER, INTENT(IN) :: length, array_of_partitions(length)
  TYPE(MPI_Request), INTENT(IN) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_PREADY_LIST(LENGTH, ARRAY_OF_PARTITIONS, REQUEST, IERROR)
  INTEGER LENGTH, ARRAY_OF_PARTITIONS(*), REQUEST, IERROR

```

A call to `MPI_PREADY_LIST` has the same effect as calls to `MPI_PREADY`, executed for the partitions specified in the range `array_of_partitions[0]` `, ..., array_of_partitions[count - 1]` of the `array_of_partitions`, executed in some arbitrary order. Calls to `MPI_PREADY_LIST` follow the same rules as those for `MPI_PREADY` calls.

4.2.2 Communication Completion under Partitioning

The functions `MPI_WAIT` and `MPI_TEST` (and variants) are used to complete a partitioned communication operation. The completion of a partitioned send operation indicates that the sender is now free to call `MPI_START/MPI_STARTALL` to restart the operation and subsequently `MPI_PREADY`, `MPI_PREADY_RANGE` or `MPI_PREADY_LIST`. Alternatively, the user can safely free the partitioned communication request after the completion of the partitioned operation. For the sending process, completion of the partitioned send operation does not indicate that the partitions of the message have all been received.

The completion of a partitioned receive operation through `MPI_WAIT` or `MPI_TEST` indicates that the receive buffer contains all of the partitions. A function for probing the partial reception of the receive buffer is provided by `MPI_PARRIVED`. The `MPI_PARRIVED` function can be used to determine if the message data for the indicated partition has been received into the receive buffer. Upon success, the receiver becomes free to access the indicated partition (as well as any others that previously completed for that operation).

```

MPI_PARRIVED(request, partition, flag)

```

IN	request	partitioned communication request (handle)
IN	partition	partition to be tested (non-negative integer)
OUT	flag	true if operation completed on the specified partition, false if not (logical)

C binding

```

int MPI_Parrived(MPI_Request request, int partition, int *flag)

```

Fortran 2008 binding

```

MPI_Parrived(request, partition, flag, ierror)
  TYPE(MPI_Request), INTENT(IN) :: request
  INTEGER, INTENT(IN) :: partition
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_PARRIVED(REQUEST, PARTITION, FLAG, IERROR)

```

INTEGER REQUEST, PARTITION, IERROR
LOGICAL FLAG

The function `MPI_PARRIVED` can be used to test partial completion of partitioned receive operations. A call to `MPI_PARRIVED` on an active partitioned communication request returns `flag = true` if the operation identified by `request` for the specified `partition` is complete. The request is not marked as complete/inactive by this procedure. A subsequent call to an MPI completing procedure (e.g., `MPI_TEST/MPI_WAIT`) is required to complete the operation, as described in Chapter 3. `MPI_PARRIVED` may be called multiple times for a partition. `MPI_PARRIVED` may be called with a null or inactive `request` argument. In either case, the operation returns with `flag = true`. Calling `MPI_PARRIVED` on a request that does not correspond to a partitioned receive operation is erroneous.

Repeated calls to `MPI_PARRIVED` with the same `request` and `partition` arguments will eventually return `flag = true` if the corresponding partitioned send operation has been started and all send partitions have been marked as ready. For additional information on MPI *progress* see Sections 2.9 and 3.7.4.

Advice to implementors. A high quality implementation will eventually return `flag = true` from `MPI_PARRIVED` after all of the corresponding `MPI_PREADY` calls have been made for a receive-side partition, even if other send partitions are not yet marked as ready. (*End of advice to implementors.*)

4.2.3 Semantics of Communications in Partitioned Mode

The semantics of nonblocking partitioned communication are defined by suitably extending the definitions in Section 3.5.

Interpretation of count and datatype for partitioned communication. Partitioned communication uses the `count` and `datatype` arguments in the partitioned communication initialization functions to describe a single partition. The argument `partitions` specifies how many equal partitions of a number (`count`) of objects of `datatypes` make up the entire buffer to be transferred in the partitioned communication. As partitioned communication describes many partitions, using absolute displacements in datatypes (e.g., `MPI_BOTTOM`) is not supported. Partitions are contiguous in memory, there is no padding in between them. Once a partitioned send operation is started, each partition must be marked as ready using `MPI_PREADY` and the operation must be completed using a completion function, such as `MPI_TEST` or `MPI_WAIT`.

Order. Matching follows the same MPI matching rules as for point-to-point communication (see Chapter 3) with communicator, tag, and source dictating message matching. In the event that the communicator, tag, and source do not uniquely identify the message, the order in which partitioned communication initialization calls are made is the order in which they will eventually match.

4.3 Partitioned Communication Examples

This section provides concrete examples of the utility of partitioned communication in realistic settings.

4.3.1 Partition Communication with Threads/Tasks Using OpenMP 4.0 or later

The equal partitioning on send-side and receive-side in Example 4.1 is shown using threads. In this case, the receive-side uses the same number of partitions as the send-side as in the previous example, but this example uses multiple threads on the send-side. Note that the `MPI_PSEND_INIT` and `MPI_PRECV_INIT` functions match each other like in the previous example.

Example 4.2. Equal partitioning on send-side and receive-side using threads.

```

1  #include <stdlib.h>
2  #include "mpi.h"
3  #define NUM_THREADS 8
4  #define PARTITIONS 8
5  #define PARTLENGTH 16
6  int main(int argc, char *argv[]) /* same send/recv partitioning */
7  {
8      double message[PARTITIONS*PARTLENGTH];
9      int partitions = PARTITIONS;
10     int partlength = PARTLENGTH;
11     int count = 1, source = 0, dest = 1, tag = 1, flag = 0;
12     int myrank;
13     int provided;
14     MPI_Request request;
15     MPI_Info info = MPI_INFO_NULL;
16     MPI_Datatype xfer_type;
17     MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
18     if (provided < MPI_THREAD_MULTIPLE)
19         MPI_Abort(MPI_COMM_WORLD, EXIT_FAILURE);
20     MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
21     MPI_Type_contiguous(partlength, MPI_DOUBLE, &xfer_type);
22     MPI_Type_commit(&xfer_type);
23     if (myrank == 0)
24     {
25         MPI_Psend_init(message, partitions, count, xfer_type, dest, tag,
26                        MPI_COMM_WORLD, info, &request);
27         MPI_Start(&request);
28
29         #pragma omp parallel for shared(request) num_threads(NUM_THREADS)
30         for (int i=0; i<partitions; i++)
31         {
32             /* compute and fill partition #i, then mark ready: */
33             MPI_Pready(i, request);
34         }
35         while(!flag)
36         {
37             /* Do useful work */
38             MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
39             /* Do useful work */
40         }
41         MPI_Request_free(&request);
42     }
43     else if (myrank == 1)
44     {
45

```

```

1  MPI_Precv_init(message, partitions, count, xfer_type, source, tag,
2  MPI_COMM_WORLD, info, &request);
3  MPI_Start(&request);
4  while(!flag)
5  {
6  /* Do useful work */
7  MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
8  /* Do useful work */
9  }
10 MPI_Request_free(&request);
11 }
12 MPI_Finalize();
13 return 0;
14 }

```

4.3.2 Send-only Partitioning Example with Tasks and OpenMP version 4.0 or later

The previous example is tailored specifically for send-side partitioning using threads. This is an example where parallel task producers produce input to part of an overall buffer; they complete in any order and contribute to the overall buffer.

Example 4.3. Parallel task producers for partitioned communication using threads.

```

22 #include <stdlib.h>
23 #include "mpi.h"
24 #define NUM_THREADS 8
25 #define NUM_TASKS 64
26 #define PARTITIONS NUM_TASKS
27 #define PARTLENGTH 16
28 #define MESSAGE_LENGTH PARTITIONS*PARTLENGTH
29 int main(int argc, char *argv[]) /* send-side partitioning */
30 {
31     double message[MESSAGE_LENGTH];
32     int send_partitions = PARTITIONS,
33     send_partlength = PARTLENGTH,
34     recv_partitions = 1,
35     recv_partlength = PARTITIONS*PARTLENGTH;
36     int count = 1, source = 0, dest = 1, tag = 1, flag = 0;
37     int myrank;
38     int provided;
39     MPI_Request request;
40     MPI_Info info = MPI_INFO_NULL;
41     MPI_Datatype send_type;
42     MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
43     if (provided < MPI_THREAD_MULTIPLE)
44         MPI_Abort(MPI_COMM_WORLD, EXIT_FAILURE);
45     MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
46     MPI_Type_contiguous(send_partlength, MPI_DOUBLE, &send_type);
47     MPI_Type_commit(&send_type);
48
49     if (myrank == 0)
50     {
51         MPI_Psend_init(message, send_partitions, count, send_type, dest, tag,
52         MPI_COMM_WORLD, info, &request);
53         MPI_Start(&request);

```

```

1
2     #pragma omp parallel shared(request) num_threads(NUM_THREADS)
3     {
4         #pragma omp single
5         {
6             /* single thread creates 64 tasks to be executed by 8 threads */
7             for (int partition_num=0;partition_num<NUM_TASKS;partition_num++)
8             {
9                 #pragma omp task firstprivate(partition_num)
10                {
11                    /* compute and fill partition #partition_num, then mark
12                     ready: */
13                    /* buffer is filled in arbitrary order from each task */
14                    MPI_Pready(partition_num, request);
15                } /*end task*/
16            } /* end for */
17        } /* end single */
18    } /* end parallel */
19    while(!flag)
20    {
21        /* Do useful work */
22        MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
23        /* Do useful work */
24    }
25    MPI_Request_free(&request);
26
27    else if (myrank == 1)
28    {
29        MPI_Precv_init(message, recv_partitions, recv_partlength, MPI_DOUBLE,
30                      source, tag, MPI_COMM_WORLD, info, &request);
31
32        MPI_Start(&request);
33        while(!flag)
34        {
35            /* Do useful work */
36            MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
37            /* Do useful work */
38        }
39        MPI_Request_free(&request);
40    }
41    MPI_Finalize();
42    return 0;
43 }
44
45
46
47
48

```

4.3.3 Send and Receive Partitioning Example with OpenMP version 4.0 or later

This example demonstrates receive-side partial completion notification using more than one partition per receive-side thread. It uses a naive flag based method to test for multiple completed partitions per thread. Note that this means that some threads may be busy polling for completion of assigned partitions when partitions are available to work on that were not assigned to the polling threads in this example. More advanced work stealing methods could be employed for greater efficiency. Like previous examples, it also demonstrates send-side production of input to part of an overall buffer. This example also uses different send-side and receive-side partitioning.

Example 4.4. Partitioned communication receive-side partial completion.

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
#include <stdlib.h>
#include "mpi.h"
#define NUM_THREADS 64
#define PARTITIONS NUM_THREADS
#define PARTLENGTH 16
#define MESSAGE_LENGTH PARTITIONS*PARTLENGTH
int main(int argc, char *argv[]) /* send-side partitioning */
{
    double message[MESSAGE_LENGTH];
    int send_partitions = PARTITIONS,
        send_partlength = PARTLENGTH,
        recv_partitions = PARTITIONS*2,
        recv_partlength = PARTLENGTH/2;
    int source = 0, dest = 1, tag = 1, flag = 0;
    int myrank;
    int provided;
    MPI_Request request;
    MPI_Info info = MPI_INFO_NULL;
    MPI_Datatype send_type;
    MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
    if (provided < MPI_THREAD_MULTIPLE)
        MPI_Abort(MPI_COMM_WORLD, EXIT_FAILURE);
    MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
    MPI_Type_contiguous(send_partlength, MPI_DOUBLE, &send_type);
    MPI_Type_commit(&send_type);

    if (myrank == 0)
    {
        MPI_Psend_init(message, send_partitions, 1, send_type, dest, tag,
                      MPI_COMM_WORLD, info, &request);
        MPI_Start(&request);
        #pragma omp parallel for shared(request) \
                               firstprivate(send_partitions) \
                               num_threads(NUM_THREADS)
        for (int i=0; i<send_partitions; i++)
        {
            /* compute and fill partition #i, then mark ready: */
            MPI_Pready(i, request);
        }
        while(!flag)
        {
            /* Do useful work */
            MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
            /* Do useful work */
        }
        MPI_Request_free(&request);
    }
    else if (myrank == 1)
    {
        MPI_Precv_init(message, recv_partitions, recv_partlength,

```

```
1      MPI_DOUBLE, source, tag, MPI_COMM_WORLD, info,
2      &request);
3  MPI_Start(&request);
4  #pragma omp parallel for shared(request) \
5      firstprivate(recv_partitions) \
6      num_threads(NUM_THREADS)
7  for (int j=0; j<recv_partitions; j+=2)
8  {
9      int part_flag = 0;
10     int part1_complete = 0;
11     int part2_complete = 0;
12     while(part1_complete == 0 || part2_complete == 0)
13     {
14         /* test partition #j and #j+1 */
15         MPI_Parrived(request, j, &part_flag);
16         if(part_flag && part1_complete == 0)
17         {
18             part1_complete++;
19             /* Do work using partition j data */
20         }
21         MPI_Parrived(request, j+1, &part_flag);
22         if(part_flag && part2_complete == 0)
23         {
24             part2_complete++;
25             /* Do work using partition j+1 */
26         }
27     }
28     while(!flag)
29     {
30         /* Do useful work */
31         MPI_Test(&request, &flag, MPI_STATUS_IGNORE);
32         /* Do useful work */
33     }
34     MPI_Request_free(&request);
35 }
36
37
38
39
40
41
42
43
44
45
46
47
48
```


Chapter 5

Datatypes

Basic datatypes are introduced in Section 3.2.2 and in Section 3.3. In this chapter, this model is extended to describe any data layout. We consider general datatypes that allow one to transfer efficiently heterogeneous and noncontiguous data. We conclude with the description of calls for explicit packing and unpacking of messages.

5.1 Derived Datatypes

Up to here, all point-to-point communications have involved only buffers containing a sequence of identical basic datatypes. This is too constraining on two accounts. One often wants to pass messages that contain values with different datatypes (e.g., an integer count, followed by a sequence of real numbers); and one often wants to send noncontiguous data (e.g., a sub-block of a matrix). One solution is to pack noncontiguous data into a contiguous buffer at the sender site and unpack it at the receiver site. This has the disadvantage of requiring additional memory-to-memory copy operations at both sites, even when the communication subsystem has scatter-gather capabilities. Instead, MPI provides mechanisms to specify more general, mixed, and noncontiguous communication buffers. It is up to the implementation to decide whether data should be first packed in a contiguous buffer before being transmitted, or whether it can be collected directly from where it resides.

The general mechanisms provided here allow one to transfer directly, without copying, objects of various shapes and sizes. It is not assumed that the MPI library is cognizant of the objects declared in the host language. Thus, if one wants to transfer a structure, or an array section, it will be necessary to provide in MPI a definition of a communication buffer that mimics the definition of the structure or array section in question. These facilities can be used by library designers to define communication functions that can transfer objects defined in the host language—by decoding their definitions as available in a symbol table or a dope vector. Such higher-level communication functions are not part of MPI.

More general communication buffers are specified by replacing the basic datatypes that have been used so far with derived datatypes that are constructed from basic datatypes using the constructors described in this section. These methods of constructing derived datatypes can be applied recursively.

A **general datatype** is an opaque object that specifies two things:

- A sequence of basic datatypes.
- A sequence of integer (byte) displacements.

The displacements are not required to be positive, distinct, or in increasing order. Therefore, the order of items need not coincide with their order in store, and an item may appear more than once. We call such a pair of sequences (or sequence of pairs) a **type**

map. The sequence of basic datatypes (displacements ignored) is the **type signature** of the datatype.

Let

$$Typemap = \{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

be such a type map, where $type_i$ are basic types, and $disp_i$ are displacements. Let

$$Typesig = \{type_0, \dots, type_{n-1}\}$$

be the associated type signature. This type map, together with a base address **buf**, specifies a communication buffer: the communication buffer that consists of n entries, where the i -th entry is at address $buf + disp_i$ and has type $type_i$. A message assembled from such a communication buffer will consist of n values, of the types defined by *Typesig*.

Most datatype constructors have replication count or block length arguments. Allowed values are nonnegative integers. If the value is zero, no elements are generated in the type map and there is no effect on datatype bounds or extent.

We can use a handle to a general datatype as an argument in a send or receive operation, instead of a basic datatype argument. The operation `MPI_SEND(buf, 1, datatype, ...)` will use the send buffer defined by the base address **buf** and the general datatype associated with **datatype**; it will generate a message with the type signature determined by the **datatype** argument. `MPI_RECV(buf, 1, datatype, ...)` will use the receive buffer defined by the base address **buf** and the general datatype associated with **datatype**.

General datatypes can be used in all send and receive operations. We discuss, in Section 5.1.11, the case where the second argument **count** has value > 1 .

The basic datatypes presented in Section 3.2.2 are particular cases of a general datatype, and are predefined. Thus, `MPI_INT` is a predefined handle to a datatype with type map $\{(int, 0)\}$, with one entry of type `int` and displacement zero. The other basic datatypes are similar.

The **extent** of a datatype is defined to be the span from the first byte to the last byte occupied by entries in this datatype, rounded up to satisfy alignment requirements. That is, if

$$Typemap = \{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

then

$$\begin{aligned} lb(Typemap) &= \min_j disp_j, \\ ub(Typemap) &= \max_j (disp_j + \text{sizeof}(type_j)) + \epsilon, \text{ and} \\ extent(Typemap) &= ub(Typemap) - lb(Typemap). \end{aligned} \tag{5.1}$$

If $type_j$ requires alignment to a byte address that is a multiple of k_j , then ϵ is the least nonnegative increment needed to round $extent(Typemap)$ to the next multiple of $\max_j k_j$. In Fortran, it is implementation dependent whether the MPI implementation computes the alignments k_j according to the alignments used by the compiler in common blocks, SEQUENCE derived types, BIND(C) derived types, or derived types that are neither SEQUENCE nor BIND(C). The complete definition of **extent** is given by Equation 5.1 Section 5.1.

Example 5.1. Assume that $Type = \{(double, 0), (char, 8)\}$ (a double at displacement zero, followed by a char at displacement eight). Assume, furthermore, that doubles have to be strictly aligned at addresses that are multiples of eight. Then, the extent of this datatype is 16 (9 rounded to the next multiple of 8). A datatype that consists of a character immediately followed by a double will also have an extent of 16.

Rationale. The definition of extent is motivated by the assumption that the amount of padding added at the end of each structure in an array of structures is the least needed to fulfill alignment constraints. More explicit control of the extent is provided in Section 5.1.6. Such explicit control is needed in cases where the assumption does not hold, for example, where union types are used. In Fortran, structures can be expressed with several language features, e.g., common blocks, SEQUENCE derived types, or BIND(C) derived types. The compiler may use different alignments, and therefore, it is recommended to use MPI_TYPE_CREATE_RESIZED for arrays of structures if an alignment may cause an alignment-gap at the end of a structure as described in Section 5.1.6 and in Section 19.1.15. (*End of rationale.*)

5.1.1 Type Constructors with Explicit Addresses

In Fortran, the procedures MPI_TYPE_CREATE_HVECTOR, MPI_TYPE_CREATE_HINDEXED, MPI_TYPE_CREATE_HINDEXED_BLOCK, MPI_TYPE_CREATE_STRUCT, and MPI_GET_ADDRESS accept arguments of type INTEGER(KIND=MPI_ADDRESS_KIND), wherever arguments of type MPI_Aint are used in C. For Fortran compilers that do not support the Fortran 90 KIND notation, and where addresses are 64 bits whereas default INTEGERS are 32 bits, these arguments will be of type INTEGER*8 (assuming the Fortran compiler accepts the common extension of INTEGER*8 for eight-byte integers).

For the large count versions of three datatype constructors with explicit addresses, MPI_TYPE_CREATE_HINDEXED, MPI_TYPE_CREATE_HINDEXED_BLOCK, and MPI_TYPE_CREATE_STRUCT, absolute addresses shall not be used to specify byte displacements since the parameter is of type MPI_COUNT instead of type MPI_AINT.

5.1.2 Datatype Constructors

Contiguous. The simplest datatype constructor is MPI_TYPE_CONTIGUOUS, which allows replication of a datatype into contiguous locations.

MPI_TYPE_CONTIGUOUS(count, oldtype, newtype)

IN	count	replication count (non-negative integer)
IN	oldtype	old datatype (handle)
OUT	newtype	new datatype (handle)

C binding

```
int MPI_Type_contiguous(int count, MPI_Datatype oldtype, MPI_Datatype *newtype)
```

```

1 int MPI_Type_contiguous_c(MPI_Count count, MPI_Datatype oldtype,
2     MPI_Datatype *newtype)
3

```

Fortran 2008 binding

```

4 MPI_Type_contiguous(count, oldtype, newtype, ierror)
5     INTEGER, INTENT(IN) :: count
6     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
7     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
8     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
9
10 MPI_Type_contiguous(count, oldtype, newtype, ierror) !(_c)
11     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
12     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
13     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
14     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15

```

Fortran binding

```

16 MPI_TYPE_CONTIGUOUS(COUNT, OLDDTYPE, NEWTYPE, IERROR)
17     INTEGER COUNT, OLDDTYPE, NEWTYPE, IERROR
18

```

newtype is the datatype obtained by concatenating count copies of oldtype. Concatenation is defined using *extent* as the size of the concatenated copies.

Example 5.2. Let oldtype have type map $\{(double, 0), (char, 8)\}$, with extent 16, and let $count = 3$. The type map of the datatype returned by newtype is

$$\{(double, 0), (char, 8), (double, 16), (char, 24), (double, 32), (char, 40)\};$$

i.e., alternating double and char elements, with displacements 0, 8, 16, 24, 32, 40.

In general, assume that the type map of oldtype is

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

with extent ex . Then newtype has a type map with $count \cdot n$ entries defined by:

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1}), (type_0, disp_0 + ex), \dots, (type_{n-1}, disp_{n-1} + ex), \\ \dots, (type_0, disp_0 + ex \cdot (count - 1)), \dots, (type_{n-1}, disp_{n-1} + ex \cdot (count - 1))\}.$$

Vector. The procedure MPI_TYPE_VECTOR is a more general constructor that allows replication of a datatype into locations that consist of equally spaced blocks. Each block is obtained by concatenating the same number of copies of the old datatype. The spacing between blocks is a multiple of the extent of the old datatype.

MPI_TYPE_VECTOR(count, blocklength, stride, oldtype, newtype)			1
IN	count	number of blocks (non-negative integer)	2
IN	blocklength	number of elements in each block (non-negative integer)	3
IN	stride	number of elements between start of each block (integer)	4
IN	oldtype	old datatype (handle)	5
OUT	newtype	new datatype (handle)	6

C binding

```
int MPI_Type_vector(int count, int blocklength, int stride,
                   MPI_Datatype oldtype, MPI_Datatype *newtype)
```

```
int MPI_Type_vector_c(MPI_Count count, MPI_Count blocklength, MPI_Count stride,
                     MPI_Datatype oldtype, MPI_Datatype *newtype)
```

Fortran 2008 binding

```
MPI_Type_vector(count, blocklength, stride, oldtype, newtype, ierror)
  INTEGER, INTENT(IN) :: count, blocklength, stride
  TYPE(MPI_Datatype), INTENT(IN) :: oldtype
  TYPE(MPI_Datatype), INTENT(OUT) :: newtype
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Type_vector(count, blocklength, stride, oldtype, newtype, ierror) !(_c)
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength, stride
  TYPE(MPI_Datatype), INTENT(IN) :: oldtype
  TYPE(MPI_Datatype), INTENT(OUT) :: newtype
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_VECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR)
  INTEGER COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR
```

Example 5.3. Assume, again, that oldtype has type map $\{(double, 0), (char, 8)\}$, with extent 16. A call to MPI_TYPE_VECTOR(2, 3, 4, oldtype, newtype) will create the datatype with type map,

$$\{(double, 0), (char, 8), (double, 16), (char, 24), (double, 32), (char, 40), \\ (double, 64), (char, 72), (double, 80), (char, 88), (double, 96), (char, 104)\}.$$

That is, two blocks with three copies each of the old type, with a stride of 4 elements ($4 \cdot 16$ bytes) between the the start of each block.

Example 5.4. A call to MPI_TYPE_VECTOR(3, 1, -2, oldtype, newtype) will create the datatype,

$$\{(double, 0), (char, 8), (double, -32), (char, -24), (double, -64), (char, -56)\}.$$

In general, assume that `oldtype` has type map,

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

with extent ex . Let `bl` be the `blocklength`. The newly created datatype has a type map with `count · bl · n` entries:

$$\begin{aligned} &\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1}), \\ &(type_0, disp_0 + ex), \dots, (type_{n-1}, disp_{n-1} + ex), \dots, \\ &(type_0, disp_0 + (bl - 1) \cdot ex), \dots, (type_{n-1}, disp_{n-1} + (bl - 1) \cdot ex), \\ &(type_0, disp_0 + stride \cdot ex), \dots, (type_{n-1}, disp_{n-1} + stride \cdot ex), \dots, \\ &(type_0, disp_0 + (stride + bl - 1) \cdot ex), \dots, (type_{n-1}, disp_{n-1} + (stride + bl - 1) \cdot ex), \dots, \\ &(type_0, disp_0 + stride \cdot (count - 1) \cdot ex), \dots, \\ &(type_{n-1}, disp_{n-1} + stride \cdot (count - 1) \cdot ex), \dots, \\ &(type_0, disp_0 + (stride \cdot (count - 1) + bl - 1) \cdot ex), \dots, \\ &(type_{n-1}, disp_{n-1} + (stride \cdot (count - 1) + bl - 1) \cdot ex)\}. \end{aligned}$$

A call to `MPI_TYPE_CONTIGUOUS(count, oldtype, newtype)` is equivalent to a call to `MPI_TYPE_VECTOR(count, 1, 1, oldtype, newtype)`, or to a call to `MPI_TYPE_VECTOR(1, count, n, oldtype, newtype)`, where `n` is an arbitrary integer value.

Hvector. The procedure `MPI_TYPE_CREATE_HVECTOR` is identical to `MPI_TYPE_VECTOR`, except that `stride` is given in bytes, rather than in elements. The use for both types of vector constructors is illustrated in Section 5.1.14. (H stands for “heterogeneous”).

`MPI_TYPE_CREATE_HVECTOR(count, blocklength, stride, oldtype, newtype)`

IN	count	number of blocks (non-negative integer)
IN	blocklength	number of elements in each block (non-negative integer)
IN	stride	number of bytes between start of each block (integer)
IN	oldtype	old datatype (handle)
OUT	newtype	new datatype (handle)

C binding

```
int MPI_Type_create_hvector(int count, int blocklength, MPI_Aint stride,
    MPI_Datatype oldtype, MPI_Datatype *newtype)
int MPI_Type_create_hvector_c(MPI_Count count, MPI_Count blocklength,
    MPI_Count stride, MPI_Datatype oldtype, MPI_Datatype *newtype)
```

Fortran 2008 binding

```

MPI_Type_create_hvector(count, blocklength, stride, oldtype, newtype, ierror)
  INTEGER, INTENT(IN) :: count, blocklength
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: stride
  TYPE(MPI_Datatype), INTENT(IN) :: oldtype
  TYPE(MPI_Datatype), INTENT(OUT) :: newtype
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_create_hvector(count, blocklength, stride, oldtype, newtype, ierror)
  !(_c)
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength, stride
  TYPE(MPI_Datatype), INTENT(IN) :: oldtype
  TYPE(MPI_Datatype), INTENT(OUT) :: newtype
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TYPE_CREATE_HVECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR)
  INTEGER COUNT, BLOCKLENGTH, OLDTYPE, NEWTYPE, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) STRIDE

```

Assume that `oldtype` has type map,

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

with extent ex . Let bl be the blocklength. The newly created datatype has a type map with $count \cdot bl \cdot n$ entries:

$$\begin{aligned}
&\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1}), \\
&(type_0, disp_0 + ex), \dots, (type_{n-1}, disp_{n-1} + ex), \dots, \\
&(type_0, disp_0 + (bl - 1) \cdot ex), \dots, (type_{n-1}, disp_{n-1} + (bl - 1) \cdot ex), \\
&(type_0, disp_0 + stride), \dots, (type_{n-1}, disp_{n-1} + stride), \dots, \\
&(type_0, disp_0 + stride + (bl - 1) \cdot ex), \dots, \\
&(type_{n-1}, disp_{n-1} + stride + (bl - 1) \cdot ex), \dots, \\
&(type_0, disp_0 + stride \cdot (count - 1)), \dots, (type_{n-1}, disp_{n-1} + stride \cdot (count - 1)), \dots, \\
&(type_0, disp_0 + stride \cdot (count - 1) + (bl - 1) \cdot ex), \dots, \\
&(type_{n-1}, disp_{n-1} + stride \cdot (count - 1) + (bl - 1) \cdot ex)\}.
\end{aligned}$$

Indexed. The procedure `MPI_TYPE_INDEXED` allows replication of an old datatype into a sequence of blocks (each block is a concatenation of the old datatype), where each block can contain a different number of copies and have a different displacement. All block displacements are multiples of the old type extent.

```

1 MPI_TYPE_INDEXED(count, array_of_blocklengths, array_of_displacements, oldtype,
2     newtype)
3
4 IN     count           number of blocks—also number of entries in
5     array_of_displacements and array_of_blocklengths
6     (non-negative integer)
7
8 IN     array_of_blocklengths  number of elements per block (array of non-negative
9     integers)
10
11 IN     array_of_displacements displacement for each block, in multiples of oldtype
12     (array of integers)
13
14 IN     oldtype         old datatype (handle)
15
16 OUT    newtype        new datatype (handle)

```

C binding

```

17 int MPI_Type_indexed(int count, const int array_of_blocklengths[],
18     const int array_of_displacements[], MPI_Datatype oldtype,
19     MPI_Datatype *newtype)
20
21 int MPI_Type_indexed_c(MPI_Count count,
22     const MPI_Count array_of_blocklengths[],
23     const MPI_Count array_of_displacements[], MPI_Datatype oldtype,
24     MPI_Datatype *newtype)

```

Fortran 2008 binding

```

25 MPI_Type_indexed(count, array_of_blocklengths, array_of_displacements, oldtype,
26     newtype, ierror)
27     INTEGER, INTENT(IN) :: count, array_of_blocklengths(count),
28     array_of_displacements(count)
29     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
30     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Type_indexed(count, array_of_blocklengths, array_of_displacements, oldtype,
34     newtype, ierror) !(_c)
35     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count,
36     array_of_blocklengths(count), array_of_displacements(count)
37     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
38     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
39     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

40 MPI_TYPE_INDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS, OLDDTYPE,
41     NEWTYPE, IERROR)
42     INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*),
43     OLDDTYPE, NEWTYPE, IERROR

```

Example 5.5. Let oldtype have type map $\{(double, 0), (char, 8)\}$, with extent 16. Let $B = (3, 1)$ and let $D = (4, 0)$. A call to `MPI_TYPE_INDEXED(2, B, D, oldtype, newtype)` returns

a datatype with type map,

$$\{(\text{double}, 64), (\text{char}, 72), (\text{double}, 80), (\text{char}, 88), (\text{double}, 96), (\text{char}, 104),$$

$$(\text{double}, 0), (\text{char}, 8)\}.$$

That is, three copies of the old type starting at displacement 64, and one copy starting at displacement 0.

In general, assume that `oldtype` has type map,

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

with extent ex . Let `B` be the `array_of_blocklengths` argument and `D` be the `array_of_displacements` argument. The newly created datatype has $n \cdot \sum_{i=0}^{\text{count}-1} B[i]$ entries:

$$\{(type_0, disp_0 + D[0] \cdot ex), \dots, (type_{n-1}, disp_{n-1} + D[0] \cdot ex), \dots,$$

$$(type_0, disp_0 + (D[0] + B[0] - 1) \cdot ex), \dots,$$

$$(type_{n-1}, disp_{n-1} + (D[0] + B[0] - 1) \cdot ex), \dots,$$

$$(type_0, disp_0 + D[\text{count}-1] \cdot ex), \dots, (type_{n-1}, disp_{n-1} + D[\text{count}-1] \cdot ex), \dots,$$

$$(type_0, disp_0 + (D[\text{count}-1] + B[\text{count}-1] - 1) \cdot ex), \dots,$$

$$(type_{n-1}, disp_{n-1} + (D[\text{count}-1] + B[\text{count}-1] - 1) \cdot ex)\}.$$

A call to `MPI_TYPE_VECTOR(count, blocklength, stride, oldtype, newtype)` is equivalent to a call to `MPI_TYPE_INDEXED(count, B, D, oldtype, newtype)` where

$$D[j] = j \cdot \text{stride}, \quad j = 0, \dots, \text{count} - 1,$$

and

$$B[j] = \text{blocklength}, \quad j = 0, \dots, \text{count} - 1.$$

Hindexed. The procedure `MPI_TYPE_CREATE_HINDEXED` is identical to `MPI_TYPE_INDEXED`, except that block displacements in `array_of_displacements` are specified in bytes, rather than in multiples of the `oldtype` extent.

```

1 MPI_TYPE_CREATE_HINDEXED(count, array_of_blocklengths, array_of_displacements,
2     oldtype, newtype)
3
4     IN      count      number of blocks—also number of entries in
5     array_of_displacements and array_of_blocklengths
6     (non-negative integer)
7
8     IN      array_of_blocklengths  number of elements in each block (array of
9     non-negative integers)
10
11    IN      array_of_displacements  byte displacement of each block (array of integers)
12
13    IN      oldtype      old datatype (handle)
14
15    OUT     newtype      new datatype (handle)

```

C binding

```

16 int MPI_Type_create_hindexed(int count, const int array_of_blocklengths[],
17     const MPI_Aint array_of_displacements[], MPI_Datatype oldtype,
18     MPI_Datatype *newtype)
19
20 int MPI_Type_create_hindexed_c(MPI_Count count,
21     const MPI_Count array_of_blocklengths[],
22     const MPI_Count array_of_displacements[], MPI_Datatype oldtype,
23     MPI_Datatype *newtype)

```

Fortran 2008 binding

```

24 MPI_Type_create_hindexed(count, array_of_blocklengths, array_of_displacements,
25     oldtype, newtype, ierror)
26     INTEGER, INTENT(IN) :: count, array_of_blocklengths(count)
27     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: array_of_displacements(count)
28     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
29     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
30     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Type_create_hindexed(count, array_of_blocklengths, array_of_displacements,
33     oldtype, newtype, ierror) !(_c)
34     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count,
35     array_of_blocklengths(count), array_of_displacements(count)
36     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
37     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
38     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

39 MPI_TYPE_CREATE_HINDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
40     OLDTYPE, NEWTYPE, IERROR)
41     INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), OLDTYPE, NEWTYPE, IERROR
42     INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_DISPLACEMENTS(*)
43
44     Assume that oldtype has type map,
45
46      $\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\}$ ,
47
48

```

with extent ex . Let B be the `array_of_blocklengths` argument and D be the `array_of_displacements` argument. The newly created datatype has a type map with $n \cdot \sum_{i=0}^{\text{count}-1} B[i]$ entries:

$$\begin{aligned} & \{(type_0, disp_0 + D[0]), \dots, (type_{n-1}, disp_{n-1} + D[0]), \dots, \\ & (type_0, disp_0 + D[0] + (B[0] - 1) \cdot ex), \dots, \\ & (type_{n-1}, disp_{n-1} + D[0] + (B[0] - 1) \cdot ex), \dots, \\ & (type_0, disp_0 + D[\text{count}-1]), \dots, (type_{n-1}, disp_{n-1} + D[\text{count}-1]), \dots, \\ & (type_0, disp_0 + D[\text{count}-1] + (B[\text{count}-1] - 1) \cdot ex), \dots, \\ & (type_{n-1}, disp_{n-1} + D[\text{count}-1] + (B[\text{count}-1] - 1) \cdot ex)\}. \end{aligned}$$

Indexed_block. This procedure is the same as `MPI_TYPE_INDEXED` except that the blocklength is the same for all blocks. There are many codes using indirect addressing arising from unstructured grids where the blocksize is always 1 (gather/scatter). The following convenience procedure allows for constant blocksize and arbitrary displacements.

`MPI_TYPE_CREATE_INDEXED_BLOCK(count, blocklength, array_of_displacements, oldtype, newtype)`

IN	count	number of blocks—also number of entries in <code>array_of_displacements</code> (non-negative integer)
IN	blocklength	number of elements in each block (non-negative integer)
IN	array_of_displacements	array of displacements, in multiples of oldtype (array of integers)
IN	oldtype	old datatype (handle)
OUT	newtype	new datatype (handle)

C binding

```
int MPI_Type_create_indexed_block(int count, int blocklength,
    const int array_of_displacements[], MPI_Datatype oldtype,
    MPI_Datatype *newtype)
```

```
int MPI_Type_create_indexed_block_c(MPI_Count count, MPI_Count blocklength,
    const MPI_Count array_of_displacements[], MPI_Datatype oldtype,
    MPI_Datatype *newtype)
```

Fortran 2008 binding

```
MPI_Type_create_indexed_block(count, blocklength, array_of_displacements,
    oldtype, newtype, ierror)
    INTEGER, INTENT(IN) :: count, blocklength, array_of_displacements(count)
    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
```

```

1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2
3 MPI_Type_create_indexed_block(count, blocklength, array_of_displacements,
4     oldtype, newtype, ierror) !(_c)
5     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength,
6     array_of_displacements(count)
7     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
8     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10
11 Fortran binding
12 MPI_TYPE_CREATE_INDEXED_BLOCK(COUNT, BLOCKLENGTH, ARRAY_OF_DISPLACEMENTS,
13     OLDTYPE, NEWTYPE, IERROR)
14     INTEGER COUNT, BLOCKLENGTH, ARRAY_OF_DISPLACEMENTS(*), OLDTYPE, NEWTYPE,
15     IERROR
16
17 Hindexed_block. The procedure MPI_TYPE_CREATE_HINDEXED_BLOCK is identical
18 to MPI_TYPE_CREATE_INDEXED_BLOCK, except that block displacements in
19 array_of_displacements are specified in bytes, rather than in multiples of the oldtype extent.
20
21 MPI_TYPE_CREATE_HINDEXED_BLOCK(count, blocklength, array_of_displacements,
22     oldtype, newtype)
23
24 IN     count                number of blocks—also number of entries in
25                                array_of_displacements (non-negative integer)
26
27 IN     blocklength          number of elements in each block (non-negative
28                                integer)
29
30 IN     array_of_displacements byte displacement of each block (array of integers)
31
32 IN     oldtype              old datatype (handle)
33
34 OUT    newtype             new datatype (handle)
35
36
37 C binding
38 int MPI_Type_create_hindexed_block(int count, int blocklength,
39     const MPI_Aint array_of_displacements[], MPI_Datatype oldtype,
40     MPI_Datatype *newtype)
41
42 int MPI_Type_create_hindexed_block_c(MPI_Count count, MPI_Count blocklength,
43     const MPI_Count array_of_displacements[], MPI_Datatype oldtype,
44     MPI_Datatype *newtype)
45
46 Fortran 2008 binding
47 MPI_Type_create_hindexed_block(count, blocklength, array_of_displacements,
48     oldtype, newtype, ierror)
49     INTEGER, INTENT(IN) :: count, blocklength
50     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: array_of_displacements(count)
51     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
52     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
53     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Type_create_hindexed_block(count, blocklength, array_of_displacements,
                               oldtype, newtype, ierror) !(_c)
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength,
        array_of_displacements(count)
    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TYPE_CREATE_HINDEXED_BLOCK(COUNT, BLOCKLENGTH, ARRAY_OF_DISPLACEMENTS,
                                OLDDTYPE, NEWTYPE, IERROR)
    INTEGER COUNT, BLOCKLENGTH, OLDDTYPE, NEWTYPE, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_DISPLACEMENTS(*)

```

Struct. MPI_TYPE_CREATE_STRUCT is the most general type constructor. It further generalizes MPI_TYPE_CREATE_HINDEXED in that it allows each block to consist of replications of different datatypes.

```

MPI_TYPE_CREATE_STRUCT(count, array_of_blocklengths, array_of_displacements,
                       array_of_types, newtype)

IN      count          number of blocks—also number of entries in arrays
                       array_of_types, array_of_displacements, and
                       array_of_blocklengths (non-negative integer)

IN      array_of_blocklengths  number of elements in each block (array of
                               non-negative integers)

IN      array_of_displacements  byte displacement of each block (array of integers)

IN      array_of_types         type of elements in each block (array of handles)

OUT     newtype              new datatype (handle)

```

C binding

```

int MPI_Type_create_struct(int count, const int array_of_blocklengths[],
                           const MPI_Aint array_of_displacements[],
                           const MPI_Datatype array_of_types[], MPI_Datatype *newtype)

int MPI_Type_create_struct_c(MPI_Count count,
                              const MPI_Count array_of_blocklengths[],
                              const MPI_Count array_of_displacements[],
                              const MPI_Datatype array_of_types[], MPI_Datatype *newtype)

```

Fortran 2008 binding

```

MPI_Type_create_struct(count, array_of_blocklengths, array_of_displacements,
                       array_of_types, newtype, ierror)
    INTEGER, INTENT(IN) :: count, array_of_blocklengths(count)
    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: array_of_displacements(count)
    TYPE(MPI_Datatype), INTENT(IN) :: array_of_types(count)
    TYPE(MPI_Datatype), INTENT(OUT) :: newtype

```

```

1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2
3     MPI_Type_create_struct(count, array_of_blocklengths, array_of_displacements,
4         array_of_types, newtype, ierror) !(_c)
5     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count,
6         array_of_blocklengths(count), array_of_displacements(count)
7     TYPE(MPI_Datatype), INTENT(IN) :: array_of_types(count)
8     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

10    MPI_TYPE_CREATE_STRUCT(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
11        ARRAY_OF_TYPES, NEWTYPE, IERROR)
12
13    INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_TYPES(*), NEWTYPE, IERROR
14    INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_DISPLACEMENTS(*)

```

Example 5.6. Let `type1` have type map,

$$\{(\text{double}, 0), (\text{char}, 8)\},$$

with extent 16. Let $B = (2, 1, 3)$, $D = (0, 16, 26)$, and $T = (\text{MPI_FLOAT}, \text{type1}, \text{MPI_CHAR})$. Then a call to `MPI_TYPE_CREATE_STRUCT(3, B, D, T, newtype)` returns a datatype with type map,

$$\{(\text{float}, 0), (\text{float}, 4), (\text{double}, 16), (\text{char}, 24), (\text{char}, 26), (\text{char}, 27), (\text{char}, 28)\}.$$

That is, two copies of `MPI_FLOAT` starting at 0, followed by one copy of `type1` starting at 16, followed by three copies of `MPI_CHAR`, starting at 26. In this example, we assume that a float occupies four bytes.

In general, let T be the `array_of_types` argument, where $T[i]$ is a handle to,

$$\text{typemap}_i = \{(type_0^i, disp_0^i), \dots, (type_{n_i-1}^i, disp_{n_i-1}^i)\},$$

with extent ex_i . Let B be the `array_of_blocklength` argument and D be the `array_of_displacements` argument. Let c be the count argument. Then the newly created datatype has a type map with $\sum_{i=0}^{c-1} B[i] \cdot n_i$ entries:

$$\begin{aligned} & \{(type_0^0, disp_0^0 + D[0]), \dots, (type_{n_0}^0, disp_{n_0}^0 + D[0]), \dots, \\ & (type_0^0, disp_0^0 + D[0] + (B[0] - 1) \cdot ex_0), \dots, (type_{n_0}^0, disp_{n_0}^0 + D[0] + (B[0]-1) \cdot ex_0), \dots, \\ & (type_0^{c-1}, disp_0^{c-1} + D[c-1]), \dots, (type_{n_{c-1}-1}^{c-1}, disp_{n_{c-1}-1}^{c-1} + D[c-1]), \dots, \\ & (type_0^{c-1}, disp_0^{c-1} + D[c-1] + (B[c-1] - 1) \cdot ex_{c-1}), \dots, \\ & (type_{n_{c-1}-1}^{c-1}, disp_{n_{c-1}-1}^{c-1} + D[c-1] + (B[c-1]-1) \cdot ex_{c-1})\}. \end{aligned}$$

A call to `MPI_TYPE_CREATE_HINDEXED(count, B, D, oldtype, newtype)` is equivalent to a call to `MPI_TYPE_CREATE_STRUCT(count, B, D, T, newtype)`, where each entry of T is equal to `oldtype`.

5.1.3 Subarray Datatype Constructor

MPI_TYPE_CREATE_SUBARRAY(ndims, array_of_sizes, array_of_subsizes, array_of_starts,			1
order, oldtype, newtype)			2
			3
			4
			5
			6
IN	ndims	number of array dimensions (positive integer)	7
IN	array_of_sizes	number of elements of type oldtype in each dimension	8
		of the full array (array of positive integers)	9
IN	array_of_subsizes	number of elements of type oldtype in each dimension	10
		of the subarray (array of positive integers)	11
IN	array_of_starts	starting coordinates of the subarray in each	12
		dimension (array of non-negative integers)	13
			14
IN	order	array storage order flag (state)	15
IN	oldtype	old datatype (handle)	16
OUT	newtype	new datatype (handle)	17
			18

C binding

```

int MPI_Type_create_subarray(int ndims, const int array_of_sizes[],
                             const int array_of_subsizes[], const int array_of_starts[],
                             int order, MPI_Datatype oldtype, MPI_Datatype *newtype)
int MPI_Type_create_subarray_c(int ndims, const MPI_Count array_of_sizes[],
                               const MPI_Count array_of_subsizes[],
                               const MPI_Count array_of_starts[], int order,
                               MPI_Datatype oldtype, MPI_Datatype *newtype)

```

Fortran 2008 binding

```

MPI_Type_create_subarray(ndims, array_of_sizes, array_of_subsizes,
                          array_of_starts, order, oldtype, newtype, ierror)
INTEGER, INTENT(IN) :: ndims, array_of_sizes(ndims),
                          array_of_subsizes(ndims), array_of_starts(ndims), order
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_subarray(ndims, array_of_sizes, array_of_subsizes,
                          array_of_starts, order, oldtype, newtype, ierror) !(_c)
INTEGER, INTENT(IN) :: ndims, order
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: array_of_sizes(ndims),
                          array_of_subsizes(ndims), array_of_starts(ndims)
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TYPE_CREATE_SUBARRAY(NDIMS, ARRAY_OF_SIZES, ARRAY_OF_SUBSIZES,
                          ARRAY_OF_STARTS, ORDER, OLDTYPE, NEWTYPE, IERROR)

```

1 INTEGER NDIMS, ARRAY_OF_SIZES(*), ARRAY_OF_SUBSIZES(*), ARRAY_OF_STARTS(*),
 2 ORDER, OLDTYPE, NEWTYPE, IERROR

3
 4 The subarray type constructor creates an MPI datatype describing an n -dimensional
 5 subarray of an n -dimensional array. The subarray may be situated anywhere within the
 6 full array, and may be of any nonzero size up to the size of the larger array as long as it
 7 is confined within this array. This type constructor facilitates creating filetypes to access
 8 arrays distributed in blocks among processes to a single file that contains the global array,
 9 see MPI I/O, especially Section 14.1.1.

10 This type constructor can handle arrays with an arbitrary number of dimensions and
 11 works for both C and Fortran ordered matrices (i.e., row-major or column-major). Note
 12 that a C program may use Fortran order and a Fortran program may use C order.

13 The `ndims` parameter specifies the number of dimensions in the full data array and
 14 gives the number of elements in `array_of_sizes`, `array_of_subsizes`, and `array_of_starts`.

15 The number of elements of type `oldtype` in each dimension of the n -dimensional array
 16 and the requested subarray are specified by `array_of_sizes` and `array_of_subsizes`, respectively.
 17 For any dimension i , it is erroneous to specify `array_of_subsizes[i] < 1` or `array_of_subsizes[i]`
 18 `> array_of_sizes[i]`.

19 The `array_of_starts` contains the starting coordinates of each dimension of the subarray.
 20 Arrays are assumed to be indexed starting from zero. For any dimension i , it is erroneous to
 21 specify `array_of_starts[i] < 0` or `array_of_starts[i] > (array_of_sizes[i] - array_of_subsizes[i])`.

22 *Advice to users.* In a Fortran program with arrays indexed starting from 1, if the
 23 starting coordinate of a particular dimension of the subarray is n , then the entry in
 24 `array_of_starts` for that dimension is $n-1$. (*End of advice to users.*)

26 The `order` argument specifies the storage order for the subarray as well as the full array.
 27 It must be set to one of the following:

28 MPI_ORDER_C The ordering used by C arrays, (i.e., row-major
 29 order).
 30 MPI_ORDER_FORTRAN The ordering used by Fortran arrays, (i.e., column-
 31 major order).

32
 33 A `ndims`-dimensional subarray (`newtype`) with no extra padding can be defined by the
 34 function `Subarray()` as follows:

35 newtype = Subarray(`ndims`, {`size`₀, `size`₁, ..., `size`_{`ndims`-1}},
 36 {`subsize`₀, `subsize`₁, ..., `subsize`_{`ndims`-1}},
 37 {`start`₀, `start`₁, ..., `start`_{`ndims`-1}}, `oldtype`)

39 Let the typemap of `oldtype` have the form:

40
 41 {(`type`₀, `disp`₀), (`type`₁, `disp`₁), ..., (`type`_{`n`-1}, `disp`_{`n`-1})}

42
 43 where `type` _{i} is a predefined MPI datatype, and let `ex` be the extent of `oldtype`. Then we define
 44 the `Subarray()` function recursively using the following three equations. Equation 5.2 defines
 45 the base step. Equation 5.3 defines the recursion step when `order = MPI_ORDER_FORTRAN`,
 46 and Equation 5.4 defines the recursion step when `order = MPI_ORDER_C`. These equations
 47 use the conceptual datatypes `lb_marker` and `ub_marker`; see Section 5.1.6 for details.

$$\begin{aligned}
& \text{Subarray}(1, \{size_0\}, \{subsize_0\}, \{start_0\}, \\
& \quad \{(type_0, disp_0), (type_1, disp_1), \dots, (type_{n-1}, disp_{n-1})\}) \\
& = \{(\text{lb_marker}, 0), \\
& \quad (type_0, disp_0 + start_0 \times ex), \dots, (type_{n-1}, disp_{n-1} + start_0 \times ex), \\
& \quad (type_0, disp_0 + (start_0 + 1) \times ex), \dots, (type_{n-1}, \\
& \quad \quad disp_{n-1} + (start_0 + 1) \times ex), \dots \\
& \quad (type_0, disp_0 + (start_0 + subsize_0 - 1) \times ex), \dots, \\
& \quad \quad (type_{n-1}, disp_{n-1} + (start_0 + subsize_0 - 1) \times ex), \\
& \quad (\text{ub_marker}, size_0 \times ex)\}
\end{aligned} \tag{5.2}$$

$$\begin{aligned}
& \text{Subarray}(ndims, \{size_0, size_1, \dots, size_{ndims-1}\}, \\
& \quad \{subsize_0, subsize_1, \dots, subsize_{ndims-1}\}, \\
& \quad \{start_0, start_1, \dots, start_{ndims-1}\}, \text{oldtype}) \\
& = \text{Subarray}(ndims - 1, \{size_1, size_2, \dots, size_{ndims-1}\}, \\
& \quad \{subsize_1, subsize_2, \dots, subsize_{ndims-1}\}, \\
& \quad \{start_1, start_2, \dots, start_{ndims-1}\}, \\
& \quad \text{Subarray}(1, \{size_0\}, \{subsize_0\}, \{start_0\}, \text{oldtype}))
\end{aligned} \tag{5.3}$$

$$\begin{aligned}
& \text{Subarray}(ndims, \{size_0, size_1, \dots, size_{ndims-1}\}, \\
& \quad \{subsize_0, subsize_1, \dots, subsize_{ndims-1}\}, \\
& \quad \{start_0, start_1, \dots, start_{ndims-1}\}, \text{oldtype}) \\
& = \text{Subarray}(ndims - 1, \{size_0, size_1, \dots, size_{ndims-2}\}, \\
& \quad \{subsize_0, subsize_1, \dots, subsize_{ndims-2}\}, \\
& \quad \{start_0, start_1, \dots, start_{ndims-2}\}, \\
& \quad \text{Subarray}(1, \{size_{ndims-1}\}, \{subsize_{ndims-1}\}, \{start_{ndims-1}\}, \text{oldtype}))
\end{aligned} \tag{5.4}$$

For an example use of `MPI_TYPE_CREATE_SUBARRAY` in the context of I/O see Section 14.9.2.

5.1.4 Distributed Array Datatype Constructor

The distributed array type constructor supports HPF-like [49] data distributions. However, unlike in HPF, the storage order may be specified for C arrays as well as for Fortran arrays.

Advice to users. One can create an HPF-like file view using this type constructor as follows. Complementary filetypes are created by having every process of a group call this constructor with identical arguments (with the exception of `rank` which should be set appropriately). These filetypes (along with identical `disp` and `etype`) are then used to define the view (via `MPI_FILE_SET_VIEW`), see MPI I/O, especially Section 14.1.1 and Section 14.3. Using this view, a collective data access operation (with identical offsets) will yield an HPF-like distribution pattern. (*End of advice to users.*)

```

1 MPI_TYPE_CREATE_DARRAY(size, rank, ndims, array_of_gsizes, array_of_distribs,
2     array_of_dargs, array_of_psizes, order, oldtype, newtype)
3
4     IN     size                size of process group (positive integer)
5     IN     rank                rank in process group (non-negative integer)
6     IN     ndims              number of array dimensions as well as process grid
7     dimensions (positive integer)
8     IN     array_of_gsizes     number of elements of type oldtype in each dimension
9     of global array (array of positive integers)
10
11    IN     array_of_distribs    distribution of array in each dimension (array of
12    states)
13
14    IN     array_of_dargs       distribution argument in each dimension (array of
15    positive integers)
16
17    IN     array_of_psizes      size of process grid in each dimension (array of
18    positive integers)
19
20    IN     order                array storage order flag (state)
21
22    IN     oldtype              old datatype (handle)
23
24    OUT    newtype              new datatype (handle)

```

C binding

```

25 int MPI_Type_create_darray(int size, int rank, int ndims,
26     const int array_of_gsizes[], const int array_of_distribs[],
27     const int array_of_dargs[], const int array_of_psizes[],
28     int order, MPI_Datatype oldtype, MPI_Datatype *newtype)
29
30 int MPI_Type_create_darray_c(int size, int rank, int ndims,
31     const MPI_Count array_of_gsizes[], const int array_of_distribs[],
32     const int array_of_dargs[], const int array_of_psizes[],
33     int order, MPI_Datatype oldtype, MPI_Datatype *newtype)

```

Fortran 2008 binding

```

34 MPI_Type_create_darray(size, rank, ndims, array_of_gsizes, array_of_distribs,
35     array_of_dargs, array_of_psizes, order, oldtype, newtype, ierror)
36     INTEGER, INTENT(IN) :: size, rank, ndims, array_of_gsizes(ndims),
37     array_of_distribs(ndims), array_of_dargs(ndims),
38     array_of_psizes(ndims), order
39     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
40     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
41     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Type_create_darray(size, rank, ndims, array_of_gsizes, array_of_distribs,
44     array_of_dargs, array_of_psizes, order, oldtype, newtype, ierror)
45     !(_c)
46     INTEGER, INTENT(IN) :: size, rank, ndims, array_of_distribs(ndims),
47     array_of_dargs(ndims), array_of_psizes(ndims), order
48     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: array_of_gsizes(ndims)
49     TYPE(MPI_Datatype), INTENT(IN) :: oldtype

```

```

TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TYPE_CREATE_DARRAY(SIZE, RANK, NDIMS, ARRAY_OF_GSIZES, ARRAY_OF_DISTRIBS,
    ARRAY_OF_DARGS, ARRAY_OF_PSIZEs, ORDER, OLDTYPE, NEWTYPE, IERROR)
INTEGER SIZE, RANK, NDIMS, ARRAY_OF_GSIZES(*), ARRAY_OF_DISTRIBS(*),
    ARRAY_OF_DARGS(*), ARRAY_OF_PSIZEs(*), ORDER, OLDTYPE, NEWTYPE,
    IERROR

```

MPI_TYPE_CREATE_DARRAY can be used to generate the datatypes corresponding to the distribution of an `ndims`-dimensional array of `oldtype` elements onto an `ndims`-dimensional grid of logical processes. Unused dimensions of `array_of_psize`s should be set to 1 (see Example 5.7). For a call to MPI_TYPE_CREATE_DARRAY to be correct, the equation $\prod_{i=0}^{ndims-1} array_of_psizes[i] = size$ must be satisfied. The ordering of processes in the process grid is assumed to be row-major, as in the case of virtual Cartesian process topologies.

Advice to users. For both Fortran and C arrays, the ordering of processes in the process grid is assumed to be row-major. This is consistent with the ordering used in virtual Cartesian process topologies in MPI. To create such virtual process topologies, or to find the coordinates of a process in the process grid, etc., users may use the corresponding process topology procedures, see Chapter 8. (*End of advice to users.*)

Each dimension of the array can be distributed in one of three ways:

MPI_DISTRIBUTE_BLOCK	Block distribution.
MPI_DISTRIBUTE_CYCLIC	Cyclic distribution.
MPI_DISTRIBUTE_NONE	Dimension not distributed.

The constant MPI_DISTRIBUTE_DFLT_DARG specifies a default distribution argument. The distribution argument for a dimension that is not distributed is ignored. For any dimension `i` in which the distribution is MPI_DISTRIBUTE_BLOCK, it is erroneous to specify `array_of_dargs[i] * array_of_psize`s[i] < `array_of_gsize`s[i].

For example, the HPF layout ARRAY(CYCLIC(15)) corresponds to MPI_DISTRIBUTE_CYCLIC with a distribution argument of 15, and the HPF layout ARRAY(BLOCK) corresponds to MPI_DISTRIBUTE_BLOCK with a distribution argument of MPI_DISTRIBUTE_DFLT_DARG.

The `order` argument is used as in MPI_TYPE_CREATE_SUBARRAY to specify the storage order. Therefore, arrays described by this type constructor may be stored in Fortran (column-major) or C (row-major) order. Valid values for `order` are MPI_ORDER_FORTRAN and MPI_ORDER_C.

This routine creates a new MPI datatype with a typemap defined in terms of a function called “cyclic()” (see below).

Without loss of generality, it suffices to define the typemap for the MPI_DISTRIBUTE_CYCLIC case where MPI_DISTRIBUTE_DFLT_DARG is not used.

MPI_DISTRIBUTE_BLOCK and MPI_DISTRIBUTE_NONE can be reduced to the MPI_DISTRIBUTE_CYCLIC case for dimension `i` as follows.

MPI_DISTRIBUTE_BLOCK with `array_of_dargs[i]` equal to MPI_DISTRIBUTE_DFLT_DARG is equivalent to MPI_DISTRIBUTE_CYCLIC with `array_of_dargs[i]` set to

$$(\text{array_of_gsizes}[i] + \text{array_of_psizes}[i] - 1) / \text{array_of_psizes}[i].$$

1 If `array_of_dargs[i]` is not `MPI_DISTRIBUTE_DFLT_DARG`, then `MPI_DISTRIBUTE_BLOCK` and
 2 `MPI_DISTRIBUTE_CYCLIC` are equivalent.

3 `MPI_DISTRIBUTE_NONE` is equivalent to `MPI_DISTRIBUTE_CYCLIC` with `array_of_dargs[i]`
 4 set to `array_of_gsizes[i]`.

5 Finally, `MPI_DISTRIBUTE_CYCLIC` with `array_of_dargs[i]` equal to
 6 `MPI_DISTRIBUTE_DFLT_DARG` is equivalent to `MPI_DISTRIBUTE_CYCLIC` with
 7 `array_of_dargs[i]` set to 1.

8 For `MPI_ORDER_FORTRAN`, an `ndims`-dimensional distributed array (`newtype`) is defined
 9 by the following code fragment:

```
10 oldtypes[0] = oldtype;
11 for (i = 0; i < ndims; i++) {
12     oldtypes[i+1] = cyclic(array_of_dargs[i],
13                          array_of_gsizes[i],
14                          r[i],
15                          array_of_psizes[i],
16                          oldtypes[i]);
17 }
18 newtype = oldtypes[ndims];
```

19 For `MPI_ORDER_C`, the code is:

```
20 oldtypes[0] = oldtype;
21 for (i = 0; i < ndims; i++) {
22     oldtypes[i+1] = cyclic(array_of_dargs[ndims - i - 1],
23                          array_of_gsizes[ndims - i - 1],
24                          r[ndims - i - 1],
25                          array_of_psizes[ndims - i - 1],
26                          oldtypes[i]);
27 }
28 newtype = oldtypes[ndims];
```

29 where `r[i]` is the position of the process (with rank `rank`) in the process grid at dimension
 30 `i`. The values of `r[i]` are given by the following code fragment:

```
31 t_rank = rank;
32 t_size = 1;
33 for (i = 0; i < ndims; i++)
34     t_size *= array_of_psizes[i];
35 for (i = 0; i < ndims; i++) {
36     t_size = t_size / array_of_psizes[i];
37     r[i] = t_rank / t_size;
38     t_rank = t_rank % t_size;
39 }
40
```

Let the typemap of `oldtype` have the form:

$$\{(type_0, disp_0), (type_1, disp_1), \dots, (type_{n-1}, disp_{n-1})\}$$

43 where `typei` is a predefined MPI datatype, and let `ex` be the extent of `oldtype`. The following
 44 function uses the conceptual datatypes `lb_marker` and `ub_marker`, see Section 5.1.6 for details.

45 Given the above, the function `cyclic()` is defined as follows:

```
46 cyclic(darg, gsize, r, psize, oldtype)
47     = {(lb_marker, 0),
48
```

```

1  (type0, disp0 + r × darg × ex), ...,
2      (typen-1, dispn-1 + r × darg × ex),
3
4  (type0, disp0 + (r × darg + 1) × ex), ...,
5      (typen-1, dispn-1 + (r × darg + 1) × ex),
6
7  ...
8  (type0, disp0 + ((r + 1) × darg - 1) × ex), ...,
9      (typen-1, dispn-1 + ((r + 1) × darg - 1) × ex),
10
11 (type0, disp0 + r × darg × ex + psize × darg × ex), ...,
12     (typen-1, dispn-1 + r × darg × ex + psize × darg × ex),
13
14 (type0, disp0 + (r × darg + 1) × ex + psize × darg × ex), ...,
15     (typen-1, dispn-1 + (r × darg + 1) × ex + psize × darg × ex),
16
17 ...
18 (type0, disp0 + ((r + 1) × darg - 1) × ex + psize × darg × ex), ...,
19     (typen-1, dispn-1 + ((r + 1) × darg - 1) × ex + psize × darg × ex),
20
21 ⋮
22 (type0, disp0 + r × darg × ex + psize × darg × ex × (count - 1)), ...,
23     (typen-1, dispn-1 + r × darg × ex + psize × darg × ex × (count - 1)),
24
25 (type0, disp0 + (r × darg + 1) × ex + psize × darg × ex × (count - 1)), ...,
26     (typen-1, dispn-1 + (r × darg + 1) × ex
27         + psize × darg × ex × (count - 1)),
28
29 ...
30 (type0, disp0 + (r × darg + darglast - 1) × ex
31     + psize × darg × ex × (count - 1)), ...,
32     (typen-1, dispn-1 + (r × darg + darglast - 1) × ex
33     + psize × darg × ex × (count - 1)),
34
35 (ub_marker, gsize * ex)

```

where *count* is defined by this code fragment:

```

36 nblocks = (gsize + (darg - 1)) / darg;
37 count = nblocks / psize;
38 left_over = nblocks - count * psize;
39 if (r < left_over)
40     count = count + 1;

```

Here, *nblocks* is the number of blocks that must be distributed among the processors. Finally, *darg_{last}* is defined by this code fragment:

```

43 if ((num_in_last_cyclic = gsize % (psize * darg)) == 0)
44     darg_last = darg;
45 else {
46     darg_last = num_in_last_cyclic - darg * r;
47     if (darg_last > darg)
48         darg_last = darg;

```

```

1   if (darg_last <= 0)
2       darg_last = darg;
3   }

```

Example 5.7. Consider generating the filetypes corresponding to the HPF distribution:

```

6   <oldtype> FILEARRAY(100, 200, 300)
7   !HPF$ PROCESSORS PROCESSES(2, 3)
8   !HPF$ DISTRIBUTE FILEARRAY(CYCLIC(10), *, BLOCK) ONTO PROCESSES

```

This can be achieved by the following Fortran code, assuming there will be six processes attached to the run:

```

12  ndims = 3
13  array_of_gsizes(1) = 100
14  array_of_distribs(1) = MPI_DISTRIBUTE_CYCLIC
15  array_of_dargs(1) = 10
16  array_of_gsizes(2) = 200
17  array_of_distribs(2) = MPI_DISTRIBUTE_NONE
18  array_of_dargs(2) = 0
19  array_of_gsizes(3) = 300
20  array_of_distribs(3) = MPI_DISTRIBUTE_BLOCK
21  array_of_dargs(3) = MPI_DISTRIBUTE_DFLT_DARG
22  array_of_psize(1) = 2
23  array_of_psize(2) = 1
24  array_of_psize(3) = 3
25  call MPI_COMM_SIZE(MPI_COMM_WORLD, size, ierr)
26  call MPI_COMM_RANK(MPI_COMM_WORLD, rank, ierr)
27  call MPI_TYPE_CREATE_DARRAY(size, rank, ndims, array_of_gsizes, &
28  array_of_distribs, array_of_dargs, array_of_psize, &
29  MPI_ORDER_FORTRAN, oldtype, newtype, ierr)

```

5.1.5 Address and Size Procedures

The displacements in a general datatype are relative to some initial buffer address. **Absolute addresses** can be substituted for these displacements: we treat them as displacements relative to “address zero,” the start of the address space. This initial address zero is indicated by the constant `MPI_BOTTOM`. Thus, a datatype can specify the absolute address of the entries in the communication buffer, in which case the `buf` argument is passed the value `MPI_BOTTOM`. Note that in Fortran `MPI_BOTTOM` is not usable for initialization or assignment, see Section 2.5.4.

The address of a location in memory can be found by invoking the procedure `MPI_GET_ADDRESS`. The **relative displacement** between two absolute addresses can be calculated with the procedure `MPI_AINT_DIFF`. A new absolute address as sum of an absolute base address and a relative displacement can be calculated with the procedure `MPI_AINT_ADD`. To ensure portability, arithmetic on absolute addresses should not be performed with the intrinsic operators “-” and “+”. See also Sections 2.5.6 and 5.1.12 on pages 21 and 161.

Rationale. Address sized integer values, i.e., `MPI_Aint` or `INTEGER(KIND=MPI_ADDRESS_KIND)` values, are signed integers, while absolute addresses

are unsigned quantities. Direct arithmetic on addresses stored in address sized signed variables can cause overflows, resulting in undefined behavior. (*End of rationale.*)

MPI_GET_ADDRESS(location, address)

IN	location	location in caller memory (choice)
OUT	address	address of location (integer)

C binding

```
int MPI_Get_address(const void *location, MPI_Aint *address)
```

Fortran 2008 binding

```
MPI_Get_address(location, address, ierror)
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: location
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: address
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GET_ADDRESS(LOCATION, ADDRESS, IERROR)
  <type> LOCATION(*)
  INTEGER(KIND=MPI_ADDRESS_KIND) ADDRESS
  INTEGER IERROR
```

Returns the (byte) address of location.

Rationale. In the `mpi_f08` module, the `location` argument is not defined with `INTENT(IN)` because existing applications may use `MPI_GET_ADDRESS` as a substitute for `MPI_F_SYNC_REG`, which was not defined before MPI-3.0. (*End of rationale.*)

Example 5.8. Using `MPI_GET_ADDRESS` for an array.

```
REAL A(100,100)
INTEGER(KIND=MPI_ADDRESS_KIND) I1, I2, DIFF
CALL MPI_GET_ADDRESS(A(1,1), I1, IERROR)
CALL MPI_GET_ADDRESS(A(10,10), I2, IERROR)
DIFF = MPI_AINT_DIFF(I2, I1)
! The value of DIFF is 909*SIZEOF(REAL); the values of I1 and I2 are
! implementation dependent.
```

Advice to users. C users may be tempted to avoid the usage of `MPI_GET_ADDRESS` and rely on the availability of the address operator `&`. Note, however, that `& cast-expression` is a pointer, not an address. ISO C does not require that the value of a pointer (or the pointer cast to `int`) be the absolute address of the object pointed at—although this is commonly the case. Furthermore, referencing may not have a unique definition on machines with a segmented address space. The use of `MPI_GET_ADDRESS` to “reference” C variables guarantees portability to such machines as well. (*End of advice to users.*)

Advice to users. To prevent problems with the argument copying and register optimization done by Fortran compilers, please note the hints in Sections 19.1.10–19.1.20. (*End of advice to users.*)

To ensure portability, arithmetic on MPI addresses must be performed using the MPI_AINT_ADD and MPI_AINT_DIFF procedures.

MPI_AINT_ADD(base, disp)

IN	base	base address (integer)
IN	disp	displacement (integer)

C binding

MPI_Aint MPI_Aint_add(MPI_Aint base, MPI_Aint disp)

Fortran 2008 binding

INTEGER(KIND=MPI_ADDRESS_KIND) MPI_Aint_add(base, disp)
 INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: base, disp

Fortran binding

INTEGER(KIND=MPI_ADDRESS_KIND) MPI_AINT_ADD(BASE, DISP)
 INTEGER(KIND=MPI_ADDRESS_KIND) BASE, DISP

MPI_AINT_ADD produces a new MPI_Aint value that is equivalent to the sum of the base and disp arguments, where base represents a base address returned by a call to MPI_GET_ADDRESS and disp represents a signed integer displacement. The resulting address is valid only at the process that generated base, and it must correspond to a location in the same object referenced by base, as described in Section 5.1.12. The addition is performed in a manner that results in the correct MPI_Aint representation of the output address, as if the process that originally produced base had called:

```
MPI_Get_address((char *) base + disp, &result);
```

MPI_AINT_DIFF(addr1, addr2)

IN	addr1	minuend address (integer)
IN	addr2	subtrahend address (integer)

C binding

MPI_Aint MPI_Aint_diff(MPI_Aint addr1, MPI_Aint addr2)

Fortran 2008 binding

INTEGER(KIND=MPI_ADDRESS_KIND) MPI_Aint_diff(addr1, addr2)
 INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: addr1, addr2

Fortran binding

INTEGER(KIND=MPI_ADDRESS_KIND) MPI_AINT_DIFF(ADDR1, ADDR2)
 INTEGER(KIND=MPI_ADDRESS_KIND) ADDR1, ADDR2

MPI_AINT_DIFF produces a new MPI_Aint value that is equivalent to the difference between addr1 and addr2 arguments, where addr1 and addr2 represent addresses returned

by calls to `MPI_GET_ADDRESS`. The resulting address is valid only at the process that generated `addr1` and `addr2`, and `addr1` and `addr2` must correspond to locations in the same object in the same process, as described in Section 5.1.12. The difference is calculated in a manner that results in the signed difference from `addr1` to `addr2`, as if the process that originally produced the addresses had called `(char *) addr1 - (char *) addr2` on the addresses initially passed to `MPI_GET_ADDRESS`.

The following auxiliary procedures provide useful information on derived datatypes.

`MPI_TYPE_SIZE(datatype, size)`

IN	<code>datatype</code>	datatype to get information on (handle)
OUT	<code>size</code>	datatype size (integer)

C binding

`int MPI_Type_size(MPI_Datatype datatype, int *size)`

`int MPI_Type_size_c(MPI_Datatype datatype, MPI_Count *size)`

Fortran 2008 binding

`MPI_Type_size(datatype, size, ierror)`
`TYPE(MPI_Datatype), INTENT(IN) :: datatype`
`INTEGER, INTENT(OUT) :: size`
`INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

`MPI_Type_size(datatype, size, ierror) !(_c)`
`TYPE(MPI_Datatype), INTENT(IN) :: datatype`
`INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size`
`INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

Fortran binding

`MPI_TYPE_SIZE(DATATYPE, SIZE, IERROR)`
`INTEGER DATATYPE, SIZE, IERROR`

`MPI_TYPE_SIZE` set the value of `size` to the total size, in bytes, of the entries in the type signature associated with `datatype`; i.e., the total size of the data in a message that would be created with this `datatype`. Entries that occur multiple times in the `datatype` are counted with their multiplicity. For both procedures, if the `OUT` parameter cannot express the value to be returned (e.g., if the parameter is too small to hold the output value), it is set to `MPI_UNDEFINED`.

5.1.6 Lower-Bound and Upper-Bound Markers

It is often convenient to define explicitly the lower bound and upper bound of a type map, and override the definition given on page 152. This allows one to define a `datatype` that has “holes” at its beginning or its end, or a `datatype` with entries that extend above the upper bound or below the lower bound. Examples of such usage are provided in Section 5.1.14. Also, the user may want to override the alignment rules that are used to compute upper bounds and extents. E.g., a C compiler may allow the user to override default alignment rules for some of the structures within a program. The user has to specify explicitly the bounds of the `datatypes` that match these structures.

To achieve this, we add two additional conceptual datatypes, **lb_marker** and **ub_marker**, that represent the lower bound and upper bound of a datatype. These conceptual datatypes occupy no space ($extent(\text{lb_marker}) = extent(\text{ub_marker}) = 0$). They do not affect the size or count of a datatype, and do not affect the content of a message created with this datatype. However, they do affect the definition of the extent of a datatype and, therefore, affect the outcome of a replication of this datatype by a datatype constructor.

Example 5.9. A call to `MPI_TYPE_CREATE_RESIZED(MPI_INT, -3, 9, type1)` creates a new datatype that has an extent of 9 (from -3 to 5, 5 included), and contains an integer at displacement 0. This is the datatype defined by the typemap $\{(\text{lb_marker}, -3), (\text{int}, 0), (\text{ub_marker}, 6)\}$. If this type is replicated twice by a call to `MPI_TYPE_CONTIGUOUS(2, type1, type2)` then the newly created type can be described by the typemap $\{(\text{lb_marker}, -3), (\text{int}, 0), (\text{int}, 9), (\text{ub_marker}, 15)\}$. (An entry of type `ub_marker` can be deleted if there is another entry of type `ub_marker` with a higher displacement; an entry of type `lb_marker` can be deleted if there is another entry of type `lb_marker` with a lower displacement.)

In general, if

$$\text{Typemap} = \{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

then the **lower bound** of *Typemap* is defined to be

$$lb(\text{Typemap}) = \begin{cases} \min_j disp_j & \text{if no entry has type } \text{lb_marker} \\ \min_j \{disp_j \text{ such that } type_j = \text{lb_marker}\} & \text{otherwise} \end{cases}$$

Similarly, the **upper bound** of *Typemap* is defined to be

$$ub(\text{Typemap}) = \begin{cases} \max_j (disp_j + \text{sizeof}(type_j)) + \epsilon & \text{if no entry has type } \text{ub_marker} \\ \max_j \{disp_j \text{ such that } type_j = \text{ub_marker}\} & \text{otherwise} \end{cases}$$

Then

$$extent(\text{Typemap}) = ub(\text{Typemap}) - lb(\text{Typemap})$$

If $type_i$ requires alignment to a byte address that is a multiple of k_i , then ϵ is the least nonnegative increment needed to round $extent(\text{Typemap})$ to the next multiple of $\max_i k_i$. In Fortran, it is implementation dependent whether the MPI implementation computes the alignments k_i according to the alignments used by the compiler in common blocks, SEQUENCE derived types, BIND(C) derived types, or derived types that are neither SEQUENCE nor BIND(C).

The formal definitions given for the various datatype constructors apply now, with the amended definition of **extent**.

Rationale. Before Fortran 2003, `MPI_TYPE_CREATE_STRUCT` could be applied to Fortran common blocks and SEQUENCE derived types. With Fortran 2003, this list was extended by BIND(C) derived types and MPI implementors have implemented the alignments k_i differently, i.e., some based on the alignments used in SEQUENCE derived types, and others according to BIND(C) derived types. (*End of rationale.*)

Advice to implementors. In Fortran, it is generally recommended to use BIND(C) derived types instead of common blocks or SEQUENCE derived types. Therefore it is recommended to calculate the alignments k_i based on BIND(C) derived types. (*End of advice to implementors.*)

Advice to users. Structures combining different basic datatypes should be defined so that there will be no gaps based on alignment rules. If such a datatype is used to create an array of structures, users should also avoid an alignment-gap at the end of the structure. In MPI communication, the content of such gaps would not be communicated into the receiver's buffer. For example, such an alignment-gap may occur between an odd number of floats or REALs before a double or DOUBLE PRECISION data. Such gaps may be added explicitly to both the structure and the MPI derived datatype handle because the communication of a contiguous derived datatype may be significantly faster than the communication of one that is noncontiguous because of such alignment-gaps.

As an example, instead of

```

TYPE, BIND(C) :: my_data
  REAL, DIMENSION(3) :: x
  ! there may be a gap of the size of one REAL
  ! if the alignment of a DOUBLE PRECISION is
  ! two times the size of a REAL
  DOUBLE PRECISION :: p
END TYPE

```

one should define

```

TYPE, BIND(C) :: my_data
  REAL, DIMENSION(3) :: x
  REAL :: gap1
  DOUBLE PRECISION :: p
END TYPE

```

and also include `gap1` in the matching MPI derived datatype. It is required that all processes in a communication add the same gaps, i.e., defined with the same basic datatype. Both the original and the modified structures are portable, but may have different performance implications for the communication and memory accesses during computation on systems with different alignment values.

In principle, a compiler may define an additional alignment rule for structures, e.g., to use at least 4 or 8 byte alignment, although the content may have a $max_i k_i$ alignment less than this structure alignment. To maintain portability, users should always resize structure derived datatype handles if used in an array of structures, see the Example in Section 19.1.15. (*End of advice to users.*)

5.1.7 Extent and Bounds of Datatypes

`MPI_Type_GET_EXTENT(datatype, lb, extent)`

IN	datatype	datatype to get information on (handle)
OUT	lb	lower bound of datatype (integer)
OUT	extent	extent of datatype (integer)

C binding

`int MPI_Type_get_extent(MPI_Datatype datatype, MPI_Aint *lb, MPI_Aint *extent)`

```

1 int MPI_Type_get_extent_c(MPI_Datatype datatype, MPI_Count *lb,
2     MPI_Count *extent)

```

Fortran 2008 binding

```

4 MPI_Type_get_extent(datatype, lb, extent, ierror)
5     TYPE(MPI_Datatype), INTENT(IN) :: datatype
6     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: lb, extent
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8

```

```

9 MPI_Type_get_extent(datatype, lb, extent, ierror) !(_c)
10    TYPE(MPI_Datatype), INTENT(IN) :: datatype
11    INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: lb, extent
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

13 MPI_TYPE_GET_EXTENT(DATATYPE, LB, EXTENT, IERROR)
14    INTEGER DATATYPE, IERROR
15    INTEGER(KIND=MPI_ADDRESS_KIND) LB, EXTENT
16

```

Returns the lower bound and the extent of `datatype` (as defined in Equation 5.1).

If either OUT parameter cannot express the value to be returned (e.g., if the parameter is too small to hold the output value), it is set to `MPI_UNDEFINED`.

MPI allows one to change the extent of a datatype, using lower bound and upper bound markers. This provides control over the stride of successive datatypes that are replicated by datatype constructors, or are replicated by the count argument in a send or receive call.

```

24
25 MPI_TYPE_CREATE_RESIZED(oldtype, lb, extent, newtype)

```

26	IN	oldtype	input datatype (handle)
27	IN	lb	new lower bound of datatype (integer)
28	IN	extent	new extent of datatype (integer)
29	OUT	newtype	output datatype (handle)
30			
31			

C binding

```

32
33 int MPI_Type_create_resized(MPI_Datatype oldtype, MPI_Aint lb, MPI_Aint extent,
34     MPI_Datatype *newtype)

```

```

35
36 int MPI_Type_create_resized_c(MPI_Datatype oldtype, MPI_Count lb,
37     MPI_Count extent, MPI_Datatype *newtype)
38

```

Fortran 2008 binding

```

39 MPI_Type_create_resized(oldtype, lb, extent, newtype, ierror)
40     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
41     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: lb, extent
42     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
43     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44

```

```

45 MPI_Type_create_resized(oldtype, lb, extent, newtype, ierror) !(_c)
46     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
47     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: lb, extent
48     TYPE(MPI_Datatype), INTENT(OUT) :: newtype

```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_CREATE_RESIZED(OLDTYPE, LB, EXTENT, NEWTYPE, IERROR)
  INTEGER OLDTYPE, NEWTYPE, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) LB, EXTENT
```

Returns in `newtype` a handle to a new datatype that is identical to `oldtype`, except that the lower bound of this new datatype is set to be `lb`, and its upper bound is set to be `lb + extent`. Any previous `lb` and `ub` markers are erased, and a new pair of lower bound and upper bound markers are put in the positions indicated by the `lb` and `extent` arguments. This affects the behavior of the datatype when used in communication operations, with `count > 1`, and when used in the construction of new derived datatypes.

5.1.8 True Extent of Datatypes

Suppose we implement gather (see also Section 6.5) as a spanning tree implemented on top of point-to-point routines. Since the receive buffer is only valid on the root process, one will need to allocate some temporary space for receiving data on intermediate nodes. However, the datatype extent cannot be used as an estimate of the amount of space that needs to be allocated, if the user has modified the extent, for example by using `MPI_TYPE_CREATE_RESIZED`. The procedure `MPI_TYPE_GET_TRUE_EXTENT` returns the true extent of the datatype.

```
MPI_TYPE_GET_TRUE_EXTENT(datatype, true_lb, true_extent)
```

IN	datatype	datatype to get information on (handle)
OUT	true_lb	true lower bound of datatype (integer)
OUT	true_extent	true extent of datatype (integer)

C binding

```
int MPI_Type_get_true_extent(MPI_Datatype datatype, MPI_Aint *true_lb,
  MPI_Aint *true_extent)
int MPI_Type_get_true_extent_c(MPI_Datatype datatype, MPI_Count *true_lb,
  MPI_Count *true_extent)
```

Fortran 2008 binding

```
MPI_Type_get_true_extent(datatype, true_lb, true_extent, ierror)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: true_lb, true_extent
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_get_true_extent(datatype, true_lb, true_extent, ierror) !(_c)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: true_lb, true_extent
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_GET_TRUE_EXTENT(DATATYPE, TRUE_LB, TRUE_EXTENT, IERROR)
  INTEGER DATATYPE, IERROR
```

1 INTEGER(KIND=MPI_ADDRESS_KIND) TRUE_LB, TRUE_EXTENT

2 *true_lb* returns the offset of the lowest unit of storage that is addressed by the datatype,
3 i.e., the lower bound of the corresponding typemap, ignoring explicit lower bound mark-
4 ers. *true_extent* returns the true size of the datatype, i.e., the extent of the correspond-
5 ing typemap, ignoring explicit lower bound and upper bound markers, and performing no
6 rounding for alignment. If the typemap associated with datatype is

$$7 \quad \text{Typemap} = \{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\}$$

8 Then

$$9 \quad \text{true_lb}(\text{Typemap}) = \min_j \{disp_j : type_j \neq \text{lb_marker}, \text{ub_marker}\},$$

$$10 \quad \text{true_ub}(\text{Typemap}) = \max_j \{disp_j + \text{sizeof}(type_j) : type_j \neq \text{lb_marker}, \text{ub_marker}\},$$

11 and

$$12 \quad \text{true_extent}(\text{Typemap}) = \text{true_ub}(\text{Typemap}) - \text{true_lb}(\text{Typemap}).$$

13 (Readers should compare this with the definitions in Section 5.1.6 and Section 5.1.7, which
14 describe the procedure `MPI_TYPE_GET_EXTENT`.)

15 The *true_extent* is the minimum number of bytes of memory necessary to hold a
16 datatype, uncompressed.

17 If either OUT parameter cannot express the value to be returned (e.g., if the parameter
18 is too small to hold the output value), it is set to `MPI_UNDEFINED`.

19 5.1.9 Commit and Free

20 A datatype object has to be **committed** before it can be used in a communication. As
21 an argument in datatype constructors, uncommitted and also committed datatypes can be
22 used. There is no need to commit basic datatypes. They are “pre-committed.”

23 `MPI_TYPE_COMMIT(datatype)`

24 INOUT datatype datatype that is committed (handle)

25 C binding

26 int MPI_Type_commit(MPI_Datatype *datatype)

27 Fortran 2008 binding

28 MPI_Type_commit(datatype, ierror)

29 TYPE(MPI_Datatype), INTENT(INOUT) :: datatype

30 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

31 Fortran binding

32 MPI_TYPE_COMMIT(DATATYPE, IERROR)

33 INTEGER DATATYPE, IERROR

34 The commit operation commits the datatype, that is, the formal description of a com-
35 munication buffer, not the content of that buffer. Thus, after a datatype has been commit-
36 ted, it can be repeatedly reused to communicate the changing content of a buffer or, indeed,
37 the content of different buffers, with different starting addresses.

Advice to implementors. The system may “compile” at commit time an internal representation for the datatype that facilitates communication, e.g., change from a compacted representation to a flat representation of the datatype, and select the most convenient transfer mechanism.

The optimizations chosen during `MPI_TYPE_COMMIT` may no longer be optimal if a session (or the World Model) is initialized or finalized. (*End of advice to implementors.*)

`MPI_TYPE_COMMIT` will accept a committed datatype; in this case, it is equivalent to a no-op.

Example 5.10. The following code fragment gives examples of using `MPI_TYPE_COMMIT`.

```

INTEGER type1, type2
CALL MPI_TYPE_CONTIGUOUS(5, MPI_REAL, type1, ierr)
    ! new type object created
CALL MPI_TYPE_COMMIT(type1, ierr)
    ! now type1 can be used for communication
type2 = type1
    ! type2 can be used for communication
    ! (it is a handle to same object as type1)
CALL MPI_TYPE_VECTOR(3, 5, 4, MPI_REAL, type1, ierr)
    ! new uncommitted type object created
CALL MPI_TYPE_COMMIT(type1, ierr)
    ! now type1 can be used anew for communication

```

`MPI_TYPE_FREE(datatype)`

INOUT datatype datatype that is freed (handle)

C binding

```
int MPI_Type_free(MPI_Datatype *datatype)
```

Fortran 2008 binding

```
MPI_Type_free(datatype, ierror)
    TYPE(MPI_Datatype), INTENT(INOUT) :: datatype
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_FREE(DATATYPE, IERROR)
    INTEGER DATATYPE, IERROR
```

Marks the datatype object associated with `datatype` for deallocation and sets `datatype` to `MPI_DATATYPE_NULL`. Any communication that is currently using this datatype will complete normally. Freeing a datatype does not affect any other datatype that was built from the freed datatype. The system behaves as if input datatype arguments to derived datatype constructors are passed by value.

Advice to implementors. The implementation may keep a reference count of active communications that use the datatype, in order to decide when to free it. Also, one

1 may implement constructors of derived datatypes so that they keep pointers to their
 2 datatype arguments, rather than copying them. In this case, one needs to keep track
 3 of active datatype definition references in order to know when a datatype object can
 4 be freed. (*End of advice to implementors.*)

6 5.1.10 Duplicating a Datatype

9 MPI_TYPE_DUP(oldtype, newtype)

11 IN oldtype datatype (handle)
 12 OUT newtype copy of oldtype (handle)

14 C binding

15 int MPI_Type_dup(MPI_Datatype oldtype, MPI_Datatype *newtype)

17 Fortran 2008 binding

18 MPI_Type_dup(oldtype, newtype, ierror)
 19 TYPE(MPI_Datatype), INTENT(IN) :: oldtype
 20 TYPE(MPI_Datatype), INTENT(OUT) :: newtype
 21 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

22 Fortran binding

23 MPI_TYPE_DUP(OLDTYPE, NEWTYPE, IERROR)
 24 INTEGER OLDTYPE, NEWTYPE, IERROR

26 MPI_TYPE_DUP is a type constructor that duplicates the existing `oldtype` with asso-
 27 ciated key values. For each key value, the respective copy callback function determines the
 28 attribute value associated with this key in the new datatype; one particular action that a
 29 copy callback may take is to delete the attribute from the new datatype. Returns in `newtype`
 30 a new datatype with exactly the same properties as `oldtype` and any copied cached infor-
 31 mation, see Section 7.7.4. The new datatype has identical upper bound and lower bound
 32 and yields the same net result when fully decoded with the procedures in Section 5.1.13.
 33 The `newtype` has the same committed state as the old `oldtype`.

35 5.1.11 Use of General Datatypes in Communication

36 Handles to derived datatypes can be passed to a communication call wherever a datatype
 37 argument is required. A call of the form `MPI_SEND(buf, count, datatype, ...)`, where `count >`
 38 `1`, is interpreted as if the call was passed a new datatype that is the concatenation of `count`
 39 copies of `datatype`. Thus, `MPI_SEND(buf, count, datatype, dest, tag, comm)` is equivalent to,

```
41 MPI_TYPE_CONTIGUOUS(count, datatype, newtype)
42 MPI_TYPE_COMMIT(newtype)
43 MPI_SEND(buf, 1, newtype, dest, tag, comm)
44 MPI_TYPE_FREE(newtype).
```

45 Similar statements apply to all other communication procedures that have a `count` and
 46 `datatype` argument.

Suppose that a send operation `MPI_SEND(buf, count, datatype, dest, tag, comm)` is executed, where `datatype` has type map,

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

and extent *extent*. (Explicit lower bound and upper bound markers are not listed in the type map, but they affect the value of *extent*.) The send operation sends $n \cdot \text{count}$ entries, where entry $i \cdot n + j$ is at location $addr_{i,j} = \text{buf} + \text{extent} \cdot i + disp_j$ and has type $type_j$, for $i = 0, \dots, \text{count} - 1$ and $j = 0, \dots, n - 1$. These entries need not be contiguous, nor distinct; their order can be arbitrary.

The variable stored at address $addr_{i,j}$ in the calling program should be of a type that matches $type_j$, where type matching is defined as in Section 3.3.1. The message sent contains $n \cdot \text{count}$ entries, where entry $i \cdot n + j$ has type $type_j$.

Similarly, suppose that a receive operation `MPI_RECV(buf, count, datatype, source, tag, comm, status)` is executed, where `datatype` has type map,

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

with extent *extent*. (Again, explicit lower bound and upper bound markers are not listed in the type map, but they affect the value of *extent*.) This receive operation receives $n \cdot \text{count}$ entries, where entry $i \cdot n + j$ is at location $\text{buf} + \text{extent} \cdot i + disp_j$ and has type $type_j$. If the incoming message consists of k elements, then we must have $k \leq n \cdot \text{count}$; the $i \cdot n + j$ -th element of the message should have a type that matches $type_j$.

Type matching is defined according to the type signature of the corresponding datatypes, that is, the sequence of basic type components. Type matching does not depend on some aspects of the datatype definition, such as the displacements (layout in memory) or the intermediate types used.

Example 5.11. This example shows that type matching is defined in terms of the basic types that a derived type consists of.

```

...
CALL MPI_TYPE_CONTIGUOUS(2, MPI_REAL, type2, ...)
CALL MPI_TYPE_CONTIGUOUS(4, MPI_REAL, type4, ...)
CALL MPI_TYPE_CONTIGUOUS(2, type2, type22, ...)
...
CALL MPI_SEND(a, 4, MPI_REAL, ...)
CALL MPI_SEND(a, 2, type2, ...)
CALL MPI_SEND(a, 1, type22, ...)
CALL MPI_SEND(a, 1, type4, ...)
...
CALL MPI_RECV(a, 4, MPI_REAL, ...)
CALL MPI_RECV(a, 2, type2, ...)
CALL MPI_RECV(a, 1, type22, ...)
CALL MPI_RECV(a, 1, type4, ...)

```

Each of the sends matches any of the receives.

A datatype may specify overlapping entries. The use of such a datatype in any communication in association with a buffer updated by the operation is erroneous. (This is erroneous even if the actual message received is short enough not to write any entry more than once.)

Suppose that `MPI_RECV(buf, count, datatype, dest, tag, comm, status)` is executed, where `datatype` has type map,

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\}.$$

The received message need not fill all the receive buffer, nor does it need to fill a number of locations that is a multiple of n . Any number, k , of basic elements can be received, where $0 \leq k \leq \text{count} \cdot n$. The number of basic elements received can be retrieved from `status` using the query procedure `MPI_GET_ELEMENTS`.

`MPI_GET_ELEMENTS(status, datatype, count)`

IN	<code>status</code>	return status of receive operation (status)
IN	<code>datatype</code>	datatype used by receive operation (handle)
OUT	<code>count</code>	number of received basic elements (integer)

C binding

```
int MPI_Get_elements(const MPI_Status *status, MPI_Datatype datatype,
                    int *count)
```

```
int MPI_Get_elements_c(const MPI_Status *status, MPI_Datatype datatype,
                       MPI_Count *count)
```

Fortran 2008 binding

```
MPI_Get_elements(status, datatype, count, ierror)
```

```
TYPE(MPI_Status), INTENT(IN) :: status
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(OUT) :: count
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Get_elements(status, datatype, count, ierror) !(_c)
```

```
TYPE(MPI_Status), INTENT(IN) :: status
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: count
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GET_ELEMENTS(STATUS, DATATYPE, COUNT, IERROR)
```

```
INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR
```

The `datatype` argument should match the argument provided by the receive call that set the `status` variable. For both procedures, if the `OUT` parameter cannot express the value to be returned (e.g., if the parameter is too small to hold the output value), it is set to `MPI_UNDEFINED`.

The previously defined procedure `MPI_GET_COUNT` (Section 3.2.5), has a different behavior. It returns the number of “top-level entries” received, i.e., the number of “copies” of type `datatype`. In the previous example, `MPI_GET_COUNT` may return any integer value k , where $0 \leq k \leq \text{count}$. If `MPI_GET_COUNT` returns k , then the number of basic elements received (and the value returned by `MPI_GET_ELEMENTS`) is $n \cdot k$. If the number of basic elements received is not a multiple of n , that is, if the receive operation has not received

an integral number of datatype “copies,” then `MPI_GET_COUNT` sets the value of `count` to `MPI_UNDEFINED`.

Example 5.12. Usage of `MPI_GET_COUNT` and `MPI_GET_ELEMENTS`.

```

...
CALL MPI_TYPE_CONTIGUOUS(2, MPI_REAL, Type2, ierr)
CALL MPI_TYPE_COMMIT(Type2, ierr)
...
CALL MPI_COMM_RANK(comm, rank, ierr)
IF (rank .EQ. 0) THEN
  CALL MPI_SEND(a, 2, MPI_REAL, 1, 0, comm, ierr)
  CALL MPI_SEND(a, 3, MPI_REAL, 1, 0, comm, ierr)
ELSE IF (rank .EQ. 1) THEN
  CALL MPI_RECV(a, 2, Type2, 0, 0, comm, stat, ierr)
  CALL MPI_GET_COUNT(stat, Type2, i, ierr)      ! returns i=1
  CALL MPI_GET_ELEMENTS(stat, Type2, i, ierr)  ! returns i=2
  CALL MPI_RECV(a, 2, Type2, 0, 0, comm, stat, ierr)
  CALL MPI_GET_COUNT(stat, Type2, i, ierr)      ! returns i=MPI_UNDEFINED
  CALL MPI_GET_ELEMENTS(stat, Type2, i, ierr)  ! returns i=3
END IF

```

The procedure `MPI_GET_ELEMENTS` can also be used after a probe to find the number of elements in the probed message. Note that the `MPI_GET_COUNT` and `MPI_GET_ELEMENTS` return the same values when they are used with basic datatypes as long as the limits of their respective `count` arguments are not exceeded.

Rationale. The extension given to the definition of `MPI_GET_COUNT` seems natural: one would expect this procedure to return the value of the `count` argument, when the receive buffer is filled. Sometimes `datatype` represents a basic unit of data one wants to transfer, for example, a record in an array of records (structures). One should be able to find out how many components were received without bothering to divide by the number of elements in each component. However, on other occasions, `datatype` is used to define a complex layout of data in the receiver memory, and does not represent a basic unit of data for transfers. In such cases, one needs to use the procedure `MPI_GET_ELEMENTS`. (*End of rationale.*)

Advice to implementors. The definition implies that a receive cannot change the value of storage outside the entries defined to compose the communication buffer. In particular, the definition implies that padding space in a structure should not be modified when such a structure is copied from one process to another. This would prevent the obvious optimization of copying the structure, together with the padding, as one contiguous block. The implementation is free to do this optimization when it does not impact the outcome of the computation. The user can “force” this optimization by explicitly including padding as part of the message. (*End of advice to implementors.*)

5.1.12 Correct Use of Addresses

Successively declared variables in C or Fortran are not necessarily stored at contiguous locations. Thus, care must be exercised that displacements do not cross from one variable to another. Also, in machines with a segmented address space, addresses are not unique and address arithmetic has some peculiar properties. Thus, the use of **addresses**, that is, displacements relative to the start address `MPI_BOTTOM`, has to be restricted.

1 Variables belong to the same **sequential storage** if they belong to the same array, to
2 the same COMMON block in Fortran, or to the same structure in C. Valid addresses are defined
3 recursively as follows:
4

- 5 1. The procedure MPI_GET_ADDRESS returns a valid address, when passed as argument
6 a variable of the calling program.
- 7 2. The `buf` argument of a communication procedure evaluates to a valid address, when
8 passed as argument a variable of the calling program.
- 9 3. If `v` is a valid address, and `i` is an integer, then `v+i` is a valid address, provided `v` and
10 `v+i` are in the same sequential storage.
11
12

13 A correct program uses only valid addresses to identify the locations of entries in
14 communication buffers. Furthermore, if `u` and `v` are two valid addresses, then the (integer)
15 difference `u - v` can be computed only if both `u` and `v` are in the same sequential storage.
16 No other arithmetic operations can be meaningfully executed on addresses.

17 The rules above impose no constraints on the use of derived datatypes, as long as
18 they are used to define a communication buffer that is wholly contained within the same
19 sequential storage. However, the construction of a communication buffer that contains
20 variables that are not within the same sequential storage must obey certain restrictions.
21 Basically, a communication buffer with variables that are not within the same sequential
22 storage can be used only by specifying in the communication call `buf = MPI_BOTTOM`, `count`
23 `= 1`, and using a `datatype` argument where all displacements are valid (absolute) addresses.
24

25 *Advice to users.* It is not expected that MPI implementations will be able to de-
26 tect erroneous, “out of bound” displacements—unless those overflow the user address
27 space—since the MPI call may not know the extent of the arrays and records in the
28 host program. (*End of advice to users.*)
29

30 *Advice to implementors.* There is no need to distinguish (absolute) addresses and
31 (relative) displacements on a machine with contiguous address space: `MPI_BOTTOM`
32 is zero, and both addresses and displacements are integers. On machines where the
33 distinction is required, addresses are recognized as expressions that involve
34 `MPI_BOTTOM`. (*End of advice to implementors.*)
35

36 5.1.13 Decoding a Datatype

37 MPI datatype objects allow users to specify an arbitrary layout of data in memory. There are
38 several cases where accessing the layout information in opaque datatype objects would be
39 useful. The opaque datatype object has found a number of uses outside MPI. Furthermore,
40 a number of tools wish to display internal information about a datatype. To achieve this,
41 datatype decoding procedures are provided. The two procedures in this section are used
42 together to decode datatypes to recreate the calling sequence used in their initial definition.
43 These can be used to allow a user to determine the type map and type signature of a
44 datatype.
45
46
47
48

MPI_TYPE_GET_ENVELOPE(datatype, num_integers, num_addresses, num_large_counts,			1
num_datatypes, combiner)			2
IN	datatype	datatype to decode (handle)	3
OUT	num_integers	number of input integers used in call constructing	4
		combiner (non-negative integer)	5
OUT	num_addresses	number of input addresses used in call constructing	6
		combiner (non-negative integer)	7
OUT	num_large_counts	number of input large counts used in call	8
		constructing combiner (non-negative integer, only	9
		present for large count variants)	10
OUT	num_datatypes	number of input datatypes used in call constructing	11
		combiner (non-negative integer)	12
OUT	combiner	combiner (state)	13
			14
			15
			16

C binding

```
int MPI_Type_get_envelope(MPI_Datatype datatype, int *num_integers,
    int *num_addresses, int *num_datatypes, int *combiner)
```

```
int MPI_Type_get_envelope_c(MPI_Datatype datatype, MPI_Count *num_integers,
    MPI_Count *num_addresses, MPI_Count *num_large_counts,
    MPI_Count *num_datatypes, int *combiner)
```

Fortran 2008 binding

```
MPI_Type_get_envelope(datatype, num_integers, num_addresses, num_datatypes,
    combiner, ierror)
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
INTEGER, INTENT(OUT) :: num_integers, num_addresses, num_datatypes,
    combiner
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Type_get_envelope(datatype, num_integers, num_addresses, num_large_counts,
    num_datatypes, combiner, ierror) !(_c)
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: num_integers, num_addresses,
    num_large_counts, num_datatypes
```

```
INTEGER, INTENT(OUT) :: combiner
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_GET_ENVELOPE(DATATYPE, NUM_INTEGERS, NUM_ADDRESSES, NUM_DATATYPES,
    COMBINER, IERROR)
```

```
INTEGER DATATYPE, NUM_INTEGERS, NUM_ADDRESSES, NUM_DATATYPES, COMBINER,
```

```
IERROR
```

For the given datatype, MPI_TYPE_GET_ENVELOPE returns information on the number and type of input arguments used in the call that created the datatype. The number-of-arguments values returned can be used to provide sufficiently large arrays in the decoding routine MPI_TYPE_GET_CONTENTS. This call and the meaning of the returned values is

described below. The combiner reflects the MPI datatype constructor call that was used in creating datatype.

Rationale. By requiring that the combiner reflect the constructor used in the creation of the datatype, the decoded information can be used to effectively recreate the calling sequence used in the original creation. This is the most useful information and was felt to be reasonable even though it constrains implementations to remember the original constructor sequence even if the internal representation is different.

The decoded information keeps track of datatype duplications. This is important as one needs to distinguish between a predefined datatype and a dup of a predefined datatype. The former is a constant object that cannot be freed, while the latter is a derived datatype that can be freed. (*End of rationale.*)

The list of values that can be returned from `MPI_TYPE_GET_ENVELOPE` in combiner (on the left) and the call associated with them (on the right) are as follows:

<code>MPI_COMBINER_NAMED</code>	a named predefined datatype
<code>MPI_COMBINER_DUP</code>	<code>MPI_TYPE_DUP</code>
<code>MPI_COMBINER_CONTIGUOUS</code>	<code>MPI_TYPE_CONTIGUOUS</code>
<code>MPI_COMBINER_VECTOR</code>	<code>MPI_TYPE_VECTOR</code>
<code>MPI_COMBINER_HVECTOR</code>	<code>MPI_TYPE_CREATE_HVECTOR</code>
<code>MPI_COMBINER_INDEXED</code>	<code>MPI_TYPE_INDEXED</code>
<code>MPI_COMBINER_HINDEXED</code>	<code>MPI_TYPE_CREATE_HINDEXED</code>
<code>MPI_COMBINER_INDEXED_BLOCK</code>	<code>MPI_TYPE_CREATE_INDEXED_BLOCK</code>
<code>MPI_COMBINER_HINDEXED_BLOCK</code>	<code>MPI_TYPE_CREATE_HINDEXED_BLOCK</code>
<code>MPI_COMBINER_STRUCT</code>	<code>MPI_TYPE_CREATE_STRUCT</code>
<code>MPI_COMBINER_SUBARRAY</code>	<code>MPI_TYPE_CREATE_SUBARRAY</code>
<code>MPI_COMBINER_DARRAY</code>	<code>MPI_TYPE_CREATE_DARRAY</code>
<code>MPI_COMBINER_F90_REAL</code>	<code>MPI_TYPE_CREATE_F90_REAL</code>
<code>MPI_COMBINER_F90_COMPLEX</code>	<code>MPI_TYPE_CREATE_F90_COMPLEX</code>
<code>MPI_COMBINER_F90_INTEGER</code>	<code>MPI_TYPE_CREATE_F90_INTEGER</code>
<code>MPI_COMBINER_RESIZED</code>	<code>MPI_TYPE_CREATE_RESIZED</code>
<code>MPI_COMBINER_VALUE_INDEX</code>	<code>MPI_TYPE_GET_VALUE_INDEX</code>

If combiner is `MPI_COMBINER_NAMED` then datatype is a named predefined datatype.

If the `MPI_TYPE_GET_ENVELOPE` variant without `num_large_counts` is invoked with a datatype that requires an output value of `num_large_counts > 0`, then an error of class `MPI_ERR_TYPE` is raised.

Rationale. The large count variant of this MPI procedure was added in MPI-4. It contains a new `num_large_counts` parameter. The other variant—the variant that existed before MPI-4—was not changed in order to preserve backwards compatibility. (*End of rationale.*)

The actual arguments used in the creation call for a datatype can be obtained using `MPI_TYPE_GET_CONTENTS`.

`MPI_TYPE_GET_ENVELOPE` and `MPI_TYPE_GET_CONTENTS` also support large count types in separate additional MPI procedures in C (suffixed with the “_c”) and interface polymorphism in Fortran when using `USE mpi_f08`.

```
MPI_TYPE_GET_CONTENTS(datatype, max_integers, max_addresses, max_large_counts,
    max_datatypes, array_of_integers, array_of_addresses, array_of_large_counts,
    array_of_datatypes)
```

			1
			2
			3
IN	datatype	datatype to decode (handle)	4
			5
IN	max_integers	number of elements in array_of_integers (non-negative integer)	6
			7
IN	max_addresses	number of elements in array_of_addresses (non-negative integer)	8
			9
IN	max_large_counts	number of elements in array_of_large_counts (non-negative integer, only present for large count variants)	10
			11
			12
IN	max_datatypes	number of elements in array_of_datatypes (non-negative integer)	13
			14
			15
OUT	array_of_integers	contains integer arguments used in constructing datatype (array of integers)	16
			17
OUT	array_of_addresses	contains address arguments used in constructing datatype (array of integers)	18
			19
OUT	array_of_large_counts	contains large count arguments used in constructing datatype (array of integers, only present for large count variants)	20
			21
			22
OUT	array_of_datatypes	contains datatype arguments used in constructing datatype (array of handles)	23
			24
			25

C binding

```
int MPI_Type_get_contents(MPI_Datatype datatype, int max_integers,
    int max_addresses, int max_datatypes, int array_of_integers[],
    MPI_Aint array_of_addresses[], MPI_Datatype array_of_datatypes[])
```

```
int MPI_Type_get_contents_c(MPI_Datatype datatype, MPI_Count max_integers,
    MPI_Count max_addresses, MPI_Count max_large_counts,
    MPI_Count max_datatypes, int array_of_integers[],
    MPI_Aint array_of_addresses[], MPI_Count array_of_large_counts[],
    MPI_Datatype array_of_datatypes[])
```

Fortran 2008 binding

```
MPI_Type_get_contents(datatype, max_integers, max_addresses, max_datatypes,
    array_of_integers, array_of_addresses, array_of_datatypes,
    ierror)
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: max_integers, max_addresses, max_datatypes
INTEGER, INTENT(OUT) :: array_of_integers(max_integers)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) ::
    array_of_addresses(max_addresses)
TYPE(MPI_Datatype), INTENT(OUT) :: array_of_datatypes(max_datatypes)
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```

1 MPI_Type_get_contents(datatype, max_integers, max_addresses, max_large_counts,
2     max_datatypes, array_of_integers, array_of_addresses,
3     array_of_large_counts, array_of_datatypes, ierror) !(_c)
4     TYPE(MPI_Datatype), INTENT(IN) :: datatype
5     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: max_integers, max_addresses,
6     max_large_counts, max_datatypes
7     INTEGER, INTENT(OUT) :: array_of_integers(max_integers)
8     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) ::
9     array_of_addresses(max_addresses)
10    INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) ::
11    array_of_large_counts(max_large_counts)
12    TYPE(MPI_Datatype), INTENT(OUT) :: array_of_datatypes(max_datatypes)
13    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

14
15 MPI_TYPE_GET_CONTENTS(DATATYPE, MAX_INTEGERS, MAX_ADDRESSES, MAX_DATATYPES,
16     ARRAY_OF_INTEGERS, ARRAY_OF_ADDRESSES, ARRAY_OF_DATATYPES,
17     IERROR)
18     INTEGER DATATYPE, MAX_INTEGERS, MAX_ADDRESSES, MAX_DATATYPES,
19     ARRAY_OF_INTEGERS(*), ARRAY_OF_DATATYPES(*), IERROR
20     INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_ADDRESSES(*)
21

```

`datatype` must be a predefined unnamed or a derived datatype; the call is erroneous if `datatype` is a predefined named datatype.

The values given for `max_integers`, `max_addresses`, `max_large_counts`, and `max_datatypes` must be at least as large as the value returned in `num_integers`, `num_addresses`, `num_large_counts`, and `num_datatypes`, respectively, in the call `MPI_TYPE_GET_ENVELOPE` for the same `datatype` argument.

Rationale. The arguments `max_integers`, `max_addresses`, `max_large_counts`, and `max_datatypes` allow for error checking in the call. (*End of rationale.*)

If the `MPI_TYPE_GET_CONTENTS` variant without `max_large_counts` is invoked with a `datatype` that requires > 0 values in `array_of_large_counts`, then an error of class `MPI_ERR_TYPE` is raised.

Rationale. The large count variant of this MPI procedure was added in MPI-4. It contains new `max_large_counts` and `array_of_large_counts` parameters. The other variant—the variant that existed before MPI-4—was not changed in order to preserve backwards compatibility. (*End of rationale.*)

The datatypes returned in `array_of_datatypes` are handles to datatype objects that are equivalent to the datatypes used in the original construction call. If these were derived datatypes, then the returned datatypes are new datatype objects, and the user is responsible for freeing these datatypes with `MPI_TYPE_FREE`. If these were predefined datatypes, then the returned datatype is equal to that (constant) predefined datatype and cannot be freed.

The committed state of returned derived datatypes is undefined, i.e., the datatypes may or may not be committed. Furthermore, the content of attributes of returned datatypes is undefined.

Note that `MPI_TYPE_GET_CONTENTS` can be invoked with a `datatype` argument that was constructed using `MPI_TYPE_CREATE_F90_REAL`, `MPI_TYPE_CREATE_F90_INTEGER`, or `MPI_TYPE_CREATE_F90_COMPLEX` (an unnamed predefined datatype). In such a case, an empty `array_of_datatypes` is returned.

Rationale. The definition of datatype equivalence implies that equivalent predefined datatypes are equal. By requiring the same handle for named predefined datatypes, it is possible to use the `==` or `.EQ.` comparison operator to determine the datatype involved. (*End of rationale.*)

Advice to implementors. The datatypes returned in `array_of_datatypes` must appear to the user as if each is an equivalent copy of the datatype used in the type constructor call. Whether this is done by creating a new datatype or via another mechanism such as a reference count mechanism is up to the implementation as long as the semantics are preserved. (*End of advice to implementors.*)

Rationale. The committed state and attributes of the returned datatype is deliberately left vague. The datatype used in the original construction may have been modified since its use in the constructor call. Attributes can be added, removed, or modified as well as having the datatype committed. The semantics given allow for a reference count implementation without having to track these changes. (*End of rationale.*)

In the deprecated datatype constructor calls, the address arguments in Fortran are of type `INTEGER`. In the preferred calls, the address arguments are of type `INTEGER(KIND=MPI_ADDRESS_KIND)`. The call `MPI_TYPE_GET_CONTENTS` returns all addresses in an argument of type `INTEGER(KIND=MPI_ADDRESS_KIND)`. This is true even if the deprecated calls were used. Thus, the location of values returned can be thought of as being returned by the C bindings. It can also be determined by examining the preferred calls for datatype constructors for the deprecated calls that involve addresses.

Rationale. By having all address arguments returned in the `array_of_addresses` argument, the result from a C and Fortran decoding of a `datatype` gives the result in the same argument. It is assumed that an integer of type `INTEGER(KIND=MPI_ADDRESS_KIND)` will be at least as large as the `INTEGER` argument used in datatype construction with the old MPI-1 calls so no loss of information will occur. (*End of rationale.*)

The following defines what values are placed in each entry of the returned arrays depending on the datatype constructor used for `datatype`. It also specifies the size of the arrays needed, which is the values returned by `MPI_TYPE_GET_ENVELOPE`. In Fortran, the following calls were made:

```

PARAMETER (LARGE = 1000)
INTEGER DTYPE, NI, NA, ND, COMBINER, I(LARGE), D(LARGE), IERROR
INTEGER(KIND=MPI_ADDRESS_KIND) A(LARGE)
! CONSTRUCT DATATYPE DTYPE (NOT SHOWN)
CALL MPI_TYPE_GET_ENVELOPE(DTYPE, NI, NA, ND, COMBINER, IERROR)
IF ((NI .GT. LARGE) .OR. (NA .GT. LARGE) .OR. (ND .GT. LARGE)) THEN
    WRITE (*, *) "NI, NA, OR ND = ", NI, NA, ND, &

```

```

1  " RETURNED BY MPI_TYPE_GET_ENVELOPE IS LARGER THAN LARGE = ", LARGE
2  CALL MPI_ABORT(MPI_COMM_WORLD, 99, IERROR)
3  ENDF
4  CALL MPI_TYPE_GET_CONTENTS(DTYPE, NI, NA, ND, I, A, D, IERROR)

```

or in C the analogous calls of:

```

6  #define LARGE 1000
7  int ni, na, nd, combiner, i[LARGE];
8  MPI_Aint a[LARGE];
9  MPI_Datatype dtype, d[LARGE];
10 /* construct datatype dtype (not shown) */
11 MPI_Type_get_envelope(dtype, &ni, &na, &nd, &combiner);
12 if ((ni > LARGE) || (na > LARGE) || (nd > LARGE)) {
13     fprintf(stderr, "ni, na, or nd = %d %d %d returned by ", ni, na, nd);
14     fprintf(stderr, "MPI_Type_get_envelope is larger than LARGE = %d\n",
15             LARGE);
16     MPI_Abort(MPI_COMM_WORLD, 99);
17 }
18 MPI_Type_get_contents(dtype, ni, na, nd, i, a, d);

```

The following describes the values of the arguments for each combiner. The lower case name of arguments is used. Also, the descriptions below refer to MPI datatypes created by procedures without large count arguments.

MPI_COMBINER_NAMED the datatype represent a predefined type and therefore it is erroneous to call MPI_TYPE_GET_CONTENTS.

MPI_COMBINER_DUP ni = 0, na = 0, nd = 1, and

Constructor argument	C	Fortran location
oldtype	d[0]	D(1)

MPI_COMBINER_CONTIGUOUS ni = 1, na = 0, nd = 1, and

Constructor argument	C	Fortran location
count	i[0]	I(1)
oldtype	d[0]	D(1)

MPI_COMBINER_VECTOR ni = 3, na = 0, nd = 1, and

Constructor argument	C	Fortran location
count	i[0]	I(1)
blocklength	i[1]	I(2)
stride	i[2]	I(3)
oldtype	d[0]	D(1)

MPI_COMBINER_HVECTOR ni = 2, na = 1, nd = 1, and

Constructor argument	C	Fortran location
count	i[0]	I(1)
blocklength	i[1]	I(2)
stride	a[0]	A(1)
oldtype	d[0]	D(1)

MPI_COMBINER_INDEXED $ni = 2*count+1$, $na = 0$, $nd = 1$, and

Constructor argument	C	Fortran location
count	$i[0]$	I(1)
array_of_blocklengths	$i[1]$ to $i[i[0]]$	I(2) to I(I(1)+1)
array_of_displacements	$i[i[0]+1]$ to $i[2*i[0]]$	I(I(1)+2) to I(2*I(1)+1)
oldtype	$d[0]$	D(1)

MPI_COMBINER_HINDEXED $ni = count+1$, $na = count$, $nd = 1$, and

Constructor argument	C	Fortran location
count	$i[0]$	I(1)
array_of_blocklengths	$i[1]$ to $i[i[0]]$	I(2) to I(I(1)+1)
array_of_displacements	$a[0]$ to $a[i[0]-1]$	A(1) to A(I(1))
oldtype	$d[0]$	D(1)

MPI_COMBINER_INDEXED_BLOCK $ni = count+2$, $na = 0$, $nd = 1$, and

Constructor argument	C	Fortran location
count	$i[0]$	I(1)
blocklength	$i[1]$	I(2)
array_of_displacements	$i[2]$ to $i[i[0]+1]$	I(3) to I(I(1)+2)
oldtype	$d[0]$	D(1)

MPI_COMBINER_HINDEXED_BLOCK $ni = 2$, $na = count$, $nd = 1$, and

Constructor argument	C	Fortran location
count	$i[0]$	I(1)
blocklength	$i[1]$	I(2)
array_of_displacements	$a[0]$ to $a[i[0]-1]$	A(1) to A(I(1))
oldtype	$d[0]$	D(1)

MPI_COMBINER_STRUCT $ni = count+1$, $na = count$, $nd = count$, and

Constructor argument	C	Fortran location
count	$i[0]$	I(1)
array_of_blocklengths	$i[1]$ to $i[i[0]]$	I(2) to I(I(1)+1)
array_of_displacements	$a[0]$ to $a[i[0]-1]$	A(1) to A(I(1))
array_of_types	$d[0]$ to $d[i[0]-1]$	D(1) to D(I(1))

MPI_COMBINER_SUBARRAY $ni = 3*ndims+2$, $na = 0$, $nd = 1$, and

Constructor argument	C	Fortran location
ndims	$i[0]$	I(1)
array_of_sizes	$i[1]$ to $i[i[0]]$	I(2) to I(I(1)+1)
array_of_subsizes	$i[i[0]+1]$ to $i[2*i[0]]$	I(I(1)+2) to I(2*I(1)+1)
array_of_starts	$i[2*i[0]+1]$ to $i[3*i[0]]$	I(2*I(1)+2) to I(3*I(1)+1)
order	$i[3*i[0]+1]$	I(3*I(1)+2)
oldtype	$d[0]$	D(1)

1 **MPI_COMBINER_DARRAY** $n_i = 4 * ndims + 4$, $n_a = 0$, $n_d = 1$, and

Constructor argument	C	Fortran location
size	$i[0]$	I(1)
rank	$i[1]$	I(2)
ndims	$i[2]$	I(3)
array_of_gsizes	$i[3]$ to $i[i[2]+2]$	I(4) to I(I(3)+3)
array_of_distribs	$i[i[2]+3]$ to $i[2*i[2]+2]$	I(I(3)+4) to I(2*I(3)+3)
array_of_dargs	$i[2*i[2]+3]$ to $i[3*i[2]+2]$	I(2*I(3)+4) to I(3*I(3)+3)
array_of_psizes	$i[3*i[2]+3]$ to $i[4*i[2]+2]$	I(3*I(3)+4) to I(4*I(3)+3)
order	$i[4*i[2]+3]$	I(4*I(3)+4)
oldtype	$d[0]$	D(1)

13 **MPI_COMBINER_F90_REAL** $n_i = 2$, $n_a = 0$, $n_d = 0$, and

Constructor argument	C	Fortran location
p	$i[0]$	I(1)
r	$i[1]$	I(2)

18 **MPI_COMBINER_F90_COMPLEX** $n_i = 2$, $n_a = 0$, $n_d = 0$, and

Constructor argument	C	Fortran location
p	$i[0]$	I(1)
r	$i[1]$	I(2)

23 **MPI_COMBINER_F90_INTEGER** $n_i = 1$, $n_a = 0$, $n_d = 0$, and

Constructor argument	C	Fortran location
r	$i[0]$	I(1)

27 **MPI_COMBINER_RESIZED** $n_i = 0$, $n_a = 2$, $n_d = 1$, and

Constructor argument	C	Fortran location
lb	$a[0]$	A(1)
extent	$a[1]$	A(2)
oldtype	$d[0]$	D(1)

33 **MPI_COMBINER_VALUE_INDEX** $n_i = 0$, $n_a = 0$, $n_d = 2$, and

Constructor argument	C	Fortran location
value_type	$d[0]$	D(1)
index_type	$d[1]$	D(2)

39 5.1.14 Examples

40 The following examples illustrate the use of derived datatypes.

42 **Example 5.13.** Send and receive a section of a 3D array.

```
43
44 REAL a(100,100,100), e(9,9,9)
45 INTEGER oneslice, twoslice, threeslice, myrank, ierr
46 INTEGER(KIND=MPI_ADDRESS_KIND) lb, sizeofreal
47 INTEGER status(MPI_STATUS_SIZE)
48
```

```

1  ! extract the section a(1:17:2, 3:11, 2:10)
2  ! and store it in e(:, :, :).
3
4  CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank, ierr)
5
6  CALL MPI_TYPE_GET_EXTENT(MPI_REAL, lb, sizeofreal, ierr)
7
8  ! create datatype for a 1D section
9  CALL MPI_TYPE_VECTOR(9, 1, 2, MPI_REAL, oneslice, ierr)
10
11 ! create datatype for a 2D section
12 CALL MPI_TYPE_CREATE_HVECTOR(9, 1, 100*sizeofreal, oneslice, &
13                               twoslice, ierr)
14
15 ! create datatype for the entire section
16 CALL MPI_TYPE_CREATE_HVECTOR(9, 1, 100*100*sizeofreal, twoslice, &
17                               threeslice, ierr)
18
19 CALL MPI_TYPE_COMMIT(threeslice, ierr)
20 CALL MPI_SENDRECV(a(1,3,2), 1, threeslice, myrank, 0, e, 9*9*9, &
21                   MPI_REAL, myrank, 0, MPI_COMM_WORLD, status, ierr)

```

Example 5.14. Copy the (strictly) lower triangular part of a matrix.

```

22
23 REAL a(100,100), b(100,100)
24 INTEGER disp(100), blocklen(100), ltype, myrank, ierr
25 INTEGER status(MPI_STATUS_SIZE)
26
27 ! copy lower triangular part of array a
28 ! onto lower triangular part of array b
29
30 CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank, ierr)
31
32 ! compute start and size of each column
33 DO i=1,100
34     disp(i) = 100*(i-1) + i
35     blocklen(i) = 100-i
36 END DO
37
38 ! create datatype for lower triangular part
39 CALL MPI_TYPE_INDEXED(100, blocklen, disp, MPI_REAL, ltype, ierr)
40
41 CALL MPI_TYPE_COMMIT(ltype, ierr)
42 CALL MPI_SENDRECV(a, 1, ltype, myrank, 0, b, 1, &
43                   ltype, myrank, 0, MPI_COMM_WORLD, status, ierr)

```

Example 5.15. Transpose a matrix.

```

44
45 REAL a(100,100), b(100,100)
46 INTEGER row, xpose, myrank, ierr
47 INTEGER(KIND=MPI_ADDRESS_KIND) lb, sizeofreal

```

```

1  INTEGER status(MPI_STATUS_SIZE)
2
3  ! transpose matrix a onto b
4
5  CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank, ierr)
6
7  CALL MPI_TYPE_GET_EXTENT(MPI_REAL, lb, sizeofreal, ierr)
8
9  ! create datatype for one row
10 CALL MPI_TYPE_VECTOR(100, 1, 100, MPI_REAL, row, ierr)
11
12 ! create datatype for matrix in row-major order
13 CALL MPI_TYPE_CREATE_HVECTOR(100, 1, sizeofreal, row, xpose, ierr)
14
15 CALL MPI_TYPE_COMMIT(xpose, ierr)
16
17 ! send matrix in row-major order and receive in column major order
18 CALL MPI_SENDRECV(a, 1, xpose, myrank, 0, b, 100*100, &
19                   MPI_REAL, myrank, 0, MPI_COMM_WORLD, status, ierr)

```

Example 5.16. Another approach to the transpose problem:

```

21 REAL a(100,100), b(100,100)
22 INTEGER row, row1
23 INTEGER(KIND=MPI_ADDRESS_KIND) lb, sizeofreal
24 INTEGER myrank, ierr
25 INTEGER status(MPI_STATUS_SIZE)
26
27 CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank, ierr)
28
29 ! transpose matrix a onto b
30
31 CALL MPI_TYPE_GET_EXTENT(MPI_REAL, lb, sizeofreal, ierr)
32
33 ! create datatype for one row
34 CALL MPI_TYPE_VECTOR(100, 1, 100, MPI_REAL, row, ierr)
35
36 ! create datatype for one row, with the extent of one real number
37 lb = 0
38 CALL MPI_TYPE_CREATE_RESIZED(row, lb, sizeofreal, row1, ierr)
39
40 CALL MPI_TYPE_COMMIT(row1, ierr)
41
42 ! send 100 rows and receive in column major order
43 CALL MPI_SENDRECV(a, 100, row1, myrank, 0, b, 100*100, &
44                   MPI_REAL, myrank, 0, MPI_COMM_WORLD, status, ierr)

```

Example 5.17. Use of MPI datatypes to manipulate an array of structures.

```

46 struct Partstruct
47 {

```

```

1   int    type; /* particle type */
2   double d[6]; /* particle coordinates */
3   char   b[7]; /* some additional information */
4 };
5
6 struct Partstruct    particle[1000];
7
8 int                i, dest, tag;
9 MPI_Comm          comm;
10
11 /* build datatype describing structure */
12
13 MPI_Datatype Particlestruct, Particletype;
14 MPI_Datatype type[3] = {MPI_INT, MPI_DOUBLE, MPI_CHAR};
15 int          blocklen[3] = {1, 6, 7};
16 MPI_Aint     disp[3];
17 MPI_Aint     base, lb, sizeofentry;
18
19 /* compute displacements of structure components */
20
21 MPI_Get_address(particle, disp);
22 MPI_Get_address(particle[0].d, disp+1);
23 MPI_Get_address(particle[0].b, disp+2);
24 base = disp[0];
25 for (i=0; i < 3; i++) disp[i] = MPI_Aint_diff(disp[i], base);
26
27 MPI_Type_create_struct(3, blocklen, disp, type, &Particlestruct);
28
29 /* Since the compiler may pad the structure, it is best to explicitly
30    set the extent of the MPI datatype for a structure element using
31    MPI_Type_create_resized */
32
33 /* compute extent of the structure */
34 MPI_Get_address(particle+1, &sizeofentry);
35 sizeofentry = MPI_Aint_diff(sizeofentry, base);
36
37 /* build datatype describing structure */
38 MPI_Type_create_resized(Particlestruct, 0, sizeofentry, &Particletype);
39
40 /* 4.1: send the entire array */
41
42 MPI_Type_commit(&Particletype);
43 MPI_Send(particle, 1000, Particletype, dest, tag, comm);
44
45 /* 4.2: send only the entries of type zero particles,
46    preceded by the number of such entries */
47
48 MPI_Datatype Zparticles; /* datatype describing all particles
                           with type zero (needs to be recomputed

```

```

1                                     if types change) */
2 MPI_Datatype Ztype;
3
4 int          zdisp[1000];
5 int          zblock[1000], j, k;
6 int          zzblock[2] = {1,1};
7 MPI_Aint     zzdisp[2];
8 MPI_Datatype zztype[2];
9
10 /* compute displacements of type zero particles */
11 j = 0;
12 for (i=0; i < 1000; i++)
13     if (particle[i].type == 0)
14     {
15         zdisp[j] = i;
16         zblock[j] = 1;
17         j++;
18     }
19
20 /* create datatype for type zero particles */
21 MPI_Type_indexed(j, zblock, zdisp, Particletype, &Zparticles);
22
23 /* prepend particle count */
24 MPI_Get_address(&j, zzdisp);
25 MPI_Get_address(particle, zzdisp+1);
26 zztype[0] = MPI_INT;
27 zztype[1] = Zparticles;
28 MPI_Type_create_struct(2, zzblock, zzdisp, zztype, &Ztype);
29
30 MPI_Type_commit(&Ztype);
31 MPI_Send(MPI_BOTTOM, 1, Ztype, dest, tag, comm);
32
33 /* A probably more efficient way of defining Zparticles */
34
35 /* consecutive particles with index zero are handled as one block */
36 j=0;
37 for (i=0; i < 1000; i++)
38     if (particle[i].type == 0)
39     {
40         for (k=i+1; (k < 1000)&&(particle[k].type == 0); k++);
41         zdisp[j] = i;
42         zblock[j] = k-i;
43         j++;
44         i = k;
45     }
46 MPI_Type_indexed(j, zblock, zdisp, Particletype, &Zparticles);
47
48 /* 4.3: send the first two coordinates of all entries */
49 MPI_Datatype Allpairs;      /* datatype for all pairs of coordinates */

```



```

1 MPI_Type_get_extent(Particletype, &lb, &sizeofentry);
2
3 /* sizeofentry can also be computed by subtracting the address
4   of particle[0] from the address of particle[1] */
5
6 MPI_Type_create_hvector(1000, 2, sizeofentry, MPI_DOUBLE, &Allpairs);
7 MPI_Type_commit(&Allpairs);
8 MPI_Send(particle[0].d, 1, Allpairs, dest, tag, comm);
9
10 /* an alternative solution to 4.3 */
11
12 MPI_Datatype Twodouble;
13
14 MPI_Type_contiguous(2, MPI_DOUBLE, &Twodouble);
15
16 MPI_Datatype Onepair; /* datatype for one pair of coordinates, with
17                       the extent of one particle entry */
18
19 MPI_Type_create_resized(Twodouble, 0, sizeofentry, &Onepair );
20 MPI_Type_commit(&Onepair);
21 MPI_Send(particle[0].d, 1000, Onepair, dest, tag, comm);

```

Example 5.18. The same manipulations as in the previous example, but use absolute addresses in datatypes.

```

22
23
24
25 struct Partstruct
26 {
27     int    type;
28     double d[6];
29     char   b[7];
30 };
31
32 struct Partstruct particle[1000];
33
34 /* build datatype describing first array entry */
35
36 MPI_Datatype Particletype;
37 MPI_Datatype type[3] = {MPI_INT, MPI_DOUBLE, MPI_CHAR};
38 int          block[3] = {1, 6, 7};
39 MPI_Aint     disp[3];
40
41 MPI_Get_address(particle, disp);
42 MPI_Get_address(particle[0].d, disp+1);
43 MPI_Get_address(particle[0].b, disp+2);
44 MPI_Type_create_struct(3, block, disp, type, &Particletype);
45
46 /* Particletype describes first array entry -- using absolute
47   addresses */
48
49 /* 5.1: send the entire array */
50
51 MPI_Type_commit(&Particletype);

```

```

1  MPI_Send(MPI_BOTTOM, 1000, Particletype, dest, tag, comm);
2
3
4  /* 5.2: send the entries of type zero,
5     preceded by the number of such entries */
6
7  MPI_Datatype Zparticles, Ztype;
8
9  int          zdisp[1000];
10 int          zblock[1000], i, j, k;
11 int          zzbblock[2] = {1,1};
12 MPI_Datatype zztype[2];
13 MPI_Aint     zzdisp[2];
14
15 j=0;
16 for (i=0; i < 1000; i++)
17     if (particle[i].type == 0)
18     {
19         for (k=i+1; (k < 1000)&&(particle[k].type == 0); k++);
20         zdisp[j] = i;
21         zblock[j] = k-i;
22         j++;
23         i = k;
24     }
25 MPI_Type_indexed(j, zblock, zdisp, Particletype, &Zparticles);
26 /* Zparticles describe particles with type zero, using
27    their absolute addresses*/
28
29 /* prepend particle count */
30 MPI_Get_address(&j, zzdisp);
31 zzdisp[1] = (MPI_Aint)0;
32 zztype[0] = MPI_INT;
33 zztype[1] = Zparticles;
34 MPI_Type_create_struct(2, zzbblock, zzdisp, zztype, &Ztype);
35
36 MPI_Type_commit(&Ztype);
37 MPI_Send(MPI_BOTTOM, 1, Ztype, dest, tag, comm);
38
39
40
41
42
43
44
45
46
47
48

```

Example 5.19. This example shows how datatypes can be used to handle unions.

```

37 union {
38     int    ival;
39     float fval;
40 } u[1000];
41
42 int    i, utype;
43
44 /* All entries of u have identical type; variable
45    utype keeps track of their current type */
46
47 MPI_Datatype mpi_utype[2];
48 MPI_Aint     ubase, extent;

```

```

1  /* compute an MPI datatype for each possible union type;
2     assume values are left-aligned in union storage. */
3
4
5  MPI_Get_address(u, &ubase);
6  MPI_Get_address(u+1, &extent);
7  extent = MPI_Aint_diff(extent, ubase);
8
9  MPI_Type_create_resized(MPI_INT, 0, extent, &mpi_utype[0]);
10
11 MPI_Type_create_resized(MPI_FLOAT, 0, extent, &mpi_utype[1]);
12
13 for(i=0; i<2; i++) MPI_Type_commit(&mpi_utype[i]);
14
15 /* actual communication */
16 MPI_Send(u, 1000, mpi_utype[utype], dest, tag, comm);

```

Example 5.20. This example shows how a datatype can be decoded. The routine `printdatatype` prints out the elements of the datatype. Note the use of `MPI_Type_free` for datatypes that are not predefined.

```

17 /*
18  Example of decoding a datatype.
19
20  Returns 0 if the datatype is predefined, 1 otherwise
21  */
22 #include <stdio.h>
23 #include <stdlib.h>
24 #include "mpi.h"
25 int printdatatype(MPI_Datatype datatype)
26 {
27     int *array_of_ints;
28     MPI_Aint *array_of_adds;
29     MPI_Datatype *array_of_dtypes;
30     int num_ints, num_adds, num_dtypes, combiner;
31     int i;
32
33     MPI_Type_get_envelope(datatype,
34                           &num_ints, &num_adds, &num_dtypes, &combiner);
35     switch (combiner) {
36     case MPI_COMBINER_NAMED:
37         printf("Datatype is named:");
38         /* To print the specific type, we can match against the
39            predefined forms. We can NOT use a switch statement here
40            We could also use MPI_TYPE_GET_NAME if we preferred to use
41            names that the user may have changed.
42            */
43         if (datatype == MPI_INT) printf("MPI_INT\n");
44         else if (datatype == MPI_DOUBLE) printf("MPI_DOUBLE\n");
45         ... else test for other types ...
46         return 0;
47         break;
48

```

```

1  case MPI_COMBINER_STRUCT:
2      printf("Datatype is struct containing");
3      array_of_ints = (int *)malloc(num_ints * sizeof(int));
4      array_of_adds =
5          (MPI_Aint *) malloc(num_adds * sizeof(MPI_Aint));
6      array_of_dtypes = (MPI_Datatype *)
7          malloc(num_dtypes * sizeof(MPI_Datatype));
8      MPI_Type_get_contents(datatype, num_ints, num_adds, num_dtypes,
9          array_of_ints, array_of_adds, array_of_dtypes);
10     printf(" %d datatypes:\n", array_of_ints[0]);
11     for (i=0; i<array_of_ints[0]; i++) {
12         printf("blocklength %d, displacement %ld, type:\n",
13             array_of_ints[i+1], (long)array_of_adds[i]);
14         if (printdatatype(array_of_dtypes[i])) {
15             /* Note that we free the type ONLY if it
16              is not predefined */
17             MPI_Type_free(&array_of_dtypes[i]);
18         }
19     }
20     free(array_of_ints);
21     free(array_of_adds);
22     free(array_of_dtypes);
23     break;
24     ... other combiner values ...
25     default:
26         printf("Unrecognized combiner type\n");
27     }
28     return 1;
29 }

```

5.2 Pack and Unpack

Some existing communication libraries provide pack/unpack procedures for sending noncontiguous data. In these, the user explicitly packs data into a contiguous buffer before sending it, and unpacks it from a contiguous buffer after receiving it. Derived datatypes, which are described in Section 5.1, allow one, in most cases, to avoid explicit packing and unpacking. The user specifies the layout of the data to be sent or received, and the communication library directly accesses a noncontiguous buffer. The pack/unpack routines are provided for compatibility with previous libraries. Also, they provide some functionality that is not otherwise available in MPI. For instance, a message can be received in several parts, where the receive operation done on a later part may depend on the content of a former part. Another use is that outgoing messages may be explicitly buffered in user supplied space, thus overriding the system buffering policy. Finally, the availability of pack and unpack operations facilitates the development of additional communication libraries layered on top of MPI.

MPI_PACK(inbuf, incount, datatype, outbuf, outsize, position, comm)

IN inbuf input buffer start (choice)

IN	incount	number of input data items (non-negative integer)	1
IN	datatype	datatype of each input data item (handle)	2
OUT	outbuf	output buffer start (choice)	3
IN	outside	output buffer size, in bytes (non-negative integer)	4
INOUT	position	current position in buffer, in bytes (integer)	5
IN	comm	communicator for packed message (handle)	6

C binding

```

int MPI_Pack(const void *inbuf, int incount, MPI_Datatype datatype,
            void *outbuf, int outside, int *position, MPI_Comm comm)
int MPI_Pack_c(const void *inbuf, MPI_Count incount, MPI_Datatype datatype,
              void *outbuf, MPI_Count outside, MPI_Count *position,
              MPI_Comm comm)

```

Fortran 2008 binding

```

MPI_Pack(inbuf, incount, datatype, outbuf, outside, position, comm, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
  INTEGER, INTENT(IN) :: incount, outside
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(*), DIMENSION(..) :: outbuf
  INTEGER, INTENT(INOUT) :: position
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Pack(inbuf, incount, datatype, outbuf, outside, position, comm, ierror)
  !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incount, outside
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(*), DIMENSION(..) :: outbuf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_PACK(INBUF, INCOUNT, DATATYPE, OUTBUF, OUTSIZE, POSITION, COMM, IERROR)
  <type> INBUF(*), OUTBUF(*)
  INTEGER INCOUNT, DATATYPE, OUTSIZE, POSITION, COMM, IERROR

```

Packs the message in the send buffer specified by `inbuf`, `incount`, `datatype` into the buffer space specified by `outbuf` and `outside`. The input buffer can be any communication buffer allowed in `MPI_SEND`. The output buffer is a contiguous storage area containing `outside` bytes, starting at the address `outbuf` (length is counted in *bytes*, not elements, as if it were a communication buffer for a message of type `MPI_PACKED`).

The input value of `position` is the first location in the output buffer to be used for packing. `position` is incremented by the size of the packed message, and the output value of `position` is the first location in the output buffer following the locations occupied by the

1 packed message. The `comm` argument is the communicator that will be subsequently used
 2 for sending the packed message.

3
 4
 5 `MPI_UNPACK(inbuf, insize, position, outbuf, outcount, datatype, comm)`

6	IN	inbuf	input buffer start (choice)
7	IN	insize	size of input buffer, in bytes (non-negative integer)
8	INOUT	position	current position in bytes (integer)
9	OUT	outbuf	output buffer start (choice)
10	IN	outcount	number of items to be unpacked (integer)
11	IN	datatype	datatype of each output data item (handle)
12	IN	comm	communicator for packed message (handle)

16 C binding

17 `int MPI_Unpack(const void *inbuf, int insize, int *position, void *outbuf,`
 18 `int outcount, MPI_Datatype datatype, MPI_Comm comm)`

19
 20 `int MPI_Unpack_c(const void *inbuf, MPI_Count insize, MPI_Count *position,`
 21 `void *outbuf, MPI_Count outcount, MPI_Datatype datatype,`
 22 `MPI_Comm comm)`

23 Fortran 2008 binding

24 `MPI_Unpack(inbuf, insize, position, outbuf, outcount, datatype, comm, ierror)`
 25 `TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf`
 26 `INTEGER, INTENT(IN) :: insize, outcount`
 27 `INTEGER, INTENT(INOUT) :: position`
 28 `TYPE(*), DIMENSION(..) :: outbuf`
 29 `TYPE(MPI_Datatype), INTENT(IN) :: datatype`
 30 `TYPE(MPI_Comm), INTENT(IN) :: comm`
 31 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

32
 33 `MPI_Unpack(inbuf, insize, position, outbuf, outcount, datatype, comm, ierror)`
 34 `!(_c)`

35 `TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf`
 36 `INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: insize, outcount`
 37 `INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position`
 38 `TYPE(*), DIMENSION(..) :: outbuf`
 39 `TYPE(MPI_Datatype), INTENT(IN) :: datatype`
 40 `TYPE(MPI_Comm), INTENT(IN) :: comm`
 41 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

42 Fortran binding

43 `MPI_UNPACK(INBUF, INSIZE, POSITION, OUTBUF, OUTCOUNT, DATATYPE, COMM, IERROR)`
 44 `<type> INBUF(*), OUTBUF(*)`
 45 `INTEGER INSIZE, POSITION, OUTCOUNT, DATATYPE, COMM, IERROR`

46 Unpacks a message into the receive buffer specified by `outbuf`, `outcount`, `datatype` from
 47 the buffer space specified by `inbuf` and `insize`. The output buffer can be any communication
 48

buffer allowed in `MPI_RECV`. The input buffer is a contiguous storage area containing `insize` bytes, starting at address `inbuf`. The input value of `position` is the first location in the input buffer occupied by the packed message. `position` is incremented by the size of the packed message, so that the output value of `position` is the first location in the input buffer after the locations occupied by the message that was unpacked. `comm` is the communicator used to receive the packed message.

Advice to users. Note the difference between `MPI_RECV` and `MPI_UNPACK`: in `MPI_RECV`, the `count` argument specifies the maximum number of items that can be received. The actual number of items received is determined by the length of the incoming message. In `MPI_UNPACK`, the `count` argument specifies the actual number of items that are unpacked; the “size” of the corresponding message is the increment in `position`. The reason for this change is that the “incoming message size” is not predetermined since the user decides how much to unpack; nor is it easy to determine the “message size” from the number of items to be unpacked. In fact, in a heterogeneous system, this number may not be determined *a priori*. (*End of advice to users.*)

To understand the behavior of pack and unpack, it is convenient to think of the data part of a message as being the sequence obtained by concatenating the successive values sent in that message. The pack operation stores this sequence in the buffer space, as if sending the message to that buffer. The unpack operation retrieves this sequence from buffer space, as if receiving a message from that buffer. (It is helpful to think of internal Fortran files or `sscanf` in C, for a similar function.)

Several messages can be successively packed into one **packing unit**. This is effected by several successive **related** calls to `MPI_PACK`, where the first call provides `position = 0`, and each successive call inputs the value of `position` that was output by the previous call, and the same values for `outbuf`, `outcount` and `comm`. This packing unit now contains the equivalent information that would have been stored in a message by one send call with a send buffer that is the “concatenation” of the individual send buffers.

A packing unit can be sent using type `MPI_PACKED`. Any point-to-point or collective communication operation can be used to move the sequence of bytes that forms the packing unit from one process to another. This packing unit can now be received using any receive operation, with any datatype: the type matching rules are relaxed for messages sent with type `MPI_PACKED`.

A message sent with any type (including `MPI_PACKED`) can be received using the type `MPI_PACKED`. Such a message can then be unpacked by calls to `MPI_UNPACK`.

A packing unit (or a message created by a regular, “typed” send) can be unpacked into several successive messages. This is effected by several successive related calls to `MPI_UNPACK`, where the first call provides `position = 0`, and each successive call inputs the value of `position` that was output by the previous call, and the same values for `inbuf`, `insize` and `comm`.

The concatenation of two packing units is not necessarily a packing unit; nor is a substring of a packing unit necessarily a packing unit. Thus, one cannot concatenate two packing units and then unpack the result as one packing unit; nor can one unpack a substring of a packing unit as a separate packing unit. Each packing unit, that was created by a related sequence of pack calls, or by a regular send, must be unpacked as a unit, by a sequence of related unpack calls.

Rationale. The restriction on “atomic” packing and unpacking of packing units allows the implementation to add at the head of packing units additional information, such as a description of the sender architecture (to be used for type conversion, in a heterogeneous environment) (*End of rationale.*)

The following call allows the user to find out how much space is needed to pack a message and, thus, manage space allocation for buffers.

MPI_PACK_SIZE(incount, datatype, comm, size)

IN	incount	count argument to packing call (non-negative integer)
IN	datatype	datatype argument to packing call (handle)
IN	comm	communicator argument to packing call (handle)
OUT	size	upper bound on size of packed message, in bytes (non-negative integer)

C binding

```
int MPI_Pack_size(int incount, MPI_Datatype datatype, MPI_Comm comm, int *size)
```

```
int MPI_Pack_size_c(MPI_Count incount, MPI_Datatype datatype, MPI_Comm comm,
                    MPI_Count *size)
```

Fortran 2008 binding

```
MPI_Pack_size(incount, datatype, comm, size, ierror)
  INTEGER, INTENT(IN) :: incount
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(OUT) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Pack_size(incount, datatype, comm, size, ierror) !(_c)
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incount
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_PACK_SIZE(INCOUNT, DATATYPE, COMM, SIZE, IERROR)
  INTEGER INCOUNT, DATATYPE, COMM, SIZE, IERROR
```

A call to **MPI_PACK_SIZE**(incount, datatype, comm, size) returns in size an upper bound on the increment in position that is effected by a call to **MPI_PACK**(inbuf, incount, datatype, outbuf, outcount, position, comm). If the packed size of the datatype cannot be expressed by the size parameter, then **MPI_PACK_SIZE** sets the value of size to **MPI_UNDEFINED**.

Rationale. The call returns an upper bound, rather than an exact bound, since the exact amount of space needed to pack the message may depend on the context (e.g., first message packed in a packing unit may take more space). (*End of rationale.*)

Example 5.21. An example using MPI_PACK.

```

1
2
3  int      position, i, j, a[2];
4  char      buff[1000];
5
6  MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
7  if (myrank == 0)
8  {
9      /* SENDER CODE */
10     position = 0;
11     MPI_Pack(&i, 1, MPI_INT, buff, 1000, &position, MPI_COMM_WORLD);
12     MPI_Pack(&j, 1, MPI_INT, buff, 1000, &position, MPI_COMM_WORLD);
13     MPI_Send(buff, position, MPI_PACKED, 1, 0, MPI_COMM_WORLD);
14 }
15 else /* RECEIVER CODE */
16     MPI_Recv(a, 2, MPI_INT, 0, 0, MPI_COMM_WORLD, MPI_STATUS_IGNORE);

```

Example 5.22. An elaborate example.

```

17
18
19  int      position, i = 200;
20  float    a[200];
21  char      buff[1000]; /* larger than or equal to the size returned
22                        from MPI_PACK_SIZE for 1,newtype */
23  MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
24  if (myrank == 0)
25  {
26      /* SENDER CODE */
27      int len[2];
28      MPI_Aint disp[2];
29      MPI_Datatype type[2], newtype;
30
31      /* build datatype for i followed by a[0]...a[i-1] */
32      len[0] = 1;
33      len[1] = i;
34      MPI_Get_address(&i, disp);
35      MPI_Get_address(a, disp+1);
36      type[0] = MPI_INT;
37      type[1] = MPI_FLOAT;
38      MPI_Type_create_struct(2, len, disp, type, &newtype);
39      MPI_Type_commit(&newtype);
40
41      /* Pack i followed by a[0]...a[i-1]*/
42      position = 0;
43      MPI_Pack(MPI_BOTTOM, 1, newtype, buff, 1000, &position,
44              MPI_COMM_WORLD);
45
46      /* Send */
47      MPI_Send(buff, position, MPI_PACKED, 1, 0,
48              MPI_COMM_WORLD);
49
50  /* *****
51     One can replace the last three lines with

```

```

1  MPI_Send(MPI_BOTTOM, 1, newtype, 1, 0, MPI_COMM_WORLD);
2  ***** */
3  }
4  else if (myrank == 1)
5  {
6      /* RECEIVER CODE */
7      MPI_Status status;
8
9      /* Receive */
10     MPI_Recv(buff, 1000, MPI_PACKED, 0, 0, MPI_COMM_WORLD, &status);
11
12     /* Unpack i */
13     position = 0;
14     MPI_Unpack(buff, 1000, &position, &i, 1, MPI_INT, MPI_COMM_WORLD);
15
16     /* Unpack a[0]...a[i-1] */
17     MPI_Unpack(buff, 1000, &position, a, i, MPI_FLOAT, MPI_COMM_WORLD);
18 }

```

Example 5.23. Each process sends a count, followed by count characters to the root; the root concatenates all characters into one string.

```

19
20
21
22 int count, gsize, counts[64], totalcount, k1, k2, k,
23     displs[64], position, concat_pos;
24 char chr[100], *lbuf, *rbuf, *cbuf;
25
26 MPI_Comm_size(comm, &gsize);
27 MPI_Comm_rank(comm, &myrank);
28
29 /* allocate local pack buffer */
30 MPI_Pack_size(1, MPI_INT, comm, &k1);
31 MPI_Pack_size(count, MPI_CHAR, comm, &k2);
32 k = k1+k2;
33 lbuf = (char *)malloc(k);
34
35 /* pack count, followed by count characters */
36 position = 0;
37 MPI_Pack(&count, 1, MPI_INT, lbuf, k, &position, comm);
38 MPI_Pack(chr, count, MPI_CHAR, lbuf, k, &position, comm);
39
40 if (myrank != root) {
41     /* gather at root sizes of all packed messages */
42     MPI_Gather(&position, 1, MPI_INT, NULL, 0,
43              MPI_DATATYPE_NULL, root, comm);
44
45     /* gather at root packed messages */
46     MPI_Gatherv(lbuf, position, MPI_PACKED, NULL,
47                NULL, NULL, MPI_DATATYPE_NULL, root, comm);
48 } else { /* root code */
49     /* gather sizes of all packed messages */
50     MPI_Gather(&position, 1, MPI_INT, counts, 1,

```

```

        MPI_INT, root, comm);

/* gather all packed messages */
displs[0] = 0;
for (i=1; i < gsize; i++)
    displs[i] = displs[i-1] + counts[i-1];
totalcount = displs[gsize-1] + counts[gsize-1];
rbuf = (char *)malloc(totalcount);
cbuf = (char *)malloc(totalcount);
MPI_Gatherv(lbuf, position, MPI_PACKED, rbuf,
            counts, displs, MPI_PACKED, root, comm);

/* unpack all messages and concatenate strings */
concat_pos = 0;
for (i=0; i < gsize; i++) {
    position = 0;
    MPI_Unpack(rbuf+displs[i], totalcount-displs[i],
               &position, &count, 1, MPI_INT, comm);
    MPI_Unpack(rbuf+displs[i], totalcount-displs[i],
               &position, cbuf+concat_pos, count, MPI_CHAR, comm);
    concat_pos += count;
}
cbuf[concat_pos] = '\0';
}

```

5.3 Canonical MPI_PACK and MPI_UNPACK

These procedures read/write data to/from the buffer in the "external32" data format specified in Section 14.5.2, and calculate the size needed for packing. Their first arguments specify the data format, for future extensibility, but currently the only valid value of the `datarep` argument is "external32".

Advice to users. These procedures could be used, for example, to send typed data in a portable format from one MPI implementation to another. (*End of advice to users.*)

The buffer will contain exactly the packed data, without headers. `MPI_BYTE` should be used to send and receive data that is packed using `MPI_PACK_EXTERNAL`.

Rationale. `MPI_PACK_EXTERNAL` specifies that there is no header on the message and further specifies the exact format of the data. Since `MPI_PACK` may (and is allowed to) use a header, the datatype `MPI_PACKED` cannot be used for data packed with `MPI_PACK_EXTERNAL`. (*End of rationale.*)

`MPI_PACK_EXTERNAL(datarep, inbuf, incount, datatype, outbuf, outsize, position)`

IN	<code>datarep</code>	data representation (string)
IN	<code>inbuf</code>	input buffer start (choice)
IN	<code>incount</code>	number of input data items (integer)
IN	<code>datatype</code>	datatype of each input data item (handle)

```

1      OUT      outbuf          output buffer start (choice)
2      IN       outsize        output buffer size, in bytes (integer)
3
4      INOUT    position       current position in buffer, in bytes (integer)
5

```

C binding

```

7      int MPI_Pack_external(const char datarep[], const void *inbuf, int incout,
8                          MPI_Datatype datatype, void *outbuf, MPI_Aint outsize,
9                          MPI_Aint *position)
10
11     int MPI_Pack_external_c(const char datarep[], const void *inbuf,
12                           MPI_Count incout, MPI_Datatype datatype, void *outbuf,
13                           MPI_Count outsize, MPI_Count *position)
14

```

Fortran 2008 binding

```

15     MPI_Pack_external(datarep, inbuf, incout, datatype, outbuf, outsize, position,
16                      ierror)
17     CHARACTER(LEN=*), INTENT(IN) :: datarep
18     TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
19     INTEGER, INTENT(IN) :: incout
20     TYPE(MPI_Datatype), INTENT(IN) :: datatype
21     TYPE(*), DIMENSION(..) :: outbuf
22     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: outsize
23     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(INOUT) :: position
24     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
25
26     MPI_Pack_external(datarep, inbuf, incout, datatype, outbuf, outsize, position,
27                      ierror) !(_c)
28     CHARACTER(LEN=*), INTENT(IN) :: datarep
29     TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
30     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incout, outsize
31     TYPE(MPI_Datatype), INTENT(IN) :: datatype
32     TYPE(*), DIMENSION(..) :: outbuf
33     INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position
34     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
35

```

Fortran binding

```

36     MPI_PACK_EXTERNAL(DATAREP, INBUF, INCOUNT, DATATYPE, OUTBUF, OUTSIZE, POSITION,
37                      IERROR)
38     CHARACTER*(*) DATAREP
39     <type> INBUF(*), OUTBUF(*)
40     INTEGER INCOUNT, DATATYPE, IERROR
41     INTEGER(KIND=MPI_ADDRESS_KIND) OUTSIZE, POSITION
42
43
44
45
46
47
48

```

MPI_UNPACK_EXTERNAL(datarep, inbuf, insize, position, outbuf, outcount, datatype)	1
IN datarep	2 data representation (string) 3
IN inbuf	4 input buffer start (choice)
IN insize	5 input buffer size, in bytes (integer)
INOUT position	6 current position in buffer, in bytes (integer) 7
OUT outbuf	8 output buffer start (choice)
IN outcount	9 number of output data items (integer)
IN datatype	10 datatype of output data item (handle) 11

C binding

```
int MPI_Unpack_external(const char datarep[], const void *inbuf,
    MPI_Aint insize, MPI_Aint *position, void *outbuf, int outcount,
    MPI_Datatype datatype)
```

```
int MPI_Unpack_external_c(const char datarep[], const void *inbuf,
    MPI_Count insize, MPI_Count *position, void *outbuf,
    MPI_Count outcount, MPI_Datatype datatype)
```

Fortran 2008 binding

```
MPI_Unpack_external(datarep, inbuf, insize, position, outbuf, outcount,
    datatype, ierror)
```

```
CHARACTER(LEN=*), INTENT(IN) :: datarep
TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: insize
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(INOUT) :: position
TYPE(*), DIMENSION(..) :: outbuf
INTEGER, INTENT(IN) :: outcount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Unpack_external(datarep, inbuf, insize, position, outbuf, outcount,
    datatype, ierror) !(_c)
```

```
CHARACTER(LEN=*), INTENT(IN) :: datarep
TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: insize, outcount
INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position
TYPE(*), DIMENSION(..) :: outbuf
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_UNPACK_EXTERNAL(DATAREP, INBUF, INSIZE, POSITION, OUTBUF, OUTCOUNT,
    DATATYPE, IERROR)
```

```
CHARACTER*(*) DATAREP
<type> INBUF(*), OUTBUF(*)
INTEGER(KIND=MPI_ADDRESS_KIND) INSIZE, POSITION
INTEGER OUTCOUNT, DATATYPE, IERROR
```

```

1 MPI_PACK_EXTERNAL_SIZE(datarep, incount, datatype, size)
2   IN      datarep          data representation (string)
3
4   IN      incount         number of input data items (integer)
5
6   IN      datatype        datatype of each input data item (handle)
7
8   OUT     size            output buffer size, in bytes (integer)

```

C binding

```

9
10 int MPI_Pack_external_size(const char datarep[], int incount,
11                          MPI_Datatype datatype, MPI_Aint *size)
12
13 int MPI_Pack_external_size_c(const char datarep[], MPI_Count incount,
14                             MPI_Datatype datatype, MPI_Count *size)

```

Fortran 2008 binding

```

15 MPI_PACK_EXTERNAL_SIZE(datarep, incount, datatype, size, ierror)
16   CHARACTER(LEN=*), INTENT(IN) :: datarep
17   INTEGER, INTENT(IN) :: incount
18   TYPE(MPI_Datatype), INTENT(IN) :: datatype
19   INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: size
20   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_PACK_EXTERNAL_SIZE(datarep, incount, datatype, size, ierror) !(_c)
23   CHARACTER(LEN=*), INTENT(IN) :: datarep
24   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incount
25   TYPE(MPI_Datatype), INTENT(IN) :: datatype
26   INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
27   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

28
29 MPI_PACK_EXTERNAL_SIZE(DATAREP, INCOUNT, DATATYPE, SIZE, IERROR)
30   CHARACTER*(*) DATAREP
31   INTEGER INCOUNT, DATATYPE, IERROR
32   INTEGER(KIND=MPI_ADDRESS_KIND) SIZE
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

Chapter 6

Collective Communication

6.1 Introduction and Overview

Collective communication is defined as communication that involves a group or groups of MPI processes. The functions of this type provided by MPI are the following:

- `MPI_BARRIER`, `MPI_IBARRIER`, `MPI_BARRIER_INIT`: Barrier synchronization across all members of a group (Section 6.3, Section 6.12.1, and Section 6.13.1).
- `MPI_BCAST`, `MPI_IBCAST`, `MPI_BCAST_INIT`: Broadcast from one member to all members of a group (Section 6.4, Section 6.12.2, and Section 6.13.2). This is shown as “broadcast” in Figure 6.1.
- `MPI_GATHER`, `MPI_IGATHER`, `MPI_GATHER_INIT`, `MPI_GATHERV`, `MPI_IGATHERV`, `MPI_GATHERV_INIT`, : Gather data from all members of a group to one member (Section 6.5, Section 6.12.3, and Section 6.13.3). This is shown as “gather” in Figure 6.1.
- `MPI_SCATTER`, `MPI_ISCATTER`, `MPI_SCATTER_INIT`, `MPI_SCATTERV`, `MPI_ISCATTERV`, `MPI_SCATTERV_INIT`: Scatter data from one member to all members of a group (Section 6.6, Section 6.12.4, and Section 6.13.4). This is shown as “scatter” in Figure 6.1.
- `MPI_ALLGATHER`, `MPI_IALLGATHER`, `MPI_ALLGATHER_INIT`, `MPI_ALLGATHERV`, `MPI_IALLGATHERV`, `MPI_ALLGATHERV_INIT`: A variation on Gather where all members of a group receive the result (Section 6.7, Section 6.12.5, and Section 6.13.5). This is shown as “allgather” in Figure 6.1.
- `MPI_ALLTOALL`, `MPI_IALLTOALL`, `MPI_ALLTOALL_INIT`, `MPI_ALLTOALLV`, `MPI_IALLTOALLV`, `MPI_ALLTOALLV_INIT`, `MPI_ALLTOALLW`, `MPI_IALLTOALLW`, `MPI_ALLTOALLW_INIT`: Scatter/Gather data from all members to all members of a group (also called complete exchange) (Section 6.8, Section 6.12.6, and Section 6.13.6). This is shown as “complete exchange” in Figure 6.1.
- `MPI_ALLREDUCE`, `MPI_IALLREDUCE`, `MPI_ALLREDUCE_INIT`, `MPI_REDUCE`, `MPI_IREDUCE`, `MPI_REDUCE_INIT`: Global reduction operations such as sum, max, min, or user-defined functions, where the result is returned to all members of a group (Section 6.9.6, Section 6.12.8, and Section 6.13.8) and a variation where the result is returned to only one member (Section 6.9, Section 6.12.7, and Section 6.13.7).
- `MPI_REDUCE_SCATTER_BLOCK`, `MPI_IREDUCE_SCATTER_BLOCK`, `MPI_REDUCE_SCATTER_BLOCK_INIT`, `MPI_REDUCE_SCATTER`,

1 MPI_IREDUCE_SCATTER, MPI_REDUCE_SCATTER_INIT: A combined reduction and
2 scatter operation (Section 6.10, Section 6.12.9, Section 6.12.10, Section 6.13.9, and
3 Section 6.13.10).

- 4 • MPI_SCAN, MPI_ISCAN, MPI_SCAN_INIT, MPI_EXSCAN, MPI_IEXSCAN,
5 MPI_EXSCAN_INIT: Scan across all members of a group (also called prefix) (Sec-
6 tion 6.11, Section 6.11.2, Section 6.12.11, Section 6.12.12, Section 6.13.11, and Sec-
7 tion 6.13.12).

9 One of the key arguments in a call to a collective routine is a communicator that
10 defines the group or groups of participating MPI processes and provides a context for the
11 operation. This is discussed further in Section 6.2. The syntax and semantics of the
12 collective operations are defined to be consistent with the syntax and semantics of the
13 point-to-point operations. Thus, general datatypes are allowed and must match between
14 sending and receiving MPI processes as specified in Chapter 5. Several collective routines
15 such as broadcast and gather have a single originating or receiving MPI process. Such an
16 MPI process is called the **root**. Some arguments in the collective functions are specified as
17 “significant only at root,” and are ignored for all participants except the root. The reader is
18 referred to Chapter 5 for information concerning communication buffers, general datatypes
19 and type matching rules, and to Chapter 7 for information on how to define groups and
20 create communicators.

21 The type-matching conditions for the collective operations are more strict than the cor-
22 responding conditions between sender and receiver in point-to-point. Namely, for collective
23 operations, the amount of data sent must exactly match the amount of data specified by
24 the receiver. Different type maps (the layout in memory, see Section 5.1) between sender
25 and receiver are still allowed.

26 Collective operations can (but are not required to) complete as soon as the caller’s
27 participation in the collective communication is finished. A blocking operation is complete
28 as soon as the call returns. A nonblocking (immediate) call requires a separate completion
29 call (cf. Section 3.7). The completion of a collective operation indicates that the caller
30 is free to modify locations in the communication buffer. It does not indicate that other
31 MPI processes in the group have completed or even started the operation (unless otherwise
32 implied by the description of the operation). Thus, a collective communication operation
33 may, or may not, have the effect of synchronizing all participating MPI processes.

34 Collective communication calls may use the same communicators as point-to-point
35 communication; MPI guarantees that messages generated on behalf of collective communi-
36 cation calls will not be confused with messages generated by point-to-point communication.
37 The collective operations do not have a message tag argument. A more detailed discussion
38 of correct use of collective routines is found in Section 6.14.

39
40 *Rationale.* The equal-data restriction (on type matching) was made so as to avoid
41 the complexity of providing a facility analogous to the status argument of MPI_RECV
42 for discovering the amount of data sent. Some of the collective routines would require
43 an array of status values.

44 The statements about synchronization are made so as to allow a variety of implemen-
45 tations of the collective functions.

46
47 (*End of rationale.*)




Figure 6.1: Collective move functions illustrated for a group of six MPI processes. In each case, each row of boxes represents data locations in one MPI process. Thus, in the broadcast, initially just the first MPI process contains the data A_0 , but after the broadcast all MPI processes contain it.

1 *Advice to users.* It is dangerous to rely on synchronization side-effects of the col-
2 lective operations for program correctness. For example, even though a particular
3 implementation may provide a broadcast routine with a side-effect of synchroniza-
4 tion, the standard does not require this, and a program that relies on this will not be
5 portable.

6 On the other hand, a correct, portable program must allow for the fact that a collective
7 call *may* be synchronizing. Though one cannot rely on any synchronization side-effect,
8 one must program so as to allow it. These issues are discussed further in Section 6.14.
9 (*End of advice to users.*)

11 *Advice to implementors.* While vendors may write optimized collective routines
12 matched to their architectures, a complete library of the collective communication
13 routines can be written entirely using the MPI point-to-point communication func-
14 tions and a few auxiliary functions. If implementing on top of point-to-point, a hidden,
15 special communicator might be created for the collective operation so as to avoid inter-
16 ference with any on-going point-to-point communication at the time of the collective
17 call. This is discussed further in Section 6.14. (*End of advice to implementors.*)

19 Many of the descriptions of the collective routines provide illustrations in terms of
20 blocking MPI point-to-point routines. These are intended solely to indicate what data
21 is sent or received by which MPI process. Many of these examples are *not* correct MPI
22 programs; for purposes of simplicity, they often assume infinite buffering.

24 6.2 Communicator Argument

26 The key concept of the collective functions is to have a group or groups of participating
27 MPI processes. The routines do not have group identifiers as explicit arguments. Instead,
28 there is a communicator argument. Groups and communicators are discussed in full detail
29 in Chapter 7. For the purposes of this chapter, it is sufficient to know that there are two
30 types of communicators: **intra-communicators** and **inter-communicators**. An intra-
31 communicator can be thought of as an identifier for a single group of MPI processes linked
32 with a context. An inter-communicator identifies two distinct groups of MPI processes
33 linked with a context.

35 6.2.1 Specifics for Intra-Communicator Collective Operations

37 All MPI processes in the group identified by the intra-communicator must call the collective
38 routine.

39 In many cases, collective communication can occur “in place” for intra-communicators,
40 with the output buffer being identical to the input buffer. This is specified by providing
41 a special argument value, `MPI_IN_PLACE`, instead of the send buffer or the receive buffer
42 argument, depending on the operation performed.

44 *Rationale.* The “in place” operations are provided to reduce unnecessary memory
45 motion by both the MPI implementation and by the user. Note that while the simple
46 check of testing whether the send and receive buffers have the same address will
47 work for some cases (e.g., `MPI_ALLREDUCE`), they are inadequate in others (e.g.,
48 `MPI_GATHER`, with `root` not equal to zero). Further, Fortran explicitly prohibits

aliasing of arguments; the approach of using a special value to denote “in place” operation eliminates that difficulty. (*End of rationale.*)

Advice to users. By allowing the “in place” option, the receive buffer in many of the collective calls becomes a send-and-receive buffer. For this reason, a Fortran binding that includes INTENT must mark these as INOUT, not OUT.

Note that MPI_IN_PLACE is a special kind of value; it has the same restrictions on its use that MPI_BOTTOM has (not usable in Fortran for initialization or assignment). See Section 2.5.4. (*End of advice to users.*)

6.2.2 Applying Collective Operations to Inter-Communicators

To understand how collective operations apply to inter-communicators, we can view most MPI intra-communicator collective operations as fitting one of the following categories (see, for instance, [64]):

All-To-All All MPI processes contribute to the result. All MPI processes receive the result.

- MPI_ALLGATHER, MPI_IALLGATHER, MPI_ALLGATHER_INIT, MPI_ALLGATHERV, MPI_IALLGATHERV, MPI_ALLGATHERV_INIT
- MPI_ALLTOALL, MPI_IALLTOALL, MPI_ALLTOALL_INIT, MPI_ALLTOALLV, MPI_IALLTOALLV, MPI_ALLTOALLV_INIT, MPI_ALLTOALLW, MPI_IALLTOALLW, MPI_ALLTOALLW_INIT
- MPI_ALLREDUCE, MPI_IALLREDUCE, MPI_ALLREDUCE_INIT, MPI_REDUCE_SCATTER_BLOCK, MPI_IREDUCE_SCATTER_BLOCK, MPI_REDUCE_SCATTER_BLOCK_INIT, MPI_REDUCE_SCATTER, MPI_IREDUCE_SCATTER, MPI_REDUCE_SCATTER_INIT
- MPI_BARRIER, MPI_IBARRIER, MPI_BARRIER_INIT

All-To-One All MPI processes contribute to the result. One MPI process receives the result.

- MPI_GATHER, MPI_IGATHER, MPI_GATHER_INIT, MPI_GATHERV, MPI_IGATHERV, MPI_GATHERV_INIT
- MPI_REDUCE, MPI_IREDUCE, MPI_REDUCE_INIT,

One-To-All One MPI process contributes to the result. All MPI processes receive the result.

- MPI_BCAST, MPI_IBCAST, MPI_BCAST_INIT
- MPI_SCATTER, MPI_ISCATTER, MPI_SCATTER_INIT, MPI_SCATTERV, MPI_ISCATTERV, MPI_SCATTERV_INIT

Other: Collective operations that do not fit into one of the above categories.

- MPI_SCAN, MPI_ISCAN, MPI_SCAN_INIT MPI_EXSCAN, MPI_IEXSCAN, MPI_EXSCAN_INIT

The data movement patterns of MPI_SCAN, MPI_ISCAN, MPI_SCAN_INIT, MPI_EXSCAN, MPI_IEXSCAN and MPI_EXSCAN_INIT do not fit this taxonomy.

The application of collective communication to inter-communicators is best described in terms of two groups. For example, an all-to-all `MPI_ALLGATHER` operation can be described as collecting data from all members of one group with the result appearing in all members of the other group (see Figure 6.2). As another example, a one-to-all `MPI_BCAST` operation sends data from one member of one group to all members of the other group. Collective computation operations such as `MPI_REDUCE_SCATTER` have a similar interpretation (see Figure 6.3). For intra-communicators, these two groups are the same. For inter-communicators, these two groups are distinct. For the all-to-all operations, each such operation is described in two phases, so that it has a symmetric, full-duplex behavior.

The following collective operations also apply to inter-communicators:

- `MPI_BARRIER`, `MPI_IBARRIER`, `MPI_BARRIER_INIT`,
- `MPI_BCAST`, `MPI_IBCAST`, `MPI_BCAST_INIT`,
- `MPI_GATHER`, `MPI_IGATHER`, `MPI_GATHER_INIT`, `MPI_GATHERV`,
`MPI_IGATHERV`, `MPI_GATHERV_INIT`,
- `MPI_SCATTER`, `MPI_ISCATTER`, `MPI_SCATTER_INIT`, `MPI_SCATTERV`,
`MPI_ISCATTERV`, `MPI_SCATTERV_INIT`,
- `MPI_ALLGATHER`, `MPI_IALLGATHER`, `MPI_ALLGATHER_INIT`, `MPI_ALLGATHERV`,
`MPI_IALLGATHERV`, `MPI_ALLGATHERV_INIT`,
- `MPI_ALLTOALL`, `MPI_IALLTOALL`, `MPI_ALLTOALL_INIT`, `MPI_ALLTOALLV`,
`MPI_IALLTOALLV`, `MPI_ALLTOALLV_INIT`, `MPI_ALLTOALLW`, `MPI_IALLTOALLW`,
`MPI_ALLTOALLW_INIT`,
- `MPI_ALLREDUCE`, `MPI_IALLREDUCE`, `MPI_ALLREDUCE_INIT`, `MPI_REDUCE`,
`MPI_IREDUCE`, `MPI_REDUCE_INIT`,
- `MPI_REDUCE_SCATTER_BLOCK`, `MPI_IREDUCE_SCATTER_BLOCK`,
`MPI_REDUCE_SCATTER_BLOCK_INIT`, `MPI_REDUCE_SCATTER`,
`MPI_IREDUCE_SCATTER`, `MPI_REDUCE_SCATTER_INIT`.

6.2.3 Specifics for Inter-Communicator Collective Operations

All MPI processes in both groups identified by the inter-communicator must call the collective routine.

Note that the “in place” option for intra-communicators does not apply to inter-communicators since in the inter-communicator case there is no communication from an MPI process to itself.

For inter-communicator collective communication, if the operation is in the All-To-One or One-To-All categories, then the transfer is unidirectional. The direction of the transfer is indicated by a special value of the `root` argument. In this case, for the group containing the root, all MPI processes in the group must call the routine using a special argument for the root. For this, the root uses the special value `MPI_ROOT`; all other MPI processes in the same group as the root use `MPI_PROC_NULL`. All MPI processes in the other group (the group that is the remote group relative to the root) must call the collective routine and




Figure 6.2: Inter-communicator allgather. The focus of data to one MPI process is represented, not mandated by the semantics. The two phases do allgathers in both directions.




Figure 6.3: Inter-communicator reduce-scatter. The focus of data to one MPI process is represented, not mandated by the semantics. The two phases do reduce-scatters in both directions.

1 provide the rank of the root. If the operation is in the All-To-All category, then the transfer
2 is bidirectional.

3
4 *Rationale.* Operations in the All-To-One and One-To-All categories are unidirectional
5 by nature, and there is a clear way of specifying direction. Operations in the All-To-All
6 category will often occur as part of an exchange, where it makes sense to communicate
7 in both directions at once. (*End of rationale.*)

9 6.3 Barrier Synchronization

12 MPI_BARRIER(comm)

14 IN comm communicator (handle)

16 C binding

17 int MPI_Barrier(MPI_Comm comm)

19 Fortran 2008 binding

20 MPI_Barrier(comm, ierror)

21 TYPE(MPI_Comm), INTENT(IN) :: comm

22 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

23 Fortran binding

24 MPI_BARRIER(COMM, IERROR)

25 INTEGER COMM, IERROR

26
27 If comm is an intra-communicator, MPI_BARRIER blocks the caller until all group
28 members have called it. The call returns at any MPI process only after all group members
29 have entered the call.

30
31 If comm is an inter-communicator, MPI_BARRIER involves two groups. The call returns
32 at MPI processes in one group (group A) of the inter-communicator only after all members
33 of the other group (group B) have entered the call (and vice versa). An MPI process may
34 return from the call before all MPI processes in its own group have entered the call.

35 6.4 Broadcast

38 MPI_BCAST(buffer, count, datatype, root, comm)

40 INOUT buffer starting address of buffer (choice)

41 IN count number of entries in buffer (non-negative integer)

42 IN datatype datatype of buffer (handle)

43 IN root rank of the root (integer)

44 IN comm communicator (handle)

47 C binding

48 int MPI_Bcast(void *buffer, int count, MPI_Datatype datatype, int root,

```

        MPI_Comm comm)
int MPI_Bcast_c(void *buffer, MPI_Count count, MPI_Datatype datatype, int root,
        MPI_Comm comm)

```

Fortran 2008 binding

```

MPI_Bcast(buffer, count, datatype, root, comm, ierror)
  TYPE(*), DIMENSION(..) :: buffer
  INTEGER, INTENT(IN) :: count, root
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Bcast(buffer, count, datatype, root, comm, ierror) !(_c)
  TYPE(*), DIMENSION(..) :: buffer
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER, INTENT(IN) :: root
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_BCAST(BUFFER, COUNT, DATATYPE, ROOT, COMM, IERROR)
  <type> BUFFER(*)
  INTEGER COUNT, DATATYPE, ROOT, COMM, IERROR

```

If `comm` is an intra-communicator, `MPI_BCAST` broadcasts a message from the MPI process with rank `root` to all MPI processes of the group, itself included. It is called by all members of the group using the same arguments for `comm` and `root`. On return, the content of the root's buffer is copied to all other MPI processes.

General, derived datatypes are allowed for `datatype`. The type signature of `count`, `datatype` on any MPI process must be equal to the type signature of `count`, `datatype` at the root. This implies that the amount of data sent must be equal to the amount received, pairwise between each MPI process and the root. `MPI_BCAST` and all other data-movement collective routines make this restriction. Distinct type maps between sender and receiver are still allowed.

The “in place” option is not meaningful here.

If `comm` is an inter-communicator, then the call involves all MPI processes in the inter-communicator, but with one group (group A) defining the root. All MPI processes in the other group (group B) pass the same value in argument `root`, which is the rank of the root in group A. The root passes the value `MPI_ROOT` in `root`. All other MPI processes in group A pass the value `MPI_PROC_NULL` in `root`. Data is broadcast from the root to all MPI processes in group B. The buffer arguments of the MPI processes in group B must be consistent with the buffer argument of the root.

6.4.1 Example using `MPI_BCAST`

The examples in this section use intra-communicators.

Example 6.1. Broadcast 100 ints from MPI process 0 to every MPI process in the group.

```

MPI_Comm comm;

```

```

1  int array[100];
2  int root=0;
3  ...
4  MPI_Bcast(array, 100, MPI_INT, root, comm);

```

As in many of our example code fragments, we assume that some of the variables (such as `comm` in the above) have been assigned appropriate values.

6.5 Gather

`MPI_GATHER(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root, comm)`

IN	sendbuf	starting address of send buffer (choice)
IN	sendcount	number of elements in send buffer (non-negative integer)
IN	sendtype	datatype of send buffer elements (handle)
OUT	recvbuf	address of receive buffer (choice, significant only at root)
IN	recvcount	number of elements for any single receive (non-negative integer, significant only at root)
IN	recvtype	datatype of recv buffer elements (handle, significant only at root)
IN	root	rank of receiving MPI process (integer)
IN	comm	communicator (handle)

C binding

```

30 int MPI_Gather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
31               void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
32               MPI_Comm comm)
33

```

```

34 int MPI_Gather_c(const void *sendbuf, MPI_Count sendcount,
35                 MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
36                 MPI_Datatype recvtype, int root, MPI_Comm comm)
37

```

Fortran 2008 binding

```

38 MPI_Gather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
39           comm, ierror)
40

```

```

41   TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
42   INTEGER, INTENT(IN) :: sendcount, recvcount, root
43   TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
44   TYPE(*), DIMENSION(..) :: recvbuf
45   TYPE(MPI_Comm), INTENT(IN) :: comm
46   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
47

```

```

47 MPI_Gather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
48           comm, ierror) !(_c)

```



```

TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_GATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, ROOT,
           COMM, IERROR)

```

```

<type> SENDBUF(*), RECVBUF(*)

```

```

INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, IERROR

```

If `comm` is an intra-communicator, each MPI process (the root included) sends the contents of its send buffer to the root. The root receives the messages and stores them in rank order. The outcome is *as if* each of the n MPI processes in the group (including the root) had executed a call to

```

MPI_Send(sendbuf, sendcount, sendtype, root , ...),

```

and the root had executed n calls to

```

MPI_Recv(recvbuf+i·recvcount·extent(recvtype), recvcount, recvtype, i,...),

```

where `extent(recvtype)` is the type extent obtained from a call to `MPI_Type_get_extent`.

An alternative description is that the n messages sent by the processes in the group are concatenated in rank order, and the resulting message is received by the root as if by a call to `MPI_RECV(recvbuf, recvcount·n, recvtype, ...)`.

The receive buffer is ignored for all nonroot MPI processes.

General, derived datatypes are allowed for both `sendtype` and `recvtype`. The type signature of `sendcount`, `sendtype` on each MPI process must be equal to the type signature of `recvcount`, `recvtype` at the root. This implies that the amount of data sent must be equal to the amount of data received, pairwise between each MPI process and the root. Distinct type maps between sender and receiver MPI processes are still allowed.

All arguments to the function are significant on the root, while on other MPI processes, only the arguments `sendbuf`, `sendcount`, `sendtype`, `root`, and `comm` are significant. The arguments `root` and `comm` must have identical values on all MPI processes.

The specification of counts and types should not cause any location on the root to be written more than once. Such a call is erroneous.

Note that the `recvcount` argument at the root indicates the number of items it receives from *each* MPI process, not the total number of items it receives.

The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` as the value of `sendbuf` at the root. In such a case, `sendcount` and `sendtype` are ignored, and the contribution of the root to the gathered vector is assumed to be already in the correct place in the receive buffer.

If `comm` is an inter-communicator, then the call involves all MPI processes in the inter-communicator, but with one group (group A) defining the root. All MPI processes in the other group (group B) pass the same value in argument `root`, which is the rank of the root in

group A. The root passes the value `MPI_ROOT` in `root`. All other MPI processes in group A pass the value `MPI_PROC_NULL` in `root`. Data is gathered from all MPI processes in group B to the root. The send buffer arguments of the MPI processes in group B must be consistent with the receive buffer argument of the root.

```

MPI_GATHERV(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs, recvtype, root,
            comm)
IN         sendbuf           starting address of send buffer (choice)
IN         sendcount         number of elements in send buffer (non-negative
                              integer)
IN         sendtype          datatype of send buffer elements (handle)
OUT        recvbuf           address of receive buffer (choice, significant only at
                              root)
IN         recvcnts          nonnegative integer array (of length group size)
                              containing the number of elements that are received
                              from each MPI process (significant only at root)
IN         displs            integer array (of length group size). Entry i specifies
                              the displacement relative to recvbuf at which to place
                              the incoming data from MPI process i (significant
                              only at root)
IN         recvtype          datatype of recv buffer elements (handle, significant
                              only at root)
IN         root              rank of receiving MPI process (integer)
IN         comm              communicator (handle)

```

C binding

```

int MPI_Gatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
               void *recvbuf, const int recvcnts[], const int displs[],
               MPI_Datatype recvtype, int root, MPI_Comm comm)
int MPI_Gatherv_c(const void *sendbuf, MPI_Count sendcount,
                 MPI_Datatype sendtype, void *recvbuf,
                 const MPI_Count recvcnts[], const MPI_Aint displs[],
                 MPI_Datatype recvtype, int root, MPI_Comm comm)

```

Fortran 2008 binding

```

MPI_Gatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
            recvtype, root, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcnts(*), displs(*), root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Gatherv(sendbuf, sendcount, sendtype, recvbuf, recvcoun1ts, displs,
            recvtype, root, comm, ierror) !(_c) 2
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf 3
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcoun4ts(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 5
TYPE(*), DIMENSION(..) :: recvbuf 6
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*) 7
INTEGER, INTENT(IN) :: root 8
TYPE(MPI_Comm), INTENT(IN) :: comm 9
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 10

```

Fortran binding

```

MPI_GATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
            RECVTYPE, ROOT, COMM, IERROR) 12
<type> SENDBUF(*), RECVBUF(*) 13
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT, 14
            COMM, IERROR 15

```

MPI_GATHERV extends the functionality of MPI_GATHER by allowing a varying count of data from each MPI process, since `recvcoun18ts is now an array. It also allows more flexibility as to where the data is placed on the root, by providing the new argument, displs. 19`

If `comm` is an intra-communicator, the outcome is *as if* each MPI process, including the root, sends a message to the root, ²⁰

```

MPI_Send(sendbuf, sendcount, sendtype, root, ...), 21

```

and the root executes `n` receives, ²²

```

MPI_Recv(recvbuf+displs[j]·extent(recvtype), recvcoun23ts[j], recvtype, i, ...) 24

```

The data received from MPI process `j` is placed into `recvbuf` of the root beginning at offset `displs[j]` elements (in terms of the `recvtype`). ²⁵

The receive buffer is ignored for all nonroot MPI processes. ²⁶

The type signature implied by `sendcount`, `sendtype` on MPI process `i` must be equal to the type signature implied by `recvcoun27ts[i], recvtype at the root. This implies that the amount of data sent must be equal to the amount of data received, pairwise between each MPI process and the root. Distinct type maps between sender and receiver are still allowed, as illustrated in Example 6.6. 28`

All arguments to the function are significant on the root, while on other MPI processes, only arguments `sendbuf`, `sendcount`, `sendtype`, `root`, and `comm` are significant. The arguments `root` and `comm` must have identical values on all MPI processes. ²⁹

The specification of counts, types, and displacements should not cause any location on the root to be written more than once. Such a call is erroneous. ³⁰

The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` as the value of `sendbuf` at the root. In such a case, `sendcount` and `sendtype` are ignored, and the contribution of the root to the gathered vector is assumed to be already in the correct place in the receive buffer. ³¹

If `comm` is an inter-communicator, then the call involves all MPI processes in the inter-communicator, but with one group (group A) defining the root. All MPI processes in the ³²




Figure 6.4: The root gathers 100 ints from each MPI process in the group.

other group (group B) pass the same value in argument `root`, which is the rank of the root in group A. The root passes the value `MPI_ROOT` in `root`. All other MPI processes in group A pass the value `MPI_PROC_NULL` in `root`. Data is gathered from all MPI processes in group B to the root. The send buffer arguments of the MPI processes in group B must be consistent with the receive buffer argument of the root.

6.5.1 Examples using `MPI_GATHER`, `MPI_GATHERV`

The examples in this section use intra-communicators.

Example 6.2. Gather 100 ints from every MPI process in group to the root. See Figure 6.4.

```

23 MPI_Comm comm;
24 int gsize, sendarray[100];
25 int root, *rbuf;
26 ...
27 MPI_Comm_size(comm, &gsiz);
28 rbuf = (int *)malloc(gsize*100*sizeof(int));
29 MPI_Gather(sendarray, 100, MPI_INT, rbuf, 100, MPI_INT, root, comm);

```

Example 6.3. Previous example modified—only the root allocates memory for the receive buffer.

```

33 MPI_Comm comm;
34 int gsize, sendarray[100];
35 int root, myrank, *rbuf;
36 ...
37 MPI_Comm_rank(comm, &myrank);
38 if (myrank == root) {
39     MPI_Comm_size(comm, &gsiz);
40     rbuf = (int *)malloc(gsize*100*sizeof(int));
41 }
42 MPI_Gather(sendarray, 100, MPI_INT, rbuf, 100, MPI_INT, root, comm);

```

Example 6.4. Do the same as the previous example, but use a derived datatype. Note that the type cannot be the entire set of `gsiz*100` ints since type matching is defined pairwise between the root and each MPI process in the gather.

```

47 MPI_Comm comm;

```




Figure 6.5: The root gathers 100 ints from each MPI process in the group, each set is placed `stride` ints apart.

```

int gsize, sendarray[100];
int root, *rbuf;
MPI_Datatype rtype;
...
MPI_Comm_size(comm, &gsize);
MPI_Type_contiguous(100, MPI_INT, &rtype);
MPI_Type_commit(&rtype);
rbuf = (int *)malloc(gsize*100*sizeof(int));
MPI_Gather(sendarray, 100, MPI_INT, rbuf, 1, rtype, root, comm);

```

Example 6.5. Now have each MPI process send 100 ints to the root, but place each set (of 100) `stride` ints apart at the receiving end. Use `MPI_GATHERV` and the `displs` argument to achieve this effect. Assume `stride` ≥ 100 . See Figure 6.5.

```

MPI_Comm comm;
int gsize, sendarray[100];
int root, *rbuf, stride;
int *displs, i, *rcounts;
...

MPI_Comm_size(comm, &gsize);
rbuf = (int *)malloc(gsize*stride*sizeof(int));
displs = (int *)malloc(gsize*sizeof(int));
rcounts = (int *)malloc(gsize*sizeof(int));
for (i=0; i<gsize; ++i) {
    displs[i] = i*stride;
    rcounts[i] = 100;
}
MPI_Gatherv(sendarray, 100, MPI_INT, rbuf, rcounts, displs, MPI_INT,
            root, comm);

```

Note that the program is erroneous if `stride` < 100 .

Example 6.6. Same as Example 6.5 on the receiving side, but send the 100 ints from the 0th column of a 100×150 int array, in C. See Figure 6.6.

```

MPI_Comm comm;

```




Figure 6.6: The root gathers column 0 of a 100×150 C array, and each set is placed `stride` ints apart.

```

14 int gsize, sendarray[100][150];
15 int root, *rbuf, stride;
16 MPI_Datatype stype;
17 int *displs, i, *rcounts;
18
19 ...
20
21 MPI_Comm_size(comm, &gsize);
22 rbuf = (int *)malloc(gsize*stride*sizeof(int));
23 displs = (int *)malloc(gsize*sizeof(int));
24 rcounts = (int *)malloc(gsize*sizeof(int));
25 for (i=0; i<gsize; ++i) {
26     displs[i] = i*stride;
27     rcounts[i] = 100;
28 }
29 /* Create datatype for 1 column of array
30 */
31 MPI_Type_vector(100, 1, 150, MPI_INT, &stype);
32 MPI_Type_commit(&stype);
33 MPI_Gatherv(sendarray, 1, stype, rbuf, rcounts, displs, MPI_INT,
34             root, comm);

```

Example 6.7. MPI process i sends $(100-i)$ ints from the i -th column of a 100×150 int array, in C. It is received into a buffer with stride, as in the previous two examples. See Figure 6.7.

```

38 MPI_Comm comm;
39 int gsize, sendarray[100][150], *sptr;
40 int root, *rbuf, stride, myrank;
41 MPI_Datatype stype;
42 int *displs, i, *rcounts;
43
44 ...
45
46 MPI_Comm_size(comm, &gsize);
47 MPI_Comm_rank(comm, &myrank);
48 rbuf = (int *)malloc(gsize*stride*sizeof(int));

```




Figure 6.7: The root gathers $100-i$ ints from column i of a 100×150 C array, and each set is placed $stride$ ints apart.

```

displs = (int *)malloc(gsize*sizeof(int));
rcounts = (int *)malloc(gsize*sizeof(int));
for (i=0; i<gsize; ++i) {
    displs[i] = i*stride;
    rcounts[i] = 100-i;    /* note change from previous example */
}
/* Create datatype for the column we are sending
*/
MPI_Type_vector(100-myrank, 1, 150, MPI_INT, &stype);
MPI_Type_commit(&stype);
/* sptr is the address of start of "myrank" column
*/
sptr = &sendarray[0][myrank];
MPI_Gatherv(sptr, 1, stype, rbuf, rcounts, displs, MPI_INT,
            root, comm);

```

Note that a different amount of data is received from each MPI process.

Example 6.8. Same as Example 6.7, but done in a different way at the sending end. We create a datatype that causes the correct striding at the sending end so that we read a column of a C array. A similar thing was done in Example 5.16, Section 5.1.14.

```

MPI_Comm comm;
int gsize, sendarray[100][150], *sptr;
int root, *rbuf, stride, myrank;
MPI_Datatype stype;
int *displs, i, *rcounts;
...

MPI_Comm_size(comm, &gsize);
MPI_Comm_rank(comm, &myrank);
rbuf = (int *)malloc(gsize*stride*sizeof(int));
displs = (int *)malloc(gsize*sizeof(int));
rcounts = (int *)malloc(gsize*sizeof(int));
for (i=0; i<gsize; ++i) {
    displs[i] = i*stride;
    rcounts[i] = 100-i;
}

```

```

1  }
2  /* Create datatype for one int, with extent of entire row
3  */
4  MPI_Type_create_resized(MPI_INT, 0, 150*sizeof(int), &stype);
5  MPI_Type_commit(&stype);
6  sptr = &sendarray[0][myrank];
7  MPI_Gatherv(sptr, 100-myrank, stype, rbuf, rcounts, displs, MPI_INT,
8             root, comm);
9

```

Example 6.9. Same as Example 6.7 at sending side, but at receiving side we make the stride between received blocks vary from block to block. See Figure 6.8.

```

10
11
12
13 MPI_Comm comm;
14 int gsize, sendarray[100][150], *sptr;
15 int root, *rbuf, *stride, myrank, bufsize;
16 MPI_Datatype stype;
17 int *displs, i, *rcounts, offset;
18
19 ...
20 MPI_Comm_size(comm, &gsize);
21 MPI_Comm_rank(comm, &myrank);
22
23 stride = (int *)malloc(gsize*sizeof(int));
24 ...
25 /* stride[i] for i = 0 to gsize-1 is set somehow
26 */
27
28 /* set up displs and rcounts vectors first
29 */
30 displs = (int *)malloc(gsize*sizeof(int));
31 rcounts = (int *)malloc(gsize*sizeof(int));
32 offset = 0;
33 for (i=0; i<gsize; ++i) {
34     displs[i] = offset;
35     offset += stride[i];
36     rcounts[i] = 100-i;
37 }
38
39 /* the required buffer size for rbuf is now easily obtained
40 */
41 bufsize = displs[gsize-1]+rcounts[gsize-1];
42 rbuf = (int *)malloc(bufsize*sizeof(int));
43 /* Create datatype for the column we are sending
44 */
45 MPI_Type_vector(100-myrank, 1, 150, MPI_INT, &stype);
46 MPI_Type_commit(&stype);
47 sptr = &sendarray[0][myrank];
48 MPI_Gatherv(sptr, 1, stype, rbuf, rcounts, displs, MPI_INT,
49             root, comm);

```





Figure 6.8: The root gathers $100-i$ ints from column i of a 100×150 C array, and each set is placed $\text{stride}[i]$ ints apart (a varying stride).

Example 6.10. MPI process i sends num ints from the i -th column of a 100×150 int array, in C. The complicating factor is that the various values of num are not known to *root*, so a separate gather must first be run to find these out. The data is placed contiguously at the receiving end.

```

MPI_Comm comm;
int gsize, sendarray[100][150], *sptr;
int root, *rbuf, myrank;
MPI_Datatype stype;
int *displs, i, *rcounts, num;

...

MPI_Comm_size(comm, &gsize);
MPI_Comm_rank(comm, &myrank);

/* First, gather nums to root
*/
rcounts = (int *)malloc(gsize*sizeof(int));
MPI_Gather(&num, 1, MPI_INT, rcounts, 1, MPI_INT, root, comm);
/* root now has correct rcounts, using these we set displs[] so that
* data is placed contiguously (or concatenated) at the receiving end
*/
displs = (int *)malloc(gsize*sizeof(int));
displs[0] = 0;
for (i=1; i<gsize; ++i) {
    displs[i] = displs[i-1]+rcounts[i-1];
}
/* And, create receive buffer
*/
rbuf = (int *)malloc(gsize*(displs[gsize-1]+rcounts[gsize-1])
                    *sizeof(int));
/* Create datatype for one int, with extent of entire row
*/
MPI_Type_create_resized(MPI_INT, 0, 150*sizeof(int), &stype);
MPI_Type_commit(&stype);
sptr = &sendarray[0][myrank];
MPI_Gatherv(sptr, num, stype, rbuf, rcounts, displs, MPI_INT,

```

```
1  root, comm);
```

6.6 Scatter

MPI_SCATTER(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root, comm)

9	IN	sendbuf	address of send buffer (choice, significant only at root)
10			
11	IN	sendcount	number of elements sent to each MPI process
12			(non-negative integer, significant only at root)
13			
14	IN	sendtype	datatype of send buffer elements (handle, significant
15			only at root)
16	OUT	recvbuf	address of receive buffer (choice)
17	IN	recvcount	number of elements in receive buffer (non-negative
18			integer)
19			
20	IN	recvtype	datatype of receive buffer elements (handle)
21	IN	root	rank of sending MPI process (integer)
22			
23	IN	comm	communicator (handle)

C binding

```
26 int MPI_Scatter(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
27                void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
28                MPI_Comm comm)
```

```
29 int MPI_Scatter_c(const void *sendbuf, MPI_Count sendcount,
30                  MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
31                  MPI_Datatype recvtype, int root, MPI_Comm comm)
```

Fortran 2008 binding

```
34 MPI_Scatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
35            comm, ierror)
```

```
36 TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
37 INTEGER, INTENT(IN) :: sendcount, recvcount, root
38 TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
39 TYPE(*), DIMENSION(..) :: recvbuf
40 TYPE(MPI_Comm), INTENT(IN) :: comm
41 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
42 MPI_Scatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
43            comm, ierror) !(_c)
```

```
44 TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
45 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
46 TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
47 TYPE(*), DIMENSION(..) :: recvbuf
48 INTEGER, INTENT(IN) :: root
```

```

TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_SCATTER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE, ROOT,
            COMM, IERROR)

```

```

<type> SENDBUF(*), RECVBUF(*)

```

```

INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, ROOT, COMM, IERROR

```

MPI_SCATTER is the inverse operation to MPI_GATHER.

If *comm* is an intra-communicator, the outcome is *as if* the root executed *n* send operations,

```

MPI_Send(sendbuf+i·sendcount·extent(sendtype), sendcount, sendtype, i,...),

```

and each MPI process executed a receive,

```

MPI_Recv(recvbuf, recvcount, recvtype, i,...).

```

An alternative description is that the root sends a message with MPI_Send(*sendbuf*, *sendcount*·*n*, *sendtype*, ...). This message is split into *n* equal segments, the *i*-th segment is sent to the *i*-th MPI process in the group, and each MPI process receives this message as above.

The send buffer is ignored for all nonroot MPI processes.

The type signature associated with *sendcount*, *sendtype* at the root must be equal to the type signature associated with *recvcount*, *recvtype* at all MPI processes (however, the type maps may be different). This implies that the amount of data sent must be equal to the amount of data received, pairwise between each MPI process and the root. Distinct type maps between sender and receiver are still allowed.

All arguments to the function are significant on the root, while on other MPI processes, only arguments *recvbuf*, *recvcount*, *recvtype*, *root*, and *comm* are significant. The arguments *root* and *comm* must have identical values on all MPI processes.

The specification of counts and types should not cause any location on the root to be read more than once.

Rationale. Though not needed, the last restriction is imposed so as to achieve symmetry with MPI_GATHER, where the corresponding restriction (a multiple-write restriction) is necessary. (*End of rationale.*)

The “in place” option for intra-communicators is specified by passing MPI_IN_PLACE as the value of *recvbuf* at the root. In such a case, *recvcount* and *recvtype* are ignored, and the root “sends” no data to itself. The scattered vector is still assumed to contain *n* segments, where *n* is the group size; the *root*-th segment, which root should “send to itself,” is not moved.

If *comm* is an inter-communicator, then the call involves all MPI processes in the inter-communicator, but with one group (group A) defining the root. All MPI processes in the other group (group B) pass the same value in argument *root*, which is the rank of the root in group A. The root passes the value MPI_ROOT in *root*. All other MPI processes in group A pass the value MPI_PROC_NULL in *root*. Data is scattered from the root to all MPI

processes in group B. The receive buffer arguments of the MPI processes in group B must be consistent with the send buffer argument of the root.

```

1 MPI_SCATTERV(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount, recvtype, root,
2             comm)
3
4
5
6
7     IN     sendbuf           address of send buffer (choice, significant only at
8                               root)
9
10    IN     sendcounts        nonnegative integer array (of length group size)
11                               specifying the number of elements to send to each
12                               rank (significant only at root)
13
14    IN     displs            integer array (of length group size). Entry i specifies
15                               the displacement (relative to sendbuf) from which to
16                               take the outgoing data to MPI process i (significant
17                               only at root)
18
19    IN     sendtype          datatype of send buffer elements (handle, significant
20                               only at root)
21
22    OUT    recvbuf           address of receive buffer (choice)
23
24    IN     recvcount         number of elements in receive buffer (non-negative
25                               integer)
26
27    IN     recvtype          datatype of receive buffer elements (handle)
28
29    IN     root              rank of sending MPI process (integer)
30
31    IN     comm              communicator (handle)

```

C binding

```

28 int MPI_Scatterv(const void *sendbuf, const int sendcounts[],
29                const int displs[], MPI_Datatype sendtype, void *recvbuf,
30                int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm)
31
32 int MPI_Scatterv_c(const void *sendbuf, const MPI_Count sendcounts[],
33                  const MPI_Aint displs[], MPI_Datatype sendtype, void *recvbuf,
34                  MPI_Count recvcount, MPI_Datatype recvtype, int root,
35                  MPI_Comm comm)

```

Fortran 2008 binding

```

36 MPI_Scatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
37             recvtype, root, comm, ierror)
38
39     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
40     INTEGER, INTENT(IN) :: sendcounts(*), displs(*), recvcount, root
41     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
42     TYPE(*), DIMENSION(..) :: recvbuf
43     TYPE(MPI_Comm), INTENT(IN) :: comm
44     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
45
46 MPI_Scatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
47             recvtype, root, comm, ierror) !(_c)
48     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf

```

```

INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcoun
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_SCATTERV(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,
             RECVTYPE, ROOT, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,
             COMM, IERROR

```

MPI_SCATTERV is the inverse operation to MPI_GATHERV.

MPI_SCATTERV extends the functionality of MPI_SCATTER by allowing a varying count of data to be sent to each MPI process, since `sendcounts` is now an array. It also allows more flexibility as to where the data is taken from on the root, by providing an additional argument, `displs`.

If `comm` is an intra-communicator, the outcome is as if the root executed `n` send operations,

```
MPI_Send(sendbuf+displs[i]·extent(sendtype), sendcounts[i], sendtype, i,...),
```

and each MPI process executed a receive,

```
MPI_Recv(recvbuf, recvcoun, recvtype, i,...).
```

The send buffer is ignored for all nonroot MPI processes.

The type signature implied by `sendcount[i]`, `sendtype` at the root must be equal to the type signature implied by `recvcoun`, `recvtype` at MPI process `i` (however, the type maps may be different). This implies that the amount of data sent must be equal to the amount of data received, pairwise between each MPI process and the root. Distinct type maps between sender and receiver are still allowed.

All arguments to the function are significant on the root, while on other MPI processes, only arguments `recvbuf`, `recvcoun`, `recvtype`, `root`, and `comm` are significant. The arguments `root` and `comm` must have identical values on all MPI processes.

The specification of counts, types, and displacements should not cause any location on the root to be read more than once.

The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` as the value of `recvbuf` at the root. In such a case, `recvcoun` and `recvtype` are ignored, and root “sends” no data to itself. The scattered vector is still assumed to contain n segments, where n is the group size; the root-th segment, which root should “send to itself,” is not moved.

If `comm` is an inter-communicator, then the call involves all MPI processes in the inter-communicator, but with one group (group A) defining the root. All MPI processes in the other group (group B) pass the same value in argument `root`, which is the rank of the root in group A. The root MPI process passes the value `MPI_ROOT` in `root`. All other MPI processes




Figure 6.9: The root scatters sets of 100 ints to each MPI process in the group.

in group A pass the value `MPI_PROC_NULL` in root. Data is scattered from the root to all MPI processes in group B. The receive buffer arguments of the MPI processes in group B must be consistent with the send buffer argument of the root.

6.6.1 Examples using `MPI_SCATTER`, `MPI_SCATTERV`

The examples in this section use intra-communicators.

Example 6.11. The reverse of Example 6.2. Scatter sets of 100 ints from the root to each MPI process in the group. See Figure 6.9.

```

22 MPI_Comm comm;
23 int gsize,*sendbuf;
24 int root, rbuf[100];
25 ...
26 MPI_Comm_size(comm, &gsize);
27 sendbuf = (int *)malloc(gsize*100*sizeof(int));
28 ...
29 MPI_Scatter(sendbuf, 100, MPI_INT, rbuf, 100, MPI_INT, root, comm);

```

Example 6.12. The reverse of Example 6.5. The root scatters sets of 100 ints to the other MPI processes, but the sets of 100 are stride ints apart in the sending buffer. Requires use of `MPI_SCATTERV`. Assume $stride \geq 100$. See Figure 6.10.

```

34 MPI_Comm comm;
35 int gsize,*sendbuf;
36 int root, rbuf[100], i, *displs, *scounts;
37 ...
38 ...
39 MPI_Comm_size(comm, &gsize);
40 sendbuf = (int *)malloc(gsize*stride*sizeof(int));
41 ...
42 displs = (int *)malloc(gsize*sizeof(int));
43 scounts = (int *)malloc(gsize*sizeof(int));
44 for (i=0; i<gsize; ++i) {
45     displs[i] = i*stride;
46     scounts[i] = 100;
47 }
48 MPI_Scatterv(sendbuf, scounts, displs, MPI_INT, rbuf, 100, MPI_INT,

```




Figure 6.10: The root scatters sets of 100 ints, moving by `stride` ints from send to send in the scatter.

```
root, comm);
```

Example 6.13. The reverse of Example 6.9. We have a varying stride between blocks at sending (root) end, at the receiving end we receive into the i -th column of a 100×150 C array. See Figure 6.11.

```
MPI_Comm comm;
int gsize, recvarray[100][150], *rptr;
int root, *sendbuf, myrank, *stride;
MPI_Datatype rtype;
int i, *displs, *counts, offset;
...
MPI_Comm_size(comm, &gsize);
MPI_Comm_rank(comm, &myrank);

stride = (int *)malloc(gsize*sizeof(int));
...
/* stride[i] for i = 0 to gsize-1 is set somehow
 * sendbuf comes from elsewhere
 */
...
displs = (int *)malloc(gsize*sizeof(int));
counts = (int *)malloc(gsize*sizeof(int));
offset = 0;
for (i=0; i<gsize; ++i) {
    displs[i] = offset;
    offset += stride[i];
    counts[i] = 100 - i;
}
/* Create datatype for the column we are receiving
 */
MPI_Type_vector(100-myrank, 1, 150, MPI_INT, &rtype);
MPI_Type_commit(&rtype);
rptr = &recvarray[0][myrank];
MPI_Scatterv(sendbuf, counts, displs, MPI_INT, rptr, 1, rtype,
            root, comm);
```




Figure 6.11: The root scatters blocks of $100-i$ ints into column i of a 100×150 C array. At the sending side, the blocks are $\text{stride}[i]$ ints apart.

6.7 Gather-to-all

```

MPI_ALLGATHER(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm)
  IN      sendbuf      starting address of send buffer (choice)
  IN      sendcount    number of elements in send buffer (non-negative
                        integer)
  IN      sendtype     datatype of send buffer elements (handle)
  OUT     recvbuf      address of receive buffer (choice)
  IN      recvcount    number of elements received from any MPI process
                        (non-negative integer)
  IN      recvtype     datatype of receive buffer elements (handle)
  IN      comm         communicator (handle)

```

C binding

```

int MPI_Allgather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
                 void *recvbuf, int recvcount, MPI_Datatype recvtype,
                 MPI_Comm comm)

```

```

int MPI_Allgather_c(const void *sendbuf, MPI_Count sendcount,
                   MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                   MPI_Datatype recvtype, MPI_Comm comm)

```

Fortran 2008 binding

```

MPI_Allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
              ierror)

```

```

  TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
  INTEGER, INTENT(IN) :: sendcount, recvcount
  TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
  TYPE(*), DIMENSION(..) :: recvbuf
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```



```

MPI_Allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
              ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, COMM,
              IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, IERROR

```

MPI_ALLGATHER can be thought of as MPI_GATHER, but where all MPI processes receive the result, instead of just the root. The block of data sent from the j -th MPI process is received by every MPI process and placed in the j -th block of the buffer `recvbuf`.

The type signature associated with `sendcount`, `sendtype`, at an MPI process must be equal to the type signature associated with `recvcount`, `recvtype` at any other MPI process.

If `comm` is an intra-communicator, the outcome of a call to MPI_ALLGATHER(...) is as if all MPI processes executed n calls to

```

MPI_Gather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
           recvtype, root, comm)

```

for `root = 0, ..., n-1`. The rules for correct usage of MPI_ALLGATHER can be found in the corresponding rules for MPI_GATHER (see Section 6.5).

The “in place” option for intra-communicators is specified by passing the value MPI_IN_PLACE to the argument `sendbuf` at all MPI processes. `sendcount` and `sendtype` are ignored. Then the input data of each MPI process is assumed to be in the area where that MPI process would receive its own contribution to the receive buffer.

If `comm` is an inter-communicator, then each MPI process of one group (group A) contributes `sendcount` data items; these data are concatenated and the result is stored at each MPI process in the other group (group B). Conversely the concatenation of the contributions of the MPI processes in group B is stored at each MPI process in group A. The send buffer arguments in group A must be consistent with the receive buffer arguments in group B, and vice versa.

Advice to users. In the inter-communicator case, the communication pattern of MPI_ALLGATHER need not be symmetric. The number of items sent by MPI processes in group A (as specified by the arguments `sendcount`, `sendtype` in group A and the arguments `recvcount`, `recvtype` in group B), need not equal the number of items sent by MPI processes in group B (as specified by the arguments `sendcount`, `sendtype` in group B and the arguments `recvcount`, `recvtype` in group A). In particular, one can move data in only one direction by specifying `sendcount = 0` for the communication in the reverse direction. (*End of advice to users.*)

```

1 MPI_ALLGATHERV(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs, recvtype,
2               comm)
3
4   IN      sendbuf      starting address of send buffer (choice)
5   IN      sendcount    number of elements in send buffer (non-negative
6                       integer)
7   IN      sendtype     datatype of send buffer elements (handle)
8   OUT     recvbuf      address of receive buffer (choice)
9
10  IN      recvcnts      nonnegative integer array (of length group size)
11                       containing the number of elements that are received
12                       from each MPI process
13  IN      displs       integer array (of length group size). Entry i specifies
14                       the displacement (relative to recvbuf) at which to
15                       place the incoming data from MPI process i
16
17  IN      recvtype     datatype of receive buffer elements (handle)
18  IN      comm         communicator (handle)
19

```

C binding

```

20
21 int MPI_Allgatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
22                 void *recvbuf, const int recvcnts[], const int displs[],
23                 MPI_Datatype recvtype, MPI_Comm comm)
24
25 int MPI_Allgatherv_c(const void *sendbuf, MPI_Count sendcount,
26                    MPI_Datatype sendtype, void *recvbuf,
27                    const MPI_Count recvcnts[], const MPI_Aint displs[],
28                    MPI_Datatype recvtype, MPI_Comm comm)
29

```

Fortran 2008 binding

```

30 MPI_Allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
31               recvtype, comm, ierror)
32     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
33     INTEGER, INTENT(IN) :: sendcount, recvcnts(*), displs(*)
34     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
35     TYPE(*), DIMENSION(..) :: recvbuf
36     TYPE(MPI_Comm), INTENT(IN) :: comm
37     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
38
39 MPI_Allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
40               recvtype, comm, ierror) !(_c)
41     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
42     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcnts(*)
43     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
44     TYPE(*), DIMENSION(..) :: recvbuf
45     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
46     TYPE(MPI_Comm), INTENT(IN) :: comm
47     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48

```

Fortran binding

```

MPI_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
               RECVTYPE, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
               IERROR

```

MPI_ALLGATHERV can be thought of as MPI_GATHERV, but where all processes receive the result, instead of just the root. The block of data sent from the j -th MPI process is received by every MPI process and placed in the j -th block of the buffer `recvbuf`. These blocks need not all be the same size.

The type signature associated with `sendcount`, `sendtype`, at MPI process j must be equal to the type signature associated with `recvcounts[j]`, `recvtype` at any other MPI process.

If `comm` is an intra-communicator, the outcome is as if all MPI processes executed calls to

```

MPI_Gatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
            recvtype, root, comm),

```

for `root = 0, ..., n-1`. The rules for correct usage of MPI_ALLGATHERV can be found in the corresponding rules for MPI_GATHERV (see Section 6.5).

The “in place” option for intra-communicators is specified by passing the value `MPI_IN_PLACE` to the argument `sendbuf` at all MPI processes. In such a case, `sendcount` and `sendtype` are ignored, and the input data of each MPI process is assumed to be in the area where that MPI process would receive its own contribution to the receive buffer.

If `comm` is an inter-communicator, then each MPI process of one group (group A) contributes `sendcount` data items; these data are concatenated and the result is stored at each MPI process in the other group (group B). Conversely the concatenation of the contributions of the MPI processes in group B is stored at each MPI process in group A. The send buffer arguments in group A must be consistent with the receive buffer arguments in group B, and vice versa.

6.7.1 Example using MPI_ALLGATHER

The example in this section uses intra-communicators.

Example 6.14. The all-gather version of Example 6.2. Using MPI_ALLGATHER, we will gather 100 ints from every MPI process in the group to every MPI process.

```

MPI_Comm comm;
int gsize, sendarray[100];
int *rbuf;
...
MPI_Comm_size(comm, &gsize);
rbuf = (int *)malloc(gsize*100*sizeof(int));
MPI_Allgather(sendarray, 100, MPI_INT, rbuf, 100, MPI_INT, comm);

```

After the call, every MPI process has the group-wide concatenation of the sets of data.

6.8 All-to-All Scatter/Gather

```

1 MPI_ALLTOALL(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm)
2
3
4
5 MPI_ALLTOALL(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm)
6     IN      sendbuf      starting address of send buffer (choice)
7     IN      sendcount    number of elements sent to each MPI process
8                          (non-negative integer)
9
10    IN      sendtype     datatype of send buffer elements (handle)
11    OUT     recvbuf      address of receive buffer (choice)
12    IN      recvcount    number of elements received from any MPI process
13                          (non-negative integer)
14    IN      recvtype     datatype of receive buffer elements (handle)
15    IN      comm         communicator (handle)
16
17

```

C binding

```

18
19 int MPI_Alltoall(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
20                void *recvbuf, int recvcount, MPI_Datatype recvtype,
21                MPI_Comm comm)
22

```

```

23 int MPI_Alltoall_c(const void *sendbuf, MPI_Count sendcount,
24                  MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
25                  MPI_Datatype recvtype, MPI_Comm comm)
26

```

Fortran 2008 binding

```

27 MPI_Alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
28             ierror)
29

```

```

30     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
31     INTEGER, INTENT(IN) :: sendcount, recvcount
32     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
33     TYPE(*), DIMENSION(..) :: recvbuf
34     TYPE(MPI_Comm), INTENT(IN) :: comm
35     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36

```

```

37 MPI_Alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
38             ierror) !(_c)
39

```

```

40     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
41     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
42     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
43     TYPE(*), DIMENSION(..) :: recvbuf
44     TYPE(MPI_Comm), INTENT(IN) :: comm
45     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
46

```

Fortran binding

```

47 MPI_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, COMM,
48             IERROR)
49

```

```

50     <type> SENDBUF(*), RECVBUF(*)
51     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, IERROR
52

```

MPI_ALLTOALL is an extension of MPI_ALLGATHER to the case where each MPI process sends distinct data to each of the receivers. The j -th block sent from MPI process i is received by MPI process j and is placed in the i -th block of `recvbuf`.

The type signature associated with `sendcount`, `sendtype`, at an MPI process must be equal to the type signature associated with `recvcount`, `recvtype` at any other MPI process. This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of MPI processes. As usual, however, the type maps may be different.

If `comm` is an intra-communicator, the outcome is as if each MPI process executed a send to each MPI process (itself included) with a call to,

```
MPI_Send(sendbuf+i·sendcount·extent(sendtype),sendcount,sendtype,i, ...),
```

and a receive from every other MPI process with a call to,

```
MPI_Recv(recvbuf+i·recvcount·extent(recvtype),recvcount,recvtype,i,...).
```

All arguments on all MPI processes are significant. The argument `comm` must have identical values on all MPI processes.

The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` to the argument `sendbuf` at *all* MPI processes. In such a case, `sendcount` and `sendtype` are ignored. The data to be sent is taken from the `recvbuf` and replaced by the received data. Data sent and received must have the same type map as specified by `recvcount` and `recvtype`.

Rationale. For large MPI_ALLTOALL instances, allocating both send and receive buffers may consume too much memory. The “in place” option effectively halves the application memory consumption and is useful in situations where the data to be sent will not be used by the sending MPI process after the MPI_ALLTOALL exchange (e.g., in parallel Fast Fourier Transforms). (*End of rationale.*)

Advice to implementors. Users may opt to use the “in place” option in order to conserve memory. Quality MPI implementations should thus strive to minimize system buffering. (*End of advice to implementors.*)

If `comm` is an inter-communicator, then the outcome is as if each MPI process in group A sends a message to each MPI process in group B, and vice versa. The j -th send buffer of MPI process i in group A should be consistent with the i -th receive buffer of MPI process j in group B, and vice versa.

Advice to users. When a complete exchange is executed in the inter-communicator case, then the number of data items sent from MPI processes in group A to MPI processes in group B need not equal the number of items sent in the reverse direction. In particular, one can have unidirectional communication by specifying `sendcount = 0` in the reverse direction. (*End of advice to users.*)

```

1 MPI_ALLTOALLV(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcoun
2         recvtype, comm)
3
4     IN     sendbuf           starting address of send buffer (choice)
5
6     IN     sendcounts       nonnegative integer array (of length group size)
7                             specifying the number of elements to send to each
8                             rank
9
10    IN     sdispls          integer array (of length group size). Entry j specifies
11                             the displacement (relative to sendbuf) from which to
12                             take the outgoing data destined for MPI process j
13
14    IN     sendtype         datatype of send buffer elements (handle)
15
16    OUT    recvbuf          address of receive buffer (choice)
17
18    IN     recvcoun
19         counts             nonnegative integer array (of length group size)
20                             specifying the number of elements that can be
21                             received from each rank
22
23    IN     rdispls          integer array (of length group size). Entry i specifies
24                             the displacement (relative to recvbuf) at which to
25                             place the incoming data from MPI process i
26
27    IN     recvtype         datatype of receive buffer elements (handle)
28
29    IN     comm              communicator (handle)

```

C binding

```

30 int MPI_Alltoallv(const void *sendbuf, const int sendcounts[],
31                 const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
32                 const int recvcoun
33                 counts[], const int rdispls[],
34                 MPI_Datatype recvtype, MPI_Comm comm)
35
36 int MPI_Alltoallv_c(const void *sendbuf, const MPI_Count sendcounts[],
37                    const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf,
38                    const MPI_Count recvcoun
39                    counts[], const MPI_Aint rdispls[],
40                    MPI_Datatype recvtype, MPI_Comm comm)

```

Fortran 2008 binding

```

41 MPI_Alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcoun
42         counts, rdispls, recvtype, comm, ierror)
43
44     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
45     INTEGER, INTENT(IN) :: sendcounts(*), sdispls(*), recvcoun
46     counts(*), rdispls(*)
47     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
48     TYPE(*), DIMENSION(..) :: recvbuf
49     TYPE(MPI_Comm), INTENT(IN) :: comm
50     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
51
52 MPI_Alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcoun
53         counts, rdispls, recvtype, comm, ierror) !(_c)
54
55     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
56     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcoun
57     counts(*)
58     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)

```

```

TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, REVCOUNTS,
              RDISPLS, RECVTYPE, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, REVCOUNTS(*), RDISPLS(*),
              RECVTYPE, COMM, IERROR

```

MPI_ALLTOALLV adds flexibility to MPI_ALLTOALL in that the location of data for the send is specified by `sdispls` and the location of the placement of the data on the receive side is specified by `rdispls`.

If `comm` is an intra-communicator, then the j -th block sent from MPI process i is received by MPI process j and is placed in the i -th block of `recvbuf`. These blocks need not all have the same size.

The type signature associated with `sendcounts[j]`, `sendtype` at MPI process i must be equal to the type signature associated with `recvcounts[i]`, `recvtype` at MPI process j . This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of MPI processes. Distinct type maps between sender and receiver are still allowed.

The outcome is as if each MPI process sent a message to every other MPI process with,

```
MPI_Send(sendbuf+sdispls[i]·extent(sendtype),sendcounts[i],sendtype,i,...),
```

and received a message from every other MPI process with a call to

```
MPI_Recv(recvbuf+rdispls[i]·extent(recvtype),recvcounts[i],recvtype,i,...).
```

All arguments on all MPI processes are significant. The argument `comm` must have identical values on all MPI processes.

The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` to the argument `sendbuf` at *all* MPI processes. In such a case, `sendcounts`, `sdispls` and `sendtype` are ignored. The data to be sent is taken from the `recvbuf` and replaced by the received data. Data sent and received must have the same type map as specified by the `recvcounts` array and the `recvtype`, and is taken from the locations of the receive buffer specified by `rdispls`.

Advice to users. Specifying the “in place” option (which must be given on all MPI processes) implies that the same amount and type of data is sent and received between any two MPI processes in the group of the communicator. Different pairs of MPI processes can exchange different amounts of data. Users must ensure that `recvcounts[j]` and `recvtype` on MPI process i match `recvcounts[i]` and `recvtype` on MPI process j . This symmetric exchange can be useful in applications where the data to be sent will not be used by the sending MPI process after the MPI_ALLTOALLV exchange. (*End of advice to users.*)

If `comm` is an inter-communicator, then the outcome is as if each MPI process in group A sends a message to each MPI process in group B, and vice versa. The j -th send buffer of

1 MPI process i in group A should be consistent with the i -th receive buffer of MPI process j
 2 in group B, and vice versa.

3
 4 *Rationale.* The definitions of MPI_ALLTOALL and MPI_ALLTOALLV give as much
 5 flexibility as one would achieve by specifying n independent, point-to-point communi-
 6 cations, with two exceptions: all messages use the same datatype, and messages are
 7 scattered from (or gathered to) sequential storage. (*End of rationale.*)

8
 9 *Advice to implementors.* Although the discussion of collective communication in
 10 terms of point-to-point operation implies that each message is transferred directly
 11 from sender to receiver, implementations may use a tree communication pattern.
 12 Messages can be forwarded by intermediate nodes where they are split (for scatter) or
 13 concatenated (for gather), if this is more efficient. (*End of advice to implementors.*)

14
 15
 16 MPI_ALLTOALLW(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcoun-
 17 tains, rdispls, recvtypes, comm)
 18
 19 IN sendbuf starting address of send buffer (choice)
 20 IN sendcounts nonnegative integer array (of length group size)
 21 specifying the number of elements to send to each
 22 rank
 23 IN sdispls integer array (of length group size). Entry j specifies
 24 the displacement in bytes (relative to sendbuf) from
 25 which to take the outgoing data destined for MPI
 26 process j (array of integers)
 27
 28 IN sendtypes array of datatypes (of length group size). Entry j
 29 specifies the type of data to send to MPI process j
 30 (array of handles)
 31 OUT recvbuf address of receive buffer (choice)
 32 IN recvcoun- nonnegative integer array (of length group size)
 33 tains specifying the number of elements that can be
 34 received from each rank
 35
 36 IN rdispls integer array (of length group size). Entry i specifies
 37 the displacement in bytes (relative to recvbuf) at
 38 which to place the incoming data from MPI process i
 39 (array of integers)
 40
 41 IN recvtypes array of datatypes (of length group size). Entry i
 42 specifies the type of data received from MPI process i
 43 (array of handles)
 44
 45 IN comm communicator (handle)

46 C binding

47 int MPI_Alltoallw(const void *sendbuf, const int sendcounts[],
 48 const int sdispls[], const MPI_Datatype sendtypes[],


```

void *recvbuf, const int recvcounts[], const int rdispls[],
const MPI_Datatype recvtypes[], MPI_Comm comm)
int MPI_Alltoallw_c(const void *sendbuf, const MPI_Count sendcounts[],
const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
void *recvbuf, const MPI_Count recvcounts[],
const MPI_Aint rdispls[], const MPI_Datatype recvtypes[],
MPI_Comm comm)

```

Fortran 2008 binding

```

MPI_Alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcounts,
rdispls, recvtypes, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcounts(*), sdispls(*), recvcounts(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcounts,
rdispls, recvtypes, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcounts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_ALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF, RECVCOUNTS,
RDISPLS, RECVTYPES, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPES(*), RECVCOUNTS(*), RDISPLS(*),
RECVTYPES(*), COMM, IERROR

```

MPI_ALLTOALLW is the most general form of complete exchange. Like MPI_TYPE_CREATE_STRUCT, the most general type constructor, MPI_ALLTOALLW allows separate specification of count, displacement and datatype. In addition, to allow maximum flexibility, the displacement of blocks within the send and receive buffers is specified in bytes.

If *comm* is an intra-communicator, then the *j*-th block sent from MPI process *i* is received by MPI process *j* and is placed in the *i*-th block of *recvbuf*. These blocks need not all have the same size.

The type signature associated with *sendcounts*[*j*], *sendtypes*[*j*] at MPI process *i* must be equal to the type signature associated with *recvcounts*[*i*], *recvtypes*[*i*] at MPI process *j*. This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of MPI processes. Distinct type maps between sender and receiver are still allowed.

The outcome is as if each MPI process sent a message to every other MPI process with

1 MPI_Send(sendbuf+sdispls[i],sendcounts[i],sendtypes[i] ,i,...),

2

3 and received a message from every other MPI process with a call to

4

5 MPI_Recv(recvbuf+rdispls[i],recvcounts[i],recvtypes[i] ,i,...).

6

7 All arguments on all MPI processes are significant. The argument `comm` must describe
8 the same communicator on all MPI processes.

9

10 Like for `MPI_ALLTOALLV`, the “in place” option for intra-communicators is specified
11 by passing `MPI_IN_PLACE` to the argument `sendbuf` at *all* MPI processes. In such a case,
12 `sendcounts`, `sdispls` and `sendtypes` are ignored. The data to be sent is taken from the `recvbuf`
13 and replaced by the received data. Data sent and received must have the same type map
14 as specified by the `recvcounts` and `recvtypes` arrays, and is taken from the locations of the
receive buffer specified by `rdispls`.

15

16 If `comm` is an inter-communicator, then the outcome is as if each MPI process in group
17 A sends a message to each MPI process in group B, and vice versa. The *j*-th send buffer of
18 MPI process *i* in group A should be consistent with the *i*-th receive buffer of MPI process *j*
in group B, and vice versa.

19

20 *Rationale.* The `MPI_ALLTOALLW` function generalizes several MPI functions by
21 carefully selecting the input arguments. For example, by making all but one MPI
22 process have `sendcounts[i] = 0`, this achieves an `MPI_SCATTERW` function. (*End of*
23 *rationale.*)

24

25 6.9 Global Reduction Operations

26

27 The functions in this section perform a global reduce operation (for example sum, maximum,
28 and logical and) across all members of a group. The reduction operation can be either one of
29 a predefined list of operations, or a user-defined operation. The global reduction functions
30 come in several flavors: a reduce that returns the result of the reduction to one member of a
31 group, an all-reduce that returns this result to all members of a group, and two scan (parallel
32 prefix) operations. In addition, a reduce-scatter operation combines the functionality of a
33 reduce and of a scatter operation.

34

35 6.9.1 Reduce

36

37

38 `MPI_REDUCE`(sendbuf, recvbuf, count, datatype, op, root, comm)

39

40 IN sendbuf address of send buffer (choice)

41

42 OUT recvbuf address of receive buffer (choice, significant only at
root)

43

44 IN count number of elements in send buffer (non-negative
integer)

45

46 IN datatype datatype of elements of send buffer (handle)

47

48 IN op reduce operation (handle)

49

IN root rank of the root (integer)

```

IN      comm      communicator (handle)
1
2
C binding
3
int MPI_Reduce(const void *sendbuf, void *recvbuf, int count,
4
               MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)
5
6
int MPI_Reduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,
7
                 MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)
8
Fortran 2008 binding
9
MPI_Reduce(sendbuf, recvbuf, count, datatype, op, root, comm, ierror)
10
   TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
11
   TYPE(*), DIMENSION(..) :: recvbuf
12
   INTEGER, INTENT(IN) :: count, root
13
   TYPE(MPI_Datatype), INTENT(IN) :: datatype
14
   TYPE(MPI_Op), INTENT(IN) :: op
15
   TYPE(MPI_Comm), INTENT(IN) :: comm
16
   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
17
MPI_Reduce(sendbuf, recvbuf, count, datatype, op, root, comm, ierror) !(_c)
18
   TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
19
   TYPE(*), DIMENSION(..) :: recvbuf
20
   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
21
   TYPE(MPI_Datatype), INTENT(IN) :: datatype
22
   TYPE(MPI_Op), INTENT(IN) :: op
23
   INTEGER, INTENT(IN) :: root
24
   TYPE(MPI_Comm), INTENT(IN) :: comm
25
   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27
Fortran binding
28
MPI_REDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, IERROR)
29
   <type> SENDBUF(*), RECVBUF(*)
30
   INTEGER COUNT, DATATYPE, OP, ROOT, COMM, IERROR
31

```

If `comm` is an intra-communicator, `MPI_REDUCE` combines the elements provided in the input buffer of each MPI process in the group, using the operation `op`, and returns the combined value in the output buffer of the MPI process with rank `root`. The input buffer is defined by the arguments `sendbuf`, `count` and `datatype`; the output buffer is defined by the arguments `recvbuf`, `count` and `datatype`; both have the same number of elements, with the same type. The routine is called by all group members using the same arguments for `count`, `datatype`, `op`, `root` and `comm`. Thus, all MPI processes provide input buffers of the same length, with elements of the same type as the output buffer at the root. Each MPI process can provide one element, or a sequence of elements, in which case the combine operation is executed element-wise on each entry of the sequence. For example, if the operation is `MPI_MAX` and the send buffer contains two elements that are floating point numbers (`count = 2` and `datatype = MPI_FLOAT`), then `recvbuf(1) = global max(sendbuf(1))` and `recvbuf(2) = global max(sendbuf(2))`.

Section 6.9.2, lists the set of predefined operations provided by MPI. That section also enumerates the datatypes to which each operation can be applied.

In addition, users may define their own operations that can be overloaded to operate on several datatypes, either basic or derived. This is further explained in Section 6.9.5.

1 The operation `op` is always assumed to be associative. All predefined operations are also
2 assumed to be commutative. Users may define operations that are assumed to be associative,
3 but not commutative. The “canonical” evaluation order of a reduction is determined by the
4 ranks of the MPI processes in the group. However, the implementation can take advantage of
5 associativity, or associativity and commutativity in order to change the order of evaluation.
6 This may change the result of the reduction for operations that are not strictly associative
7 and commutative, such as floating point addition.

8
9 *Advice to implementors.* It is strongly recommended that `MPI_REDUCE` be im-
10 plemented so that the same result be obtained whenever the function is applied on
11 the same arguments, appearing in the same order. Note that this may prevent op-
12 timizations that take advantage of the physical location of ranks. (*End of advice to*
13 *implementors.*)

14
15 *Advice to users.* Some applications may not be able to ignore the nonassocia-
16 tive nature of floating-point operations or may use user-defined operations (see Sec-
17 tion 6.9.5) that require a special reduction order and cannot be treated as associative.
18 Such applications should enforce the order of evaluation explicitly. For example, in
19 the case of operations that require a strict left-to-right (or right-to-left) evaluation
20 order, this could be done by gathering all operands at a single MPI process (e.g.,
21 with `MPI_GATHER`), applying the reduction operation in the desired order (e.g., with
22 `MPI_REDUCE_LOCAL`), and if needed, broadcast or scatter the result to the other
23 MPI processes (e.g., with `MPI_BCAST`). (*End of advice to users.*)

24
25 The `datatype` argument of `MPI_REDUCE` must be compatible with `op`. Predefined op-
26 erators work only with the MPI types listed in Section 6.9.2 and Section 6.9.4. Furthermore,
27 the `datatype` and `op` given for predefined operators must be the same on all MPI processes.

28 Note that it is possible for users to supply different user-defined operations to
29 `MPI_REDUCE` in each MPI process. MPI does not define which operations are used on which
30 operands in this case. User-defined operators may operate on general, derived datatypes.
31 In this case, each argument that the reduce operation is applied to is one element described
32 by such a datatype, which may contain several basic values. This is further explained in
33 Section 6.9.5.

34
35 *Advice to users.* Users should make no assumptions about how `MPI_REDUCE` is
36 implemented. It is safest to ensure that the same function is passed to `MPI_REDUCE`
37 by each MPI process. (*End of advice to users.*)

38 Overlapping datatypes are permitted in “send” buffers. Overlapping datatypes in “re-
39 ceive” buffers are erroneous and may give unpredictable results.

40 The “in place” option for intra-communicators is specified by passing the value
41 `MPI_IN_PLACE` to the argument `sendbuf` at the root. In such a case, the input data is taken
42 at the root from the receive buffer, where it will be replaced by the output data.

43 If `comm` is an inter-communicator, then the call involves all MPI processes in the inter-
44 communicator, but with one group (group A) defining the root. All MPI processes in the
45 other group (group B) pass the same value in argument `root`, which is the rank of the root
46 in group A. The root passes the value `MPI_ROOT` in `root`. All other MPI processes in group
47 A pass the value `MPI_PROC_NULL` in `root`. Only send buffer arguments are significant in
48 group B and only receive buffer arguments are significant at the root.

6.9.2 Predefined Reduction Operations

The following predefined operations are supplied for `MPI_REDUCE` and related functions `MPI_ALLREDUCE`, `MPI_REDUCE_SCATTER_BLOCK`, `MPI_REDUCE_SCATTER`, `MPI_SCAN`, `MPI_EXSCAN`, all nonblocking variants of those (see Section 6.12), and `MPI_REDUCE_LOCAL`. These operations are invoked by placing the following in `op`.

Name	Meaning
<code>MPI_MAX</code>	maximum
<code>MPI_MIN</code>	minimum
<code>MPI_SUM</code>	sum
<code>MPI_PROD</code>	product
<code>MPI_LAND</code>	logical and
<code>MPI_BAND</code>	bit-wise and
<code>MPI_LOR</code>	logical or
<code>MPI_BOR</code>	bit-wise or
<code>MPI_LXOR</code>	logical exclusive or (xor)
<code>MPI_BXOR</code>	bit-wise exclusive or (xor)
<code>MPI_MAXLOC</code>	max value and location
<code>MPI_MINLOC</code>	min value and location

The two operations `MPI_MINLOC` and `MPI_MAXLOC` are discussed separately in Section 6.9.4. For the other predefined operations, we enumerate below the allowed combinations of `op` and `datatype` arguments. First, define groups of MPI basic datatypes in the following way.

C integer:	<code>MPI_INT</code> , <code>MPI_LONG</code> , <code>MPI_SHORT</code> , <code>MPI_UNSIGNED_SHORT</code> , <code>MPI_UNSIGNED</code> , <code>MPI_UNSIGNED_LONG</code> , <code>MPI_LONG_LONG_INT</code> , <code>MPI_LONG_LONG</code> (as synonym), <code>MPI_UNSIGNED_LONG_LONG</code> , <code>MPI_SIGNED_CHAR</code> , <code>MPI_UNSIGNED_CHAR</code> , <code>MPI_INT8_T</code> , <code>MPI_INT16_T</code> , <code>MPI_INT32_T</code> , <code>MPI_INT64_T</code> , <code>MPI_UINT8_T</code> , <code>MPI_UINT16_T</code> , <code>MPI_UINT32_T</code> , and <code>MPI_UINT64_T</code>
Fortran integer:	<code>MPI_INTEGER</code> and handles returned from <code>MPI_TYPE_CREATE_F90_INTEGER</code> and, if available, <code>MPI_INTEGER1</code> , <code>MPI_INTEGER2</code> , <code>MPI_INTEGER4</code> , <code>MPI_INTEGER8</code> , and <code>MPI_INTEGER16</code>
Floating point:	<code>MPI_FLOAT</code> , <code>MPI_DOUBLE</code> , <code>MPI_REAL</code> , <code>MPI_DOUBLE_PRECISION</code> , <code>MPI_LONG_DOUBLE</code> , and handles returned from <code>MPI_TYPE_CREATE_F90_REAL</code>

1		and, if available, MPI_REAL2,
2		MPI_REAL4, MPI_REAL8, and MPI_REAL16
3	Logical:	MPI_LOGICAL, MPI_C_BOOL,
4		and MPI_CXX_BOOL
5	Complex:	MPI_COMPLEX, MPI_C_COMPLEX,
6		MPI_C_FLOAT_COMPLEX (as synonym),
7		MPI_C_DOUBLE_COMPLEX,
8		MPI_C_LONG_DOUBLE_COMPLEX,
9		MPI_CXX_FLOAT_COMPLEX,
10		MPI_CXX_DOUBLE_COMPLEX,
11		MPI_CXX_LONG_DOUBLE_COMPLEX,
12		and handles returned from
13		MPI_TYPE_CREATE_F90_COMPLEX
14		and, if available, MPI_DOUBLE_COMPLEX,
15		MPI_COMPLEX4, MPI_COMPLEX8,
16		MPI_COMPLEX16, and MPI_COMPLEX32
17	Byte:	MPI_BYTE
18	Multi-language types:	MPI_AINT, MPI_OFFSET, and MPI_COUNT

Now, the valid datatypes for each operation are specified below.

Op	Allowed Types
MPI_MAX, MPI_MIN	C integer, Fortran integer, Floating point, Multi-language types
MPI_SUM, MPI_PROD	C integer, Fortran integer, Floating point, Complex, Multi-language types
MPI_LAND, MPI_LOR, MPI_LXOR	C integer, Logical
MPI_BAND, MPI_BOR, MPI_BXOR	C integer, Fortran integer, Byte, Multi-language types

These operations together with all listed datatypes are valid in all supported programming languages, see also Reduce Operations in Section 19.3.6.

The following examples use intra-communicators.

Example 6.15. A routine that computes the dot product of two vectors that are distributed across a group of MPI processes and returns the answer at node zero.

```

SUBROUTINE PAR_BLAS1(m, a, b, c, comm)
USE MPI
REAL a(m), b(m)      ! local slice of array
REAL c              ! result (at node zero)
REAL sum
INTEGER m, comm, i, ierr

! local sum
sum = 0.0
DO i = 1, m
  sum = sum + a(i)*b(i)
END DO

```

```

! global sum
CALL MPI_REDUCE(sum, c, 1, MPI_REAL, MPI_SUM, 0, comm, ierr)
RETURN
END

```

Example 6.16. A routine that computes the product of a vector and an array that are distributed across a group of MPI processes and returns the answer at node zero.

```

SUBROUTINE PAR_BLAS2(m, n, a, b, c, comm)
USE MPI
REAL a(m), b(m,n)    ! local slice of array
REAL c(n)            ! result
REAL sum(n)
INTEGER m, n, comm, i, j, ierr

! local sum
DO j=1,n
  sum(j) = 0.0
  DO i=1,m
    sum(j) = sum(j) + a(i)*b(i,j)
  END DO
END DO

! global sum
CALL MPI_REDUCE(sum, c, n, MPI_REAL, MPI_SUM, 0, comm, ierr)

! return result at node zero (and garbage at the other nodes)
RETURN
END

```

6.9.3 Signed Characters and Reductions

The types `MPI_SIGNED_CHAR` and `MPI_UNSIGNED_CHAR` can be used in reduction operations. `MPI_CHAR`, `MPI_WCHAR`, and `MPI_CHARACTER` (which represent printable characters) cannot be used in reduction operations. In a heterogeneous environment, `MPI_CHAR`, `MPI_WCHAR`, and `MPI_CHARACTER` will be translated so as to preserve the printable character, whereas `MPI_SIGNED_CHAR` and `MPI_UNSIGNED_CHAR` will be translated so as to preserve the integer value.

Advice to users. The types `MPI_CHAR`, `MPI_WCHAR`, and `MPI_CHARACTER` are intended for characters, and so will be translated to preserve the printable representation, rather than the integer value, if sent between machines with different character codes. The types `MPI_SIGNED_CHAR` and `MPI_UNSIGNED_CHAR` should be used in C if the integer value should be preserved. (*End of advice to users.*)

6.9.4 MINLOC and MAXLOC

The operator `MPI_MINLOC` is used to compute a global minimum and also an index attached to the minimum value. `MPI_MAXLOC` similarly computes a global maximum and index. One

1 application of these is to compute a global minimum (maximum) and the rank of the MPI
 2 process containing this value.

3 The operation that defines MPI_MAXLOC is:

$$4 \quad \begin{pmatrix} u \\ i \end{pmatrix} \circ \begin{pmatrix} v \\ j \end{pmatrix} = \begin{pmatrix} w \\ k \end{pmatrix}$$

7 where

$$9 \quad w = \max(u, v)$$

11 and

$$12 \quad k = \begin{cases} i & \text{if } u > v \\ \min(i, j) & \text{if } u = v \\ j & \text{if } u < v \end{cases}$$

16 MPI_MINLOC is defined similarly:

$$17 \quad \begin{pmatrix} u \\ i \end{pmatrix} \circ \begin{pmatrix} v \\ j \end{pmatrix} = \begin{pmatrix} w \\ k \end{pmatrix}$$

21 where

$$22 \quad w = \min(u, v)$$

24 and

$$25 \quad k = \begin{cases} i & \text{if } u < v \\ \min(i, j) & \text{if } u = v \\ j & \text{if } u > v \end{cases}$$

29 Both operations are associative and commutative. Note that if MPI_MAXLOC is applied
 30 to reduce a sequence of pairs $(u_0, 0), (u_1, 1), \dots, (u_{n-1}, n-1)$, then the value returned is
 31 (u, r) , where $u = \max_i u_i$ and r is the index of the first global maximum in the sequence.
 32 Thus, if each MPI process supplies a value and its rank within the group, then a reduce
 33 operation with `op = MPI_MAXLOC` will return the maximum value and the rank of the first
 34 MPI process with that value. Similarly, MPI_MINLOC can be used to return a minimum and
 35 its index. More generally, MPI_MINLOC computes a *lexicographic minimum*, where elements
 36 are ordered according to the first component of each pair, and ties are resolved according
 37 to the second component.

38 The reduce operation is defined to operate on arguments that consist of a pair: value
 39 and index. For both Fortran and C, types are provided to describe the pair. The potentially
 40 mixed-type nature of such arguments is a problem in older versions of Fortran. The problem
 41 is circumvented there by having the MPI-provided type consist of a pair of the same type
 42 as value, and coercing the index to this type also. In C, the MPI-provided pair type has
 43 distinct types and the index is an integer type. For named predefined pair types in C the
 44 index is of type `int`. For unnamed predefined pair types, other integer types are allowed as
 45 index instead. To use pair types with distinct value and index in Fortran, these types need
 46 to be defined using `BIND(C)` and be equivalent to the corresponding C struct.

47 In order to use MPI_MINLOC and MPI_MAXLOC in a reduce operation, one must pro-
 48 vide a `datatype` argument that represents a pair (value and index). MPI provides nine

such named predefined datatypes as well as the function `MPI_TYPE_GET_VALUE_INDEX` to query named and unnamed predefined types using value type and index type. The operations `MPI_MAXLOC` and `MPI_MINLOC` can be used with each of the following named datatypes.

Fortran:

Name	Description
<code>MPI_2REAL</code>	pair of REALs
<code>MPI_2DOUBLE_PRECISION</code>	pair of DOUBLE PRECISION variables
<code>MPI_2INTEGER</code>	pair of INTEGERS

C:

Name	Description
<code>MPI_FLOAT_INT</code>	float and int
<code>MPI_DOUBLE_INT</code>	double and int
<code>MPI_LONG_INT</code>	long and int
<code>MPI_2INT</code>	pair of int
<code>MPI_SHORT_INT</code>	short and int
<code>MPI_LONG_DOUBLE_INT</code>	long double and int

The datatype `MPI_2REAL` is *as if* defined by the following (see Section 5.1).

```
MPI_Type_contiguous(2, MPI_REAL, MPI_2REAL);
```

Similar statements apply for `MPI_2INTEGER`, `MPI_2DOUBLE_PRECISION`, and `MPI_2INT`.

The datatype `MPI_SHORT_INT` is *as if* defined by the following sequence of instructions.

```
struct mystruct {
    short val;
    int rank;
};
type[0] = MPI_SHORT;
type[1] = MPI_INT;
disp[0] = 0;
disp[1] = offsetof(struct mystruct, rank);
block[0] = 1;
block[1] = 1;
MPI_Type_create_struct(2, block, disp, type, &MPI_SHORT_INT);
MPI_Type_commit(&MPI_SHORT_INT);
```

Similar statements apply for `MPI_FLOAT_INT`, `MPI_LONG_INT` and `MPI_DOUBLE_INT`.

```
MPI_Type_get_value_index(value_type, index_type, pair_type)
```

IN	value_type	datatype of the value in pair (handle)
IN	index_type	datatype of the index in pair (handle)
OUT	pair_type	datatype of the value-index pair (handle)

C binding

```
int MPI_Type_get_value_index(MPI_Datatype value_type, MPI_Datatype index_type,
                             MPI_Datatype *pair_type)
```

Fortran 2008 binding

```

1 MPI_Type_get_value_index(value_type, index_type, pair_type, ierror)
2   TYPE(MPI_Datatype), INTENT(IN) :: value_type, index_type
3   TYPE(MPI_Datatype), INTENT(OUT) :: pair_type
4   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

7 MPI_TYPE_GET_VALUE_INDEX(VALUE_TYPE, INDEX_TYPE, PAIR_TYPE, IERROR)
8   INTEGER VALUE_TYPE, INDEX_TYPE, PAIR_TYPE, IERROR

```

MPI_TYPE_GET_VALUE_INDEX returns a handle to a predefined datatype suitable for the use with MPI_MINLOC and MPI_MAXLOC if such a predefined type exists. If the provided combination of `value_type` and `index_type` does not match a predefined pair datatype (named or unnamed), the function will set `pair_type` to `MPI_DATATYPE_NULL` and return `MPI_SUCCESS`. The returned type is not a duplicate. This type cannot be freed. Types supported by the underlying compiler for which the operators `MPI_MIN` and `MPI_MAX` are defined in Section 6.9.2 are acceptable value types. Integer types supported by the underlying compiler are acceptable index types.

Advice to users. Note that a named type handle returned by `MPI_TYPE_GET_VALUE_INDEX` will yield the combiner value `MPI_COMBINER_NAMED` when queried with `MPI_TYPE_GET_ENVELOPE` to ensure backward compatibility to existing behavior, whereas all unnamed type handles returned by `MPI_TYPE_GET_VALUE_INDEX` will yield the combiner value `MPI_COMBINER_VALUE_INDEX` when queried with `MPI_TYPE_GET_ENVELOPE`. There is no observable difference between the named constant value used via its symbol name or via `MPI_TYPE_GET_VALUE_INDEX`. Code evaluating the combiner of type handles returned from `MPI_TYPE_GET_VALUE_INDEX` must therefore handle both `MPI_COMBINER_NAMED` and `MPI_COMBINER_VALUE_INDEX`. (*End of advice to users.*)

Example 6.17. An unnamed predefined value-index type is retrieved for use with the corresponding C struct. If the requested value-index pair does not exist as a predefined type `MPI_DATATYPE_NULL` is returned.

```

34 struct mystruct {
35     double val;
36     uint64_t index;
37 };
38 MPI_Datatype dtype;
39 MPI_Type_get_value_index(MPI_DOUBLE, MPI_UINT64_T, &dtype);
40
41 if (dtype == MPI_DATATYPE_NULL) {
42     // Handling for unsupported value-index type
43 }

```

Advice to users. Implementations may apply certain optimizations to operations on compound types with equally sized value and index types. Such optimizations may not be applicable to operations on compound types where value and index type are of different size. (*End of advice to users.*)

The following examples use intra-communicators.

Example 6.18. Each MPI process has an array of 30 doubles, in C. For each of the 30 locations, compute the value and rank of the MPI process containing the largest value.

```

...
/* each MPI process has an array of 30 double: ain[30]
*/
double ain[30], aout[30];
int ind[30];
struct {
    double val;
    int rank;
} in[30], out[30];
int i, myrank, root;

MPI_Comm_rank(comm, &myrank);
for (i=0; i<30; ++i) {
    in[i].val = ain[i];
    in[i].rank = myrank;
}
MPI_Reduce(in, out, 30, MPI_DOUBLE_INT, MPI_MAXLOC, root, comm);
/* At this point, the answer resides on root MPI process
*/
if (myrank == root) {
    /* read ranks out
    */
    for (i=0; i<30; ++i) {
        aout[i] = out[i].val;
        ind[i] = out[i].rank;
    }
}

```

Example 6.19. Same example, in Fortran.

```

...
! each process has an array of 30 double: ain(30)

DOUBLE PRECISION ain(30), aout(30)
INTEGER ind(30)
DOUBLE PRECISION in(2,30), out(2,30)
INTEGER i, myrank, root, ierr

CALL MPI_COMM_RANK(comm, myrank, ierr)
DO i=1,30
    in(1,i) = ain(i)
    in(2,i) = myrank ! myrank is coerced to a double
END DO

CALL MPI_REDUCE(in, out, 30, MPI_2DOUBLE_PRECISION, MPI_MAXLOC, root,&
    comm, ierr)
! At this point, the answer resides on root MPI process

```

```

1  IF (myrank .EQ. root) THEN
2      ! read ranks out
3      DO i=1,30
4          aout(i) = out(1,i)
5          ind(i) = out(2,i) ! rank is coerced back to an integer
6      END DO
7  END IF

```

Example 6.20. Each MPI process has a nonempty array of values. Find the minimum global value, the rank of the MPI process that holds it and its index on this MPI process.

```

12 #define LEN 1000
13
14 float val[LEN]; /* local array of values */
15 int count; /* local number of values */
16 int myrank, minrank, minindex;
17 float minval;
18
19 struct {
20     float value;
21     int index;
22 } in, out;
23
24 /* local minloc */
25 in.value = val[0];
26 in.index = 0;
27 for (i=1; i < count; i++) {
28     if (in.value > val[i]) {
29         in.value = val[i];
30         in.index = i;
31     }
32 }
33 /* global minloc */
34 MPI_Comm_rank(comm, &myrank);
35 in.index = myrank*LEN + in.index;
36 MPI_Reduce(&in, &out, 1, MPI_FLOAT_INT, MPI_MINLOC, root, comm);
37
38 /* At this point, the answer resides on the root */
39 if (myrank == root) {
40     /* read answer out */
41     minval = out.value;
42     minrank = out.index / LEN;
43     minindex = out.index % LEN;
44 }

```

Rationale. The definition of MPI_MINLOC and MPI_MAXLOC given here has the advantage that it does not require any special-case handling of these two operations: they are handled like any other reduce operation. By assigning a value other than myrank to the in.index field, a programmer can provide a different definition of MPI_MAXLOC and MPI_MINLOC, if so desired. The disadvantage is that values and indices have to be first interleaved, and that indices and values have to be coerced to

the same type, in Fortran. (*End of rationale.*)

6.9.5 User-Defined Reduction Operations

`MPI_OP_CREATE(user_fn, commute, op)`

IN	<code>user_fn</code>	user defined function (function)
IN	<code>commute</code>	true if commutative; false otherwise.
OUT	<code>op</code>	operation (handle)

C binding

```
int MPI_Op_create(MPI_User_function *user_fn, int commute, MPI_Op *op)
```

```
int MPI_Op_create_c(MPI_User_function_c *user_fn, int commute, MPI_Op *op)
```

Fortran 2008 binding

```
MPI_Op_create(user_fn, commute, op, ierror)
```

```
  PROCEDURE(MPI_User_function) :: user_fn
```

```
  LOGICAL, INTENT(IN) :: commute
```

```
  TYPE(MPI_Op), INTENT(OUT) :: op
```

```
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Op_create_c(user_fn, commute, op, ierror) !(_c)
```

```
  PROCEDURE(MPI_User_function_c) :: user_fn
```

```
  LOGICAL, INTENT(IN) :: commute
```

```
  TYPE(MPI_Op), INTENT(OUT) :: op
```

```
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_OP_CREATE(USER_FN, COMMUTE, OP, IERROR)
```

```
  EXTERNAL USER_FN
```

```
  LOGICAL COMMUTE
```

```
  INTEGER OP, IERROR
```

`MPI_OP_CREATE` binds a user-defined reduction operation to an `op` handle that can subsequently be used in `MPI_REDUCE`, `MPI_ALLREDUCE`, `MPI_REDUCE_SCATTER`, `MPI_REDUCE_SCATTER_BLOCK`, `MPI_SCAN`, `MPI_EXSCAN`, all nonblocking and persistent variants of those (see Section 6.12 and Section 6.13), and `MPI_REDUCE_LOCAL`. The user-defined operation is assumed to be associative. If `commute = true`, then the operation should be both commutative and associative. If `commute = false`, then the order of operands is fixed and is defined to be in ascending, process rank order, beginning with MPI process with rank 0 in the communicator `comm`. The order of evaluation can be changed, taking advantage of the associativity of the operation. If `commute = true` then the order of evaluation can be changed, taking advantage of commutativity and associativity.

In Fortran when using `USE mpi_f08`, the large count variant shall be called explicitly as `MPI_Op_create_c` (i.e., with suffix “_c”) because interface polymorphism cannot be used to differentiate between the two different user callback prototypes despite their different type signatures.

1 The argument `user_fn` is the user-defined function, which must have the following four
2 arguments: `invec`, `inoutvec`, `len`, and `datatype`.

3 `MPI_USER_FUNCTION` also supports large count types in separate additional MPI
4 callback function prototype declarations in C (suffixed with the “_c”) and in Fortran when
5 using `USE mpi_f08`.

6 The ISO C prototypes for the functions are the following.

```
7 typedef void MPI_User_function(void *invec, void *inoutvec, int *len,
8     MPI_Datatype *datatype);
9
10 typedef void MPI_User_function_c(void *invec, void *inoutvec, MPI_Count *len,
11     MPI_Datatype *datatype);
```

12 The Fortran declarations of the user-defined function `user_fn` appear below.

```
13 ABSTRACT INTERFACE
14   SUBROUTINE MPI_User_function(invec, inoutvec, len, datatype)
15     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
16     TYPE(C_PTR), VALUE :: invec, inoutvec
17     INTEGER :: len
18     TYPE(MPI_Datatype) :: datatype
19
20 ABSTRACT INTERFACE
21   SUBROUTINE MPI_User_function_c(invec, inoutvec, len, datatype) !(_c)
22     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
23     TYPE(C_PTR), VALUE :: invec, inoutvec
24     INTEGER(KIND=MPI_COUNT_KIND) :: len
25     TYPE(MPI_Datatype) :: datatype
26
27 SUBROUTINE USER_FUNCTION(INVEC, INOUTVEC, LEN, DATATYPE)
28   <type> INVEC(LEN), INOUTVEC(LEN)
29   INTEGER LEN, DATATYPE
```

29 The `datatype` argument is a handle to the `datatype` that was passed into the call to
30 `MPI_REDUCE`. The user reduce function should be written such that the following holds:
31 Let $u[0], \dots, u[\text{len}-1]$ be the `len` elements in the communication buffer described by the
32 arguments `invec`, `len` and `datatype` when the function is invoked; let $v[0], \dots, v[\text{len}-1]$ be `len`
33 elements in the communication buffer described by the arguments `inoutvec`, `len` and `datatype`
34 when the function is invoked; let $w[0], \dots, w[\text{len}-1]$ be `len` elements in the communication
35 buffer described by the arguments `inoutvec`, `len` and `datatype` when the function returns;
36 then $w[i] = u[i] \circ v[i]$, for $i=0, \dots, \text{len}-1$, where \circ is the reduce operation that the function
37 computes.

38 Informally, we can think of `invec` and `inoutvec` as arrays of `len` elements that `user_fn`
39 is combining. The result of the reduction over-writes values in `inoutvec`, hence the name.
40 Each invocation of the function results in the pointwise evaluation of the reduce operator
41 on `len` elements: i.e., the function returns in `inoutvec[i]` the value $\text{invec}[i] \circ \text{inoutvec}[i]$, for
42 $i=0, \dots, \text{count}-1$, where \circ is the combining operation computed by the function.

44 *Rationale.* The `len` argument allows `MPI_REDUCE` to avoid calling the function for
45 each element in the input buffer. Rather, the system can choose to apply the function
46 to chunks of input. In C, it is passed in as a reference for reasons of compatibility
47 with Fortran.

By internally comparing the value of the `datatype` argument to known, global handles, it is possible to overload the use of a single user-defined function for several, different datatypes. (*End of rationale.*)

When calling any reduction or prefix scan MPI procedure with a user-defined MPI operator, the type of the `count` parameter in the call to the reduction or prefix scan MPI procedure does not need to be identical to the type of the `len` parameter in the user function associated with the user-defined MPI operator. If the `count` parameter has a type of `int` in C or `INTEGER` in Fortran and the `len` parameter has a type of `MPI_COUNT`, then MPI will perform the appropriate widening type conversion of the `len` parameter. If the `count` parameter has a type of `MPI_COUNT` and the `len` parameter has a type of `int` in C or `INTEGER` in Fortran, then MPI will perform the appropriate narrowing type conversion of the `len` parameter. If this narrowing conversion would result in truncation of the `len` value, then MPI will call the user function multiple times with a sequence of values for `len` that sum to the value of `count`.

Advice to implementors. If the number of data items cannot be represented in `len`, the implementation may need to invoke `user_fn` multiple times. (*End of advice to implementors.*)

General datatypes may be passed to the user function. However, use of datatypes that are not contiguous is likely to lead to inefficiencies.

No MPI communication function may be called inside the user function. `MPI_ABORT` may be called inside the function in case of an error.

Advice to users. Suppose one defines a library of user-defined reduce functions that are overloaded: the `datatype` argument is used to select the right execution path at each invocation, according to the types of the operands. The user-defined reduce function cannot “decode” the `datatype` argument that it is passed, and cannot identify, by itself, the correspondence between the datatype handles and the datatype they represent. This correspondence was established when the datatypes were created. Before the library is used, a library initialization preamble must be executed. This preamble code will define the datatypes that are used by the library, and store handles to these datatypes in global, static variables that are shared by the user code and the library code.

The Fortran version of `MPI_REDUCE` will invoke a user-defined reduce function using the Fortran calling conventions and will pass a Fortran-type datatype argument; the C version will use C calling convention and the C representation of a datatype handle. Users who plan to mix languages should define their reduction functions accordingly. (*End of advice to users.*)

Advice to implementors. We outline below a naive and inefficient implementation of `MPI_REDUCE` not supporting the “in place” option and only valid for intra-communicators.

```
MPI_Comm_size(comm, &groupsize);
MPI_Comm_rank(comm, &rank);
if (rank > 0) {
    MPI_Recv(tempbuf, count, datatype, rank-1,...);
```

```

1     User_reduce(tempbuf, sendbuf, count, datatype);
2   }
3   if (rank < groupsize-1) {
4     MPI_Send(sendbuf, count, datatype, rank+1, ...);
5   }
6   /* answer now resides in MPI process groupsize-1 ...
7    * now send to root
8    */
9   if (rank == root) {
10    MPI_Irecv(recvbuf, count, datatype, groupsize-1, ..., &req);
11  }
12  if (rank == groupsize-1) {
13    MPI_Send(sendbuf, count, datatype, root, ...);
14  }
15  if (rank == root) {
16    MPI_Wait(&req, &status);
17  }

```

The reduction computation proceeds, sequentially, from MPI process with rank 0 to MPI process with rank groupsize-1. This order is chosen so as to respect the order of a possibly noncommutative operator defined by the function User_reduce(). A more efficient implementation is achieved by taking advantage of associativity and using a logarithmic tree reduction. Commutativity can be used to advantage, for those cases in which the commute argument to MPI_OP_CREATE is true. Also, the amount of temporary buffer space required can be reduced, and communication can be pipelined with computation, by transferring and reducing the elements in chunks of size len < count.

The predefined reduce operations can be implemented as a library of user-defined operations. However, better performance might be achieved if MPI_REDUCE handles these functions as a special case. (*End of advice to implementors.*)

MPI_OP_FREE(op)

INOUT op operation (handle)

C binding

int MPI_Op_free(MPI_Op *op)

Fortran 2008 binding

```

MPI_Op_free(op, ierror)
  TYPE(MPI_Op), INTENT(INOUT) :: op
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_OP_FREE(OP, IERROR)
  INTEGER OP, IERROR

```

Marks a user-defined reduction operation for deallocation and sets op to MPI_OP_NULL.

Example of User-Defined Reduce

The example in this section uses an intra-communicator.

Example 6.21. Compute the product of an array of complex numbers, in C.

```

typedef struct {
    double real, imag;
} Complex;

/* the user-defined function
*/
void myProd(void *inP, void *inoutP, int *len, MPI_Datatype *dptr)
{
    int i;
    Complex c;
    Complex *in = (Complex *)inP, *inout = (Complex *)inoutP;

    for (i=0; i < *len; ++i) {
        c.real = inout->real*in->real -
                inout->imag*in->imag;
        c.imag = inout->real*in->imag +
                inout->imag*in->real;
        *inout = c;
        in++; inout++;
    }
}

/* and, to call it...
*/
...

/* each MPI process has an array of 100 Complexes
*/
Complex a[100], answer[100];
MPI_Op myOp;
MPI_Datatype ctype;

/* explain to MPI how type Complex is defined
*/
MPI_Type_contiguous(2, MPI_DOUBLE, &ctype);
MPI_Type_commit(&ctype);
/* create the complex-product user-op
*/
MPI_Op_create(myProd, 1, &myOp);

MPI_Reduce(a, answer, 100, ctype, myOp, root, comm);

/* At this point, the answer, which consists of 100 Complexes,
* resides on root MPI process
*/

```

Example 6.22. How to use the `mpi_f08` interface of the Fortran `MPI_User_function`.

```

1  subroutine my_user_function(invec, inoutvec, len, dtype)  bind(c)
2  use, intrinsic :: iso_c_binding, only : c_ptr, c_f_pointer
3  use mpi_f08
4  type(c_ptr), value :: invec, inoutvec
5  integer :: len
6  type(MPI_Datatype) :: dtype
7  real, pointer :: invec_r(:), inoutvec_r(:)
8  if (dtype == MPI_REAL) then
9      call c_f_pointer(invec, invec_r, (/ len /))
10     call c_f_pointer(inoutvec, inoutvec_r, (/ len /))
11     inoutvec_r = invec_r + inoutvec_r
12 end if
13 end subroutine
14
15
16

```

6.9.6 All-Reduce

MPI includes a variant of the reduce operations where the result is returned to all MPI processes in a group. MPI requires that all MPI processes from the same group participating in these operations receive identical results.

`MPI_ALLREDUCE`(sendbuf, recvbuf, count, datatype, op, comm)

IN	sendbuf	starting address of send buffer (choice)
OUT	recvbuf	starting address of receive buffer (choice)
IN	count	number of elements in send buffer (non-negative integer)
IN	datatype	datatype of elements of send buffer (handle)
IN	op	operation (handle)
IN	comm	communicator (handle)

C binding

```

34 int MPI_Allreduce(const void *sendbuf, void *recvbuf, int count,
35                 MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
36

```

```

37 int MPI_Allreduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,
38                   MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
39

```

Fortran 2008 binding

```

40 MPI_Allreduce(sendbuf, recvbuf, count, datatype, op, comm, ierror)
41     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
42     TYPE(*), DIMENSION(..) :: recvbuf
43     INTEGER, INTENT(IN) :: count
44     TYPE(MPI_Datatype), INTENT(IN) :: datatype
45     TYPE(MPI_Op), INTENT(IN) :: op
46     TYPE(MPI_Comm), INTENT(IN) :: comm
47     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48

```

```

MPI_Allreduce(sendbuf, recvbuf, count, datatype, op, comm, ierror) !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
  TYPE(*), DIMENSION(..) :: recvbuf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_ALLREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
  <type> SENDBUF(*), RECVBUF(*)
  INTEGER COUNT, DATATYPE, OP, COMM, IERROR

```

If `comm` is an intra-communicator, `MPI_ALLREDUCE` behaves the same as `MPI_REDUCE` except that the result appears in the receive buffer of all the group members.

Advice to implementors. The all-reduce operations can be implemented as a reduce, followed by a broadcast. However, a direct implementation can lead to better performance. (*End of advice to implementors.*)

The “in place” option for intra-communicators is specified by passing the value `MPI_IN_PLACE` to the argument `sendbuf` at all processes. value `MPI_IN_PLACE` to the argument `sendbuf` at all MPI processes. In this case, the input data is taken at each MPI process from the receive buffer, where it will be replaced by the output data.

If `comm` is an inter-communicator, then the result of the reduction of the data provided by MPI processes in group A is stored at each MPI process in group B, and vice versa. Both groups should provide `count` and `datatype` arguments that specify the same type signature.

The following example uses an intra-communicator.

Example 6.23. A routine that computes the product of a vector and an array that are distributed across a group of MPI processes and returns the answer at all nodes (see also Example 6.16).

```

SUBROUTINE PAR_BLAS2(m, n, a, b, c, comm)
USE MPI
REAL a(m), b(m,n)      ! local slice of array
REAL c(n)              ! result
REAL sum(n)
INTEGER m, n, comm, i, j, ierr

! local sum
DO j=1,n
  sum(j) = 0.0
  DO i=1,m
    sum(j) = sum(j) + a(i)*b(i,j)
  END DO
END DO

! global sum
CALL MPI_ALLREDUCE(sum, c, n, MPI_REAL, MPI_SUM, comm, ierr)

```

```

1  ! return result at all nodes
2  RETURN
3  END

```

6.9.7 MPI Process-Local Reduction

The functions in this section are of importance to library implementors who may want to implement special reduction patterns that are otherwise not easily covered by the standard MPI operations.

The following function applies a reduction operator to local arguments.

MPI_REDUCE_LOCAL(inbuf, inoutbuf, count, datatype, op)

IN	inbuf	input buffer (choice)
INOUT	inoutbuf	combined input and output buffer (choice)
IN	count	number of elements in inbuf and inoutbuf buffers (non-negative integer)
IN	datatype	datatype of elements of inbuf and inoutbuf buffers (handle)
IN	op	operation (handle)

C binding

```
int MPI_Reduce_local(const void *inbuf, void *inoutbuf, int count,
                    MPI_Datatype datatype, MPI_Op op)
```

```
int MPI_Reduce_local_c(const void *inbuf, void *inoutbuf, MPI_Count count,
                      MPI_Datatype datatype, MPI_Op op)
```

Fortran 2008 binding

```
MPI_Reduce_local(inbuf, inoutbuf, count, datatype, op, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
  TYPE(*), DIMENSION(..) :: inoutbuf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Reduce_local(inbuf, inoutbuf, count, datatype, op, ierror) !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
  TYPE(*), DIMENSION(..) :: inoutbuf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_REDUCE_LOCAL(INBUF, INOUTBUF, COUNT, DATATYPE, OP, IERROR)
  <type> INBUF(*), INOUTBUF(*)
```

INTEGER COUNT, DATATYPE, OP, IERROR

The function applies the operation given by `op` element-wise to the elements of `inbuf` and `inoutbuf` with the result stored element-wise in `inoutbuf`, as explained for user-defined operations in Section 6.9.5. Both `inbuf` and `inoutbuf` (input as well as result) have the same number of elements given by `count` and the same datatype given by `datatype`. The `MPI_IN_PLACE` option is not allowed.

Reduction operations can be queried for their commutativity.

`MPI_OP_COMMUTATIVE(op, commute)`

IN	<code>op</code>	operation (handle)
OUT	<code>commute</code>	true if <code>op</code> is commutative, false otherwise (logical)

C binding

`int MPI_Op_commutative(MPI_Op op, int *commute)`

Fortran 2008 binding

`MPI_Op_commutative(op, commute, ierror)`
 TYPE(MPI_Op), INTENT(IN) :: op
 LOGICAL, INTENT(OUT) :: commute
 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

`MPI_OP_COMMUTATIVE(OP, COMMUTE, IERROR)`
 INTEGER OP, IERROR
 LOGICAL COMMUTE

6.10 Reduce-Scatter

MPI includes variants of the reduce operations where the result is scattered to all MPI processes in a group on return. One variant scatters equal-sized blocks to all MPI processes, while another variant scatters blocks that may vary in size for each MPI process.

6.10.1 MPI_REDUCE_SCATTER_BLOCK

`MPI_REDUCE_SCATTER_BLOCK(sendbuf, recvbuf, recvcnt, datatype, op, comm)`

IN	<code>sendbuf</code>	starting address of send buffer (choice)
OUT	<code>recvbuf</code>	starting address of receive buffer (choice)
IN	<code>recvcnt</code>	element count per block (non-negative integer)
IN	<code>datatype</code>	datatype of elements of send and receive buffers (handle)
IN	<code>op</code>	operation (handle)
IN	<code>comm</code>	communicator (handle)

C binding

```
1 int MPI_Reduce_scatter_block(const void *sendbuf, void *recvbuf, int recvcount,
2                             MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
3
```

```
4 int MPI_Reduce_scatter_block_c(const void *sendbuf, void *recvbuf,
5                               MPI_Count recvcount, MPI_Datatype datatype, MPI_Op op,
6                               MPI_Comm comm)
7
```

Fortran 2008 binding

```
8 MPI_Reduce_scatter_block(sendbuf, recvbuf, recvcount, datatype, op, comm,
9                           ierror)
10
```

```
11 TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
```

```
12 TYPE(*), DIMENSION(..) :: recvbuf
```

```
13 INTEGER, INTENT(IN) :: recvcount
```

```
14 TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
15 TYPE(MPI_Op), INTENT(IN) :: op
```

```
16 TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
17 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
18 MPI_Reduce_scatter_block(sendbuf, recvbuf, recvcount, datatype, op, comm,
19                           ierror) !(_c)
```

```
20 TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
```

```
21 TYPE(*), DIMENSION(..) :: recvbuf
```

```
22 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcount
```

```
23 TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
24 TYPE(MPI_Op), INTENT(IN) :: op
```

```
25 TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
26 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
28 MPI_REDUCE_SCATTER_BLOCK(SENDBUF, RECVBUF, RECVCOUNT, DATATYPE, OP, COMM,
29                           IERROR)
```

```
30 <type> SENDBUF(*), RECVBUF(*)
```

```
31 INTEGER RECVCOUNT, DATATYPE, OP, COMM, IERROR
```

33 If `comm` is an intra-communicator, `MPI_REDUCE_SCATTER_BLOCK` first performs a
34 global, element-wise reduction on vectors of `count = n*recvcount` elements in the send buffers
35 defined by `sendbuf`, `count` and `datatype`, using the operation `op`, where `n` is the number of
36 MPI processes in the group of `comm`. The routine is called by all group members using the
37 same arguments for `recvcount`, `datatype`, `op` and `comm`. The resulting vector is treated as
38 `n` consecutive blocks of `recvcount` elements that are scattered to the MPI processes of the
39 group. The `i`-th block is sent to MPI process `i` and stored in the receive buffer defined by
40 `recvbuf`, `recvcount`, and `datatype`.

41 *Advice to implementors.* The `MPI_REDUCE_SCATTER_BLOCK` routine is function-
42 ally equivalent to an `MPI_REDUCE` collective operation with `count` equal to
43 `recvcount*n`, followed by an `MPI_SCATTER` with `sendcount` equal to `recvcount`. How-
44 ever, a direct implementation may run faster. (*End of advice to implementors.*)

46 The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` in
47 the `sendbuf` argument on *all* MPI processes. In this case, the input data is taken from the
48 receive buffer.

If `comm` is an inter-communicator, then the result of the reduction of the data provided by MPI processes in one group (group A) is scattered among MPI processes in the other group (group B) and vice versa. Within each group, all MPI processes provide the same value for the `recvcount` argument, and provide input vectors of `count = n*recvcount` elements stored in the send buffers, where `n` is the size of the group. The number of elements `count` must be the same for the two groups. The resulting vector from the other group is scattered in blocks of `recvcount` elements among the MPI processes in the group.

Rationale. The last restriction is needed so that the length of the send buffer of one group can be determined by the local `recvcount` argument of the other group. Otherwise, communication is needed to figure out how many elements are reduced. (*End of rationale.*)

6.10.2 MPI_REDUCE_SCATTER

`MPI_REDUCE_SCATTER` extends the functionality of `MPI_REDUCE_SCATTER_BLOCK` such that the scattered blocks can vary in size. Block sizes are determined by the `recvcounts` array, such that the `i`-th block contains `recvcounts[i]` elements.

`MPI_REDUCE_SCATTER(sendbuf, recvbuf, recvcounts, datatype, op, comm)`

IN	<code>sendbuf</code>	starting address of send buffer (choice)
OUT	<code>recvbuf</code>	starting address of receive buffer (choice)
IN	<code>recvcounts</code>	nonnegative integer array (of length group size) specifying the number of elements of the result distributed to each MPI process.
IN	<code>datatype</code>	datatype of elements of send and receive buffers (handle)
IN	<code>op</code>	operation (handle)
IN	<code>comm</code>	communicator (handle)

C binding

```
int MPI_Reduce_scatter(const void *sendbuf, void *recvbuf,
                      const int recvcounts[], MPI_Datatype datatype, MPI_Op op,
                      MPI_Comm comm)
```

```
int MPI_Reduce_scatter_c(const void *sendbuf, void *recvbuf,
                        const MPI_Count recvcounts[], MPI_Datatype datatype, MPI_Op op,
                        MPI_Comm comm)
```

Fortran 2008 binding

```
MPI_Reduce_scatter(sendbuf, recvbuf, recvcounts, datatype, op, comm, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
  TYPE(*), DIMENSION(..) :: recvbuf
  INTEGER, INTENT(IN) :: recvcounts(*)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
```

```

1      TYPE(MPI_Comm), INTENT(IN) :: comm
2      INTEGER, OPTIONAL, INTENT(OUT) :: ierror
3
4      MPI_Reduce_scatter(sendbuf, recvbuf, recvcunts, datatype, op, comm, ierror)
5          !(_c)
6      TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
7      TYPE(*), DIMENSION(..) :: recvbuf
8      INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcunts(*)
9      TYPE(MPI_Datatype), INTENT(IN) :: datatype
10     TYPE(MPI_Op), INTENT(IN) :: op
11     TYPE(MPI_Comm), INTENT(IN) :: comm
12     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

13     MPI_REDUCE_SCATTER(SENDBUF, RECVBUF, RECVCOUNTS, DATATYPE, OP, COMM, IERROR)
14     <type> SENDBUF(*), RECVBUF(*)
15     INTEGER REVCOUNTS(*), DATATYPE, OP, COMM, IERROR

```

If `comm` is an intra-communicator, `MPI_REDUCE_SCATTER` first performs a global, element-wise reduction on vectors of $\text{count} = \sum_{i=0}^{n-1} \text{recvcunts}[i]$ elements in the send buffers defined by `sendbuf`, `count` and `datatype`, using the operation `op`, where `n` is the number of MPI processes in the group of `comm`. The routine is called by all group members using the same arguments for `recvcunts`, `datatype`, `op` and `comm`. The resulting vector is treated as `n` consecutive blocks where the number of elements of the `i`-th block is `recvcunts[i]`. The blocks are scattered to the MPI processes of the group. The `i`-th block is sent to MPI process `i` and stored in the receive buffer defined by `recvbuf`, `recvcunts[i]` and `datatype`.

Advice to implementors. The `MPI_REDUCE_SCATTER` routine is functionally equivalent to an `MPI_REDUCE` collective operation with `count` equal to the sum of `recvcunts[i]` followed by `MPI_SCATTERV` with `sendcounts` equal to `recvcunts`. However, a direct implementation may run faster. (*End of advice to implementors.*)

The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` in the `sendbuf` argument. In this case, the input data is taken from the receive buffer. It is not required to specify the “in place” option on all MPI processes, since the MPI processes for which `recvcunts[i] = 0` may not have allocated a receive buffer.

If `comm` is an inter-communicator, then the result of the reduction of the data provided by MPI processes in one group (group A) is scattered among MPI processes in the other group (group B), and vice versa. Within each group, all MPI processes provide the same `recvcunts` argument, and provide input vectors of $\text{count} = \sum_{i=0}^{n-1} \text{recvcunts}[i]$ elements stored in the send buffers, where `n` is the size of the group. The resulting vector from the other group is scattered in blocks of `recvcunts[i]` elements among the MPI processes in the group. The number of elements `count` must be the same for the two groups.

Rationale. The last restriction is needed so that the length of the send buffer can be determined by the sum of the local `recvcunts` entries. Otherwise, communication is needed to figure out how many elements are reduced. (*End of rationale.*)

6.11 Scan

6.11.1 Inclusive Scan

`MPI_SCAN(sendbuf, recvbuf, count, datatype, op, comm)`

IN	sendbuf	starting address of send buffer (choice)
OUT	recvbuf	starting address of receive buffer (choice)
IN	count	number of elements in input buffer (non-negative integer)
IN	datatype	datatype of elements of input buffer (handle)
IN	op	operation (handle)
IN	comm	communicator (handle)

C binding

```
int MPI_Scan(const void *sendbuf, void *recvbuf, int count,
            MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
```

```
int MPI_Scan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
              MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
```

Fortran 2008 binding

```
MPI_Scan(sendbuf, recvbuf, count, datatype, op, comm, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
```

```
TYPE(*), DIMENSION(..) :: recvbuf
```

```
INTEGER, INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Op), INTENT(IN) :: op
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Scan(sendbuf, recvbuf, count, datatype, op, comm, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
```

```
TYPE(*), DIMENSION(..) :: recvbuf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Op), INTENT(IN) :: op
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
```

```
<type> SENDBUF(*), RECVBUF(*)
```

```
INTEGER COUNT, DATATYPE, OP, COMM, IERROR
```

If `comm` is an intra-communicator, `MPI_SCAN` is used to perform a prefix reduction on data distributed across the group. The operation returns, in the receive buffer of the MPI process with rank `i`, the reduction of the values in the send buffers of MPI processes

with ranks 0,...,i (inclusive). The routine is called by all group members using the same arguments for count, datatype, op and comm, except that for user-defined operations, the same rules apply as for MPI_REDUCE. The type of operations supported, their semantics, and the constraints on send and receive buffers are as for MPI_REDUCE.

The “in place” option for intra-communicators is specified by passing MPI_IN_PLACE in the sendbuf argument. In this case, the input data is taken from the receive buffer, and replaced by the output data.

This operation is invalid for inter-communicators.

6.11.2 Exclusive Scan

MPI_EXSCAN(sendbuf, recvbuf, count, datatype, op, comm)

IN	sendbuf	starting address of send buffer (choice)
OUT	recvbuf	starting address of receive buffer (choice)
IN	count	number of elements in input buffer (non-negative integer)
IN	datatype	datatype of elements of input buffer (handle)
IN	op	operation (handle)
IN	comm	intra-communicator (handle)

C binding

```
int MPI_Exscan(const void *sendbuf, void *recvbuf, int count,
               MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
```

```
int MPI_Exscan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
                 MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
```

Fortran 2008 binding

```
MPI_Exscan(sendbuf, recvbuf, count, datatype, op, comm, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
```

```
TYPE(*), DIMENSION(..) :: recvbuf
```

```
INTEGER, INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Op), INTENT(IN) :: op
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Exscan(sendbuf, recvbuf, count, datatype, op, comm, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
```

```
TYPE(*), DIMENSION(..) :: recvbuf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Op), INTENT(IN) :: op
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_EXSCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
```

```
<type> SENDBUF(*), RECVBUF(*)
```

```
INTEGER COUNT, DATATYPE, OP, COMM, IERROR
```

If `comm` is an intra-communicator, `MPI_EXSCAN` is used to perform a prefix reduction on data distributed across the group. The value in `recvbuf` on the MPI process with rank 0 is undefined, and `recvbuf` is not significant on that MPI process. The value in `recvbuf` on the MPI process with rank 1 is defined as the value in `sendbuf` on the MPI process with rank 0. For MPI processes with rank $i > 1$, the operation returns, in the receive buffer of the MPI process with rank i , the reduction of the values in the send buffers of MPI processes with ranks $0, \dots, i - 1$ (inclusive). The routine is called by all group members using the same arguments for `count`, `datatype`, `op` and `comm`, except that for user-defined operations, the same rules apply as for `MPI_REDUCE`. The type of operations supported, their semantics, and the constraints on send and receive buffers, are as for `MPI_REDUCE`.

The “in place” option for intra-communicators is specified by passing `MPI_IN_PLACE` in the `sendbuf` argument. In this case, the input data is taken from the receive buffer, and replaced by the output data. The receive buffer on rank 0 is not changed by this operation.

This operation is invalid for inter-communicators.

Rationale. The exclusive scan is more general than the inclusive scan. Any inclusive scan operation can be achieved by using the exclusive scan and then locally combining the local contribution. Note that for noninvertable operations such as `MPI_MAX`, the exclusive scan cannot be computed with the inclusive scan. (*End of rationale.*)

6.11.3 Example using `MPI_SCAN`

The example in this section uses an intra-communicator.

Example 6.24. This example uses a user-defined operation to produce a **segmented scan**. A segmented scan takes, as input, a set of values and a set of logicals, and the logicals delineate the various segments of the scan. For example:

<i>values</i>	v_1	v_2	v_3	v_4	v_5	v_6	v_7	v_8
<i>logicals</i>	0	0	1	1	1	0	0	1
<i>result</i>	v_1	$v_1 + v_2$	v_3	$v_3 + v_4$	$v_3 + v_4 + v_5$	v_6	$v_6 + v_7$	v_8

The operator that produces this effect is

$$\begin{pmatrix} u \\ i \end{pmatrix} \circ \begin{pmatrix} v \\ j \end{pmatrix} = \begin{pmatrix} w \\ j \end{pmatrix},$$

where

$$w = \begin{cases} u + v & \text{if } i = j \\ v & \text{if } i \neq j \end{cases}.$$

Note that this is a noncommutative operator. C code that implements it is given below.

```
typedef struct {
    double val;
    int log;
```

```

1  } SegScanPair;
2
3  /* the user-defined function
4  */
5  void segScan(SegScanPair *in, SegScanPair *inout, int *len,
6              MPI_Datatype *dptr)
7  {
8      int i;
9      SegScanPair c;
10
11     for (i=0; i< *len; ++i) {
12         if (in->log == inout->log)
13             c.val = in->val + inout->val;
14         else
15             c.val = inout->val;
16         c.log = inout->log;
17         *inout = c;
18         in++; inout++;
19     }
20 }

```

Note that the `inout` argument to the user-defined function corresponds to the right-hand operand of the operator. When using this operator, we must be careful to specify that it is noncommutative, as in the following.

```

23 int i, base;
24 SegScanPair a, answer;
25 MPI_Op myOp;
26 MPI_Datatype type[2] = {MPI_DOUBLE, MPI_INT};
27 MPI_Aint disp[2];
28 int blocklen[2] = {1, 1};
29 MPI_Datatype sspair;
30
31 /* explain to MPI how type SegScanPair is defined
32 */
33 MPI_Get_address(&a, disp);
34 MPI_Get_address(&a.log, disp+1);
35 base = disp[0];
36 for (i=0; i<2; ++i) disp[i] -= base;
37 MPI_Type_create_struct(2, blocklen, disp, type, &sspair);
38 MPI_Type_commit(&sspair);
39 /* create the segmented-scan user-op
40 */
41 MPI_Op_create(segScan, 0, &myOp);
42 ...
43 MPI_Scan(&a, &answer, 1, sspair, myOp, comm);

```

6.12 Nonblocking Collective Operations

As described in Section 3.7, performance of many applications can be improved by overlapping communication and computation, and many systems enable this. Nonblocking collective operations combine the potential benefits of nonblocking point-to-point opera-

tions, to exploit overlap and to avoid synchronization, with the optimized implementation and message scheduling provided by collective operations [35, 39]. One way of doing this would be to perform a blocking collective operation in a separate thread. An alternative mechanism that often leads to better performance (e.g., avoids context switching, scheduler overheads, and thread management) is to use nonblocking collective communication [37].

The nonblocking collective communication model is similar to the model used for nonblocking point-to-point communication. A nonblocking call initiates a collective operation, which must be completed in a separate completion call. Once initiated, the operation may progress independently of any computation or other communication at participating MPI processes. In this manner, nonblocking collective operations can mitigate possible synchronizing effects of collective operations by running them in the “background.” In addition to enabling communication-computation overlap, nonblocking collective operations can perform collective operations on overlapping communicators, which would lead to deadlocks with blocking operations. Their semantic advantages can also be useful in combination with point-to-point communication.

As in the nonblocking point-to-point case, all calls are local and return immediately, irrespective of the status of other MPI processes. The call initiates the operation, which indicates that the system may start to copy data out of the send buffer and into the receive buffer. Once initiated, all associated send buffers and buffers associated with input arguments (such as arrays of counts, displacements, or datatypes in the vector versions of the collectives) should not be modified, and all associated receive buffers should not be accessed, until the collective operation completes. The call returns a request handle, which must be passed to a completion call.

All completion calls (e.g., `MPI_WAIT`) described in Section 3.7.3 are supported for nonblocking collective operations. Similarly to the blocking case, nonblocking collective operations are considered to be complete when the local part of the operation is finished, i.e., for the caller, the semantics of the operation are guaranteed and all buffers can be safely accessed and modified. Completion does not indicate that other MPI processes have completed or even started the operation (unless otherwise implied by the description of the operation). Completion of a particular nonblocking collective operation also does not indicate completion of any other posted nonblocking collective (or send-receive) operations, whether they are posted before or after the completed operation.

Advice to users. Users should be aware that implementations are allowed, but not required (with exception of `MPI_IBARRIER`), to synchronize MPI processes during the completion of a nonblocking collective operation. (*End of advice to users.*)

Upon returning from a completion call in which a nonblocking collective operation completes, the values of the `MPI_SOURCE` and `MPI_TAG` fields in the associated status object, if any, are undefined. The value of `MPI_ERROR` may be defined, if appropriate, according to the specification in Section 3.2.5. It is valid to mix different request types (i.e., any combination of collective requests, I/O requests, generalized requests, or point-to-point requests) in functions that enable multiple completions (e.g., `MPI_WAITALL`). It is erroneous to call `MPI_REQUEST_FREE` or `MPI_CANCEL` for a request associated with a nonblocking collective operation. Nonblocking collective requests created using the APIs described in this section are not persistent. However, persistent collective requests can be created using persistent collective operations described in Sections 6.13 and 8.8.

Rationale. Freeing an active nonblocking collective request could cause similar

1 problems as discussed for point-to-point requests (see Section 3.7.3). Cancelling a
2 request is not supported because the semantics of this operation are not well-defined.
3 (*End of rationale.*)
4

5 Multiple nonblocking collective operations can be outstanding on a single communica-
6 tor. If the nonblocking call causes some system resource to be exhausted, then it will fail
7 and raise an error. Quality implementations of MPI should ensure that this happens only
8 in pathological cases. That is, an MPI implementation should be able to support a large
9 number of *pending* nonblocking operations.

10 Unlike point-to-point operations, nonblocking collective operations do not match with
11 blocking collective operations, and collective operations do not have a tag argument. All
12 MPI processes must call collective operations (blocking and nonblocking) in the same order
13 per communicator. In particular, once a MPI process calls a collective operation, all other
14 MPI processes in the communicator must eventually call the same collective operation, and
15 no other collective operation with the same communicator in between. This is consistent
16 with the ordering rules for blocking collective operations in threaded environments.
17

18 *Rationale.* Matching blocking and nonblocking collective operations is not allowed
19 because the implementation might use different communication algorithms for the two
20 cases. Blocking collective operations may be optimized for minimal time to comple-
21 tion, while nonblocking collective operations may balance time to completion with
22 CPU overhead and asynchronous progress.

23 The use of tags for collective operations can prevent certain hardware optimizations.
24 (*End of rationale.*)
25

26 *Advice to users.* If program semantics require matching blocking and nonblocking
27 collective operations, then a nonblocking collective operation can be initiated and
28 immediately completed with a blocking wait to emulate blocking behavior. (*End of*
29 *advice to users.*)
30

31 In terms of data movement, each nonblocking collective operation has the same effect as
32 its blocking counterpart for intra-communicators and inter-communicators after completion.
33 Likewise, upon completion, nonblocking collective reduction operations have the same effect
34 as their blocking counterparts, and the same restrictions and recommendations on reduction
35 orders apply.

36 The use of the “in place” option is allowed exactly as described for the corresponding
37 blocking collective operations. When using the “in place” option, message buffers function
38 as both send and receive buffers. Such buffers should not be modified or accessed until the
39 operation completes.

40 The *progress* rules for nonblocking collective operations are similar to the progress rules
41 for nonblocking point-to-point operations, refer to Sections 2.9 and 3.7.4.
42

43 *Advice to implementors.* Nonblocking collective operations can be implemented with
44 local execution schedules [38] using nonblocking point-to-point communication and a
45 reserved tag-space. (*End of advice to implementors.*)
46
47
48

6.12.1 Nonblocking Barrier Synchronization

MPI_IBARRIER(comm, request)

IN	comm	communicator (handle)
OUT	request	communication request (handle)

C binding

```
int MPI_Ibarrier(MPI_Comm comm, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Ibarrier(comm, request, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_IBARRIER(COMM, REQUEST, IERROR)
  INTEGER COMM, REQUEST, IERROR
```

MPI_IBARRIER is a nonblocking version of MPI_BARRIER. By calling MPI_IBARRIER, an MPI process notifies that it has reached the barrier. The call returns immediately, independent of whether other MPI processes have called MPI_IBARRIER. The usual barrier semantics are enforced at the corresponding completion operation (test or wait), which in the intra-communicator case will complete only after all other MPI processes in the communicator have called MPI_IBARRIER. In the inter-communicator case, it will complete when all MPI processes in the remote group have called MPI_IBARRIER.

Advice to users. A nonblocking barrier can be used to hide latency. Moving independent computations between the MPI_IBARRIER and the subsequent completion call can overlap the barrier latency and therefore shorten possible waiting times. The semantic properties are also useful when mixing collective operations and point-to-point messages. (*End of advice to users.*)

6.12.2 Nonblocking Broadcast

MPI_IBCAST(buffer, count, datatype, root, comm, request)

INOUT	buffer	starting address of buffer (choice)
IN	count	number of entries in buffer (non-negative integer)
IN	datatype	datatype of buffer (handle)
IN	root	rank of broadcast root (integer)
IN	comm	communicator (handle)
OUT	request	communication request (handle)

C binding

```
int MPI_Ibcast(void *buffer, int count, MPI_Datatype datatype, int root,
```

```

1         MPI_Comm comm, MPI_Request *request)
2
3 int MPI_Ibcast_c(void *buffer, MPI_Count count, MPI_Datatype datatype,
4                 int root, MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```

5 MPI_Ibcast(buffer, count, datatype, root, comm, request, ierror)
6     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
7     INTEGER, INTENT(IN) :: count, root
8     TYPE(MPI_Datatype), INTENT(IN) :: datatype
9     TYPE(MPI_Comm), INTENT(IN) :: comm
10    TYPE(MPI_Request), INTENT(OUT) :: request
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_Ibcast(buffer, count, datatype, root, comm, request, ierror) !(_c)
14    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
15    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
16    TYPE(MPI_Datatype), INTENT(IN) :: datatype
17    INTEGER, INTENT(IN) :: root
18    TYPE(MPI_Comm), INTENT(IN) :: comm
19    TYPE(MPI_Request), INTENT(OUT) :: request
20    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21

```

Fortran binding

```

22 MPI_IBCAST(BUFFER, COUNT, DATATYPE, ROOT, COMM, REQUEST, IERROR)
23     <type> BUFFER(*)
24     INTEGER COUNT, DATATYPE, ROOT, COMM, REQUEST, IERROR

```

This call starts a nonblocking variant of MPI_BCAST (see Section 6.4).

Example using MPI_IBCAST

The example in this section uses an intra-communicator.

Example 6.25. Start a broadcast of 100 ints from MPI process 0 to every MPI process in the group, perform some computation on independent data, and then complete the outstanding broadcast operation.

```

35 MPI_Comm comm;
36 int array1[100], array2[100];
37 int root=0;
38 MPI_Request req;
39 ...
40 MPI_Ibcast(array1, 100, MPI_INT, root, comm, &req);
41 compute(array2, 100);
42 MPI_Wait(&req, MPI_STATUS_IGNORE);

```


6.12.3 Nonblocking Gather

			1
			2
			3
			4
			5
			6
IN	sendbuf	starting address of send buffer (choice)	7
IN	sendcount	number of elements in send buffer (non-negative integer)	8
			9
IN	sendtype	datatype of send buffer elements (handle)	10
			11
OUT	recvbuf	address of receive buffer (choice, significant only at root)	12
			13
IN	recvcount	number of elements for any single receive (non-negative integer, significant only at root)	14
			15
IN	recvtype	datatype of recv buffer elements (handle, significant only at root)	16
			17
IN	root	rank of receiving MPI process (integer)	18
			19
IN	comm	communicator (handle)	20
			21
OUT	request	communication request (handle)	22
			23

C binding

```
int MPI_Igather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
               void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
               MPI_Comm comm, MPI_Request *request)
```

```
int MPI_Igather_c(const void *sendbuf, MPI_Count sendcount,
                 MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                 MPI_Datatype recvtype, int root, MPI_Comm comm,
                 MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Igather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
            comm, request, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount, root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Igather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
            comm, request, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```

1     INTEGER, INTENT(IN) :: root
2     TYPE(MPI_Comm), INTENT(IN) :: comm
3     TYPE(MPI_Request), INTENT(OUT) :: request
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

6     MPI_IGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, ROOT,
7                COMM, REQUEST, IERROR)
8     <type> SENDBUF(*), RECVBUF(*)
9     INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, ROOT, COMM, REQUEST,
10            IERROR

```

This call starts a nonblocking variant of MPI_GATHER (see Section 6.5).

```

14    MPI_IGATHERV(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs, recvtype, root,
15                comm, request)

```

17	IN	sendbuf	starting address of send buffer (choice)
18	IN	sendcount	number of elements in send buffer (non-negative integer)
19			
20	IN	sendtype	datatype of send buffer elements (handle)
21			
22	OUT	recvbuf	address of receive buffer (choice, significant only at root)
23			
24	IN	recvcnts	nonnegative integer array (of length group size) containing the number of elements that are received from each MPI process (significant only at root)
25			
26			
27	IN	displs	integer array (of length group size). Entry <i>i</i> specifies the displacement relative to <i>recvbuf</i> at which to place the incoming data from MPI process <i>i</i> (significant only at root)
28			
29			
30			
31			
32	IN	recvtype	datatype of recv buffer elements (handle, significant only at root)
33			
34	IN	root	rank of receiving MPI process (integer)
35	IN	comm	communicator (handle)
36	OUT	request	communication request (handle)
37			

C binding

```

39    int MPI_Igatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
40                   void *recvbuf, const int recvcnts[], const int displs[],
41                   MPI_Datatype recvtype, int root, MPI_Comm comm,
42                   MPI_Request *request)

```

```

44    int MPI_Igatherv_c(const void *sendbuf, MPI_Count sendcount,
45                      MPI_Datatype sendtype, void *recvbuf,
46                      const MPI_Count recvcnts[], const MPI_Aint displs[],
47                      MPI_Datatype recvtype, int root, MPI_Comm comm,
48                      MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Igatherv(sendbuf, sendcount, sendtype, recvbuf, recvcoun1ts, displs,
              recvtype, root, comm, request, ierror) 2
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 3
INTEGER, INTENT(IN) :: sendcount, root 4
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 5
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 6
INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcoun7ts(*), displs(*) 8
TYPE(MPI_Comm), INTENT(IN) :: comm 9
TYPE(MPI_Request), INTENT(OUT) :: request 10
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 11

MPI_Igatherv(sendbuf, sendcount, sendtype, recvbuf, recvcoun12ts, displs,
              recvtype, root, comm, request, ierror) !(_c) 13
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 14
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount 15
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 16
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 17
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcoun18ts(*) 18
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*) 19
INTEGER, INTENT(IN) :: root 20
TYPE(MPI_Comm), INTENT(IN) :: comm 21
TYPE(MPI_Request), INTENT(OUT) :: request 22
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 23

```

Fortran binding

```

MPI_IGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
              RECVTYPE, ROOT, COMM, REQUEST, IERROR) 25
<type> SENDBUF(*), RECVBUF(*) 26
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT, 27
COMM, REQUEST, IERROR 28

```

This call starts a nonblocking variant of MPI_GATHERV (see Section 6.5).

6.12.4 Nonblocking Scatter

```

MPI_ISCATTER(sendbuf, sendcount, sendtype, recvbuf, recvcoun36t, recvtype, root, comm,
              request) 37

```

IN	sendbuf	address of send buffer (choice, significant only at root)	³⁹
IN	sendcount	number of elements sent to each MPI process (non-negative integer, significant only at root)	⁴⁰
IN	sendtype	datatype of send buffer elements (handle, significant only at root)	⁴¹
OUT	recvbuf	address of receive buffer (choice)	⁴²
IN	recvcoun ⁴³ t	number of elements in receive buffer (non-negative integer)	⁴⁴

```

1      IN      recvtype      datatype of receive buffer elements (handle)
2      IN      root         rank of sending MPI process (integer)
3
4      IN      comm         communicator (handle)
5      OUT     request      communication request (handle)
6

```

C binding

```

8      int MPI_Iscatter(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
9                    void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
10                   MPI_Comm comm, MPI_Request *request)
11

```

```

12     int MPI_Iscatter_c(const void *sendbuf, MPI_Count sendcount,
13                     MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
14                     MPI_Datatype recvtype, int root, MPI_Comm comm,
15                     MPI_Request *request)
16

```

Fortran 2008 binding

```

17     MPI_Iscatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
18                comm, request, ierror)
19     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
20     INTEGER, INTENT(IN) :: sendcount, recvcount, root
21     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
22     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
23     TYPE(MPI_Comm), INTENT(IN) :: comm
24     TYPE(MPI_Request), INTENT(OUT) :: request
25     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26

```

```

27     MPI_Iscatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
28                comm, request, ierror) !(_c)
29     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
30     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
31     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
32     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
33     INTEGER, INTENT(IN) :: root
34     TYPE(MPI_Comm), INTENT(IN) :: comm
35     TYPE(MPI_Request), INTENT(OUT) :: request
36     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37

```

Fortran binding

```

38     MPI_ISCATTER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, ROOT,
39                COMM, REQUEST, IERROR)
40     <type> SENDBUF(*), RECVBUF(*)
41     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, REQUEST,
42                IERROR
43

```

44 This call starts a nonblocking variant of MPI_SCATTER (see Section 6.6).

```

45
46
47
48

```

MPI_ISCATTERV(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount, recvtype, root, comm, request)	1
	2
IN sendbuf	3
address of send buffer (choice, significant only at root)	4
	5
IN sendcounts	6
nonnegative integer array (of length group size) specifying the number of elements to send to each MPI process (significant only at root)	7
	8
IN displs	9
integer array (of length group size). Entry i specifies the displacement (relative to sendbuf) from which to take the outgoing data to MPI process i (significant only at root)	10
	11
IN sendtype	12
datatype of send buffer elements (handle, significant only at root)	13
	14
OUT recvbuf	15
address of receive buffer (choice)	16
IN recvcount	17
number of elements in receive buffer (non-negative integer)	18
	19
IN recvtype	20
datatype of receive buffer elements (handle)	21
IN root	22
rank of sending MPI process (integer)	23
IN comm	24
communicator (handle)	25
OUT request	26
communication request (handle)	27
	28
C binding	29
int MPI_Iscatterv(const void *sendbuf, const int sendcounts[],	30
const int displs[], MPI_Datatype sendtype, void *recvbuf,	31
int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm,	32
MPI_Request *request)	33
	34
int MPI_Iscatterv_c(const void *sendbuf, const MPI_Count sendcounts[],	35
const MPI_Aint displs[], MPI_Datatype sendtype, void *recvbuf,	36
MPI_Count recvcount, MPI_Datatype recvtype, int root,	37
MPI_Comm comm, MPI_Request *request)	38
	39
Fortran 2008 binding	40
MPI_Iscatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,	41
recvtype, root, comm, request, ierror)	42
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf	43
INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), displs(*)	44
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype	45
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf	46
INTEGER, INTENT(IN) :: recvcount, root	47
TYPE(MPI_Comm), INTENT(IN) :: comm	48
TYPE(MPI_Request), INTENT(OUT) :: request	
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	
MPI_Iscatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,	
recvtype, root, comm, request, ierror) !(_c)	

```

1     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
2     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*)
3     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
4     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
5     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
6     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcnt
7     INTEGER, INTENT(IN) :: root
8     TYPE(MPI_Comm), INTENT(IN) :: comm
9     TYPE(MPI_Request), INTENT(OUT) :: request
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

12    MPI_ISCATTERV(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,
13                RECVTYPE, ROOT, COMM, REQUEST, IERROR)
14    <type> SENDBUF(*), RECVBUF(*)
15    INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,
16                COMM, REQUEST, IERROR

```

18 This call starts a nonblocking variant of MPI_SCATTERV (see Section 6.6).

6.12.5 Nonblocking Gather-to-all

```

23    MPI_IALLGATHER(sendbuf, sendcount, sendtype, recvbuf, recvcnt, recvtype, comm,
24                  request)

```

26	IN	sendbuf	starting address of send buffer (choice)
27	IN	sendcount	number of elements in send buffer (non-negative integer)
29	IN	sendtype	datatype of send buffer elements (handle)
30	OUT	recvbuf	address of receive buffer (choice)
31	IN	recvcnt	number of elements received from any MPI process (non-negative integer)
32	IN	recvtype	datatype of receive buffer elements (handle)
33	IN	comm	communicator (handle)
34	OUT	request	communication request (handle)

C binding

```

40    int MPI_Iallgather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
41                    void *recvbuf, int recvcnt, MPI_Datatype recvtype,
42                    MPI_Comm comm, MPI_Request *request)
43
44    int MPI_Iallgather_c(const void *sendbuf, MPI_Count sendcount,
45                      MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcnt,
46                      MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```
MPI_Iallgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
               comm, request, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER, INTENT(IN) :: sendcount, recvcount
```

```
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Iallgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
               comm, request, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
```

```
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_IALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
               COMM, REQUEST, IERROR)
```

```
<type> SENDBUF(*), RECVBUF(*)
```

```
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, REQUEST, IERROR
```

This call starts a nonblocking variant of MPI_ALLGATHER (see Section 6.7).

```
MPI_IALLGATHERV(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs, recvtype,
                comm, request)
```

IN	sendbuf	starting address of send buffer (choice)
IN	sendcount	number of elements in send buffer (non-negative integer)
IN	sendtype	datatype of send buffer elements (handle)
OUT	recvbuf	address of receive buffer (choice)
IN	recvcnts	nonnegative integer array (of length group size) containing the number of elements that are received from each MPI process
IN	displs	integer array (of length group size). Entry <i>i</i> specifies the displacement (relative to <i>recvbuf</i>) at which to place the incoming data from MPI process <i>i</i>
IN	recvtype	datatype of receive buffer elements (handle)
IN	comm	communicator (handle)
OUT	request	communication request (handle)

C binding

```

1 int MPI_Iallgatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
2                   void *recvbuf, const int recvcounts[], const int displs[],
3                   MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
4
5
6
7
8
9

```

```

10 int MPI_Iallgatherv_c(const void *sendbuf, MPI_Count sendcount,
11                    MPI_Datatype sendtype, void *recvbuf,
12                    const MPI_Count recvcounts[], const MPI_Aint displs[],
13                    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
14
15
16
17
18
19
20
21

```

Fortran 2008 binding

```

22 MPI_Iallgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
23               recvtype, comm, request, ierror)
24   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
25   INTEGER, INTENT(IN) :: sendcount
26   TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
27   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
28   INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*), displs(*)
29   TYPE(MPI_Comm), INTENT(IN) :: comm
30   TYPE(MPI_Request), INTENT(OUT) :: request
31   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33
34
35
36
37
38
39

```

```

40 MPI_Iallgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
41               recvtype, comm, request, ierror) !(_c)
42   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
43   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
44   TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
45   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
46   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
47   INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
48   TYPE(MPI_Comm), INTENT(IN) :: comm
49   TYPE(MPI_Request), INTENT(OUT) :: request
50   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
51
52
53
54
55
56
57
58

```

Fortran binding

```

59 MPI_IALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
60               RECVTYPE, COMM, REQUEST, IERROR)
61   <type> SENDBUF(*), RECVBUF(*)
62   INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
63   REQUEST, IERROR
64
65
66
67
68
69
70
71

```

This call starts a nonblocking variant of MPI_ALLGATHERV (see Section 6.7).

6.12.6 Nonblocking All-to-All Scatter/Gather

```

72 MPI_IALLTOALL(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm, request)
73
74
75
76
77
78
79
80

```

IN **sendbuf** starting address of send buffer (choice)

IN	sendcount	number of elements sent to each MPI process (non-negative integer)	1 2
IN	sendtype	datatype of send buffer elements (handle)	3 4
OUT	recvbuf	address of receive buffer (choice)	5
IN	recvcount	number of elements received from any MPI process (non-negative integer)	6 7
IN	recvtype	datatype of receive buffer elements (handle)	8 9
IN	comm	communicator (handle)	10
OUT	request	communication request (handle)	11 12

C binding

```
int MPI_Ialltoall(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
                 void *recvbuf, int recvcount, MPI_Datatype recvtype,
                 MPI_Comm comm, MPI_Request *request)
```

```
int MPI_Ialltoall_c(const void *sendbuf, MPI_Count sendcount,
                   MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                   MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Ialltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
              request, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER, INTENT(IN) :: sendcount, recvcount
```

```
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Ialltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
              request, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
```

```
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_IALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE, COMM,
              REQUEST, IERROR)
```

```
<type> SENDBUF(*), RECVBUF(*)
```

```
INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, REQUEST, IERROR
```

This call starts a nonblocking variant of MPI_ALLTOALL (see Section 6.8).

```

1 MPI_IALLTOALLV(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts, rdispls,
2               recvtype, comm, request)
3
4   IN      sendbuf      starting address of send buffer (choice)
5   IN      sendcounts   nonnegative integer array (of length group size)
6                       specifying the number of elements to send to each
7                       MPI process
8   IN      sdispls      integer array (of length group size). Entry j specifies
9                       the displacement (relative to sendbuf) from which to
10                      take the outgoing data destined for MPI process j
11
12  IN      sendtype      datatype of send buffer elements (handle)
13  OUT     recvbuf       address of receive buffer (choice)
14  IN      recvcounts    nonnegative integer array (of length group size)
15                      specifying the number of elements that can be
16                      received from each MPI process
17
18  IN      rdispls       integer array (of length group size). Entry i specifies
19                      the displacement (relative to recvbuf) at which to
20                      place the incoming data from MPI process i
21
22  IN      recvtype      datatype of receive buffer elements (handle)
23  IN      comm          communicator (handle)
24  OUT     request       communication request (handle)

```

C binding

```

26 int MPI_Ialltoallv(const void *sendbuf, const int sendcounts[],
27                  const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
28                  const int recvcounts[], const int rdispls[],
29                  MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
30
31 int MPI_Ialltoallv_c(const void *sendbuf, const MPI_Count sendcounts[],
32                    const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf,
33                    const MPI_Count recvcounts[], const MPI_Aint rdispls[],
34                    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```

36 MPI_Ialltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts,
37               rdispls, recvtype, comm, request, ierror)
38   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
39   INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
40   recvcounts(*), rdispls(*)
41   TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
42   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
43   TYPE(MPI_Comm), INTENT(IN) :: comm
44   TYPE(MPI_Request), INTENT(OUT) :: request
45   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
46
47 MPI_Ialltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts,
48               rdispls, recvtype, comm, request, ierror) !(_c)

```

```

TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf           1
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),  2
    recvcounts(*)                                                    3
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),  4
    rdispls(*)                                                        5
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype                6
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf                    7
TYPE(MPI_Comm), INTENT(IN) :: comm                                  8
TYPE(MPI_Request), INTENT(OUT) :: request                          9
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                           10

```

Fortran binding

```

MPI_IALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, REVCOUNTS,  12
    RDISPLS, RECVTYPE, COMM, REQUEST, IERROR)                          13
<type> SENDBUF(*), RECVBUF(*)                                         14
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, REVCOUNTS(*), RDISPLS(*),  15
    RECVTYPE, COMM, REQUEST, IERROR                                    16

```

This call starts a nonblocking variant of `MPI_ALLTOALLV` (see Section 6.8).

```

MPI_IALLTOALLW(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcounts, rdispls,  21
    recvtypes, comm, request)                                          22

```

IN	sendbuf	starting address of send buffer (choice)	23
IN	sendcounts	integer array (of length group size) specifying the number of elements to send to each MPI process (array of non-negative integers)	24 25 26 27
IN	sdispls	integer array (of length group size). Entry j specifies the displacement in bytes (relative to <code>sendbuf</code>) from which to take the outgoing data destined for MPI process j (array of integers)	28 29 30 31
IN	sendtypes	array of datatypes (of length group size). Entry j specifies the type of data to send to MPI process j (array of handles)	32 33 34
OUT	recvbuf	address of receive buffer (choice)	35 36
IN	recvcounts	integer array (of length group size) specifying the number of elements that can be received from each MPI process (array of non-negative integers)	37 38 39
IN	rdispls	integer array (of length group size). Entry i specifies the displacement in bytes (relative to <code>recvbuf</code>) at which to place the incoming data from MPI process i (array of integers)	40 41 42 43
IN	recvtypes	array of datatypes (of length group size). Entry i specifies the type of data received from MPI process i (array of handles)	44 45 46
IN	comm	communicator (handle)	47 48

```
1     OUT     request                communication request (handle)
```

```
2
```

```
3 C binding
```

```
4 int MPI_Ialltoallw(const void *sendbuf, const int sendcounts[],
5                   const int sdispls[], const MPI_Datatype sendtypes[],
6                   void *recvbuf, const int recvcounts[], const int rdispls[],
7                   const MPI_Datatype recvtypes[], MPI_Comm comm,
8                   MPI_Request *request)
```

```
9
```

```
10 int MPI_Ialltoallw_c(const void *sendbuf, const MPI_Count sendcounts[],
11                    const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
12                    void *recvbuf, const MPI_Count recvcounts[],
13                    const MPI_Aint rdispls[], const MPI_Datatype recvtypes[],
14                    MPI_Comm comm, MPI_Request *request)
```

```
15 Fortran 2008 binding
```

```
16 MPI_Ialltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcounts,
17               rdispls, recvtypes, comm, request, ierror)
18     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
19     INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
20               recvcounts(*), rdispls(*)
21     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
22     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
23     TYPE(MPI_Comm), INTENT(IN) :: comm
24     TYPE(MPI_Request), INTENT(OUT) :: request
25     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
26
```

```
27 MPI_Ialltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcounts,
28               rdispls, recvtypes, comm, request, ierror) !(_c)
29     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
30     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
31               recvcounts(*)
32     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
33               rdispls(*)
34     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
35     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
36     TYPE(MPI_Comm), INTENT(IN) :: comm
37     TYPE(MPI_Request), INTENT(OUT) :: request
38     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
39 Fortran binding
```

```
40 MPI_IALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF, REVCOUNTS,
41               RDISPLS, RECVTYPES, COMM, REQUEST, IERROR)
42     <type> SENDBUF(*), RECVBUF(*)
43     INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPES(*), REVCOUNTS(*), RDISPLS(*),
44               RECVTYPES(*), COMM, REQUEST, IERROR
```

```
45
```

```
46     This call starts a nonblocking variant of MPI_ALLTOALLW (see Section 6.8).
```

```
47
```

```
48
```

6.12.7 Nonblocking Reduce

MPI_IREDUCE(sendbuf, recvbuf, count, datatype, op, root, comm, request)

IN	sendbuf	address of send buffer (choice)	1
OUT	recvbuf	address of receive buffer (choice, significant only at root)	2
IN	count	number of elements in send buffer (non-negative integer)	3
IN	datatype	datatype of elements of send buffer (handle)	4
IN	op	reduce operation (handle)	5
IN	root	rank of the root (integer)	6
IN	comm	communicator (handle)	7
OUT	request	communication request (handle)	8

C binding

```
int MPI_Ireduce(const void *sendbuf, void *recvbuf, int count,
               MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
               MPI_Request *request)
```

```
int MPI_Ireduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,
                  MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
                  MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Ireduce(sendbuf, recvbuf, count, datatype, op, root, comm, request, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
  INTEGER, INTENT(IN) :: count, root
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Ireduce(sendbuf, recvbuf, count, datatype, op, root, comm, request, ierror)
  !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
  INTEGER, INTENT(IN) :: root
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```

MPI_IREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, REQUEST, IERROR)
  <type> SENDBUF(*), RECVBUF(*)
  INTEGER COUNT, DATATYPE, OP, ROOT, COMM, REQUEST, IERROR

```

This call starts a nonblocking variant of MPI_REDUCE (see Section 6.9.1).

Advice to implementors. The implementation is explicitly allowed to use different algorithms for blocking and nonblocking reduction operations that might change the order of evaluation of the operations. However, as for MPI_REDUCE, it is strongly recommended that MPI_IREDUCE be implemented so that the same result be obtained whenever the function is applied on the same arguments, appearing in the same order. Note that this may prevent optimizations that take advantage of the physical location of MPI processes. (*End of advice to implementors.*)

Advice to users. For operations that are not truly associative, the result delivered upon completion of the nonblocking reduction may not exactly equal the result delivered by the blocking reduction, even when specifying the same arguments in the same order. (*End of advice to users.*)

6.12.8 Nonblocking All-Reduce

```

MPI_IALLREDUCE(sendbuf, recvbuf, count, datatype, op, comm, request)

```

IN	sendbuf	starting address of send buffer (choice)
OUT	recvbuf	starting address of receive buffer (choice)
IN	count	number of elements in send buffer (non-negative integer)
IN	datatype	datatype of elements of send buffer (handle)
IN	op	operation (handle)
IN	comm	communicator (handle)
OUT	request	communication request (handle)

C binding

```

int MPI_Iallreduce(const void *sendbuf, void *recvbuf, int count,
                  MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
                  MPI_Request *request)

```

```

int MPI_Iallreduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,
                    MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
                    MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Iallreduce(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
  INTEGER, INTENT(IN) :: count

```

```

TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Iallreduce(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
!(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_IALLREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER COUNT, DATATYPE, OP, COMM, REQUEST, IERROR

```

This call starts a nonblocking variant of `MPI_ALLREDUCE` (see Section 6.9.6).

6.12.9 Nonblocking Reduce-Scatter with Equal Blocks

```

MPI_IREDUCE_SCATTER_BLOCK(sendbuf, recvbuf, recvcount, datatype, op, comm,
request)
IN      sendbuf      starting address of send buffer (choice)
OUT     recvbuf      starting address of receive buffer (choice)
IN      recvcount    element count per block (non-negative integer)
IN      datatype     datatype of elements of send and receive buffers
                        (handle)
IN      op           operation (handle)
IN      comm         communicator (handle)
OUT     request      communication request (handle)

```

C binding

```

int MPI_Ireduce_scatter_block(const void *sendbuf, void *recvbuf,
int recvcount, MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
MPI_Request *request)
int MPI_Ireduce_scatter_block_c(const void *sendbuf, void *recvbuf,
MPI_Count recvcount, MPI_Datatype datatype, MPI_Op op,
MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```

1 MPI_Ireduce_scatter_block(sendbuf, recvbuf, recvcnt, datatype, op, comm,
2     request, ierror)
3

```

```

4     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
5

```

```

6     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
7

```

```

8     INTEGER, INTENT(IN) :: recvcnt
9

```

```

10    TYPE(MPI_Datatype), INTENT(IN) :: datatype
11

```

```

12    TYPE(MPI_Op), INTENT(IN) :: op
13

```

```

14    TYPE(MPI_Comm), INTENT(IN) :: comm
15

```

```

16    TYPE(MPI_Request), INTENT(OUT) :: request
17

```

```

18    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19

```

```

20 MPI_Ireduce_scatter_block(sendbuf, recvbuf, recvcnt, datatype, op, comm,
21     request, ierror) !(_c)
22

```

```

23    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
24

```

```

25    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
26

```

```

27    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcnt
28

```

```

29    TYPE(MPI_Datatype), INTENT(IN) :: datatype
30

```

```

31    TYPE(MPI_Op), INTENT(IN) :: op
32

```

```

33    TYPE(MPI_Comm), INTENT(IN) :: comm
34

```

```

35    TYPE(MPI_Request), INTENT(OUT) :: request
36

```

```

37    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
38

```

Fortran binding

```

39 MPI_IREDUCE_SCATTER_BLOCK(SENDBUF, RECVBUF, RECVCOUNT, DATATYPE, OP, COMM,
40     REQUEST, IERROR)
41

```

```

42     <type> SENDBUF(*), RECVBUF(*)
43

```

```

44     INTEGER RECVCOUNT, DATATYPE, OP, COMM, REQUEST, IERROR
45

```

This call starts a nonblocking variant of `MPI_REDUCE_SCATTER_BLOCK` (see Section 6.10.1).

6.12.10 Nonblocking Reduce-Scatter

```

46 MPI_IREDUCE_SCATTER(sendbuf, recvbuf, recvcnts, datatype, op, comm, request)
47

```

```

48     IN     sendbuf      starting address of send buffer (choice)
49

```

```

50     OUT    recvbuf      starting address of receive buffer (choice)
51

```

```

52     IN     recvcnts     nonnegative integer array specifying the number of
53                       elements in result distributed to each MPI process.
54                       This array must be identical on all calling MPI
55                       processes.
56

```

```

57     IN     datatype     datatype of elements of input buffer (handle)
58

```

```

59     IN     op           operation (handle)
60

```

```

61     IN     comm         communicator (handle)
62

```

```

63     OUT    request      communication request (handle)
64

```


C binding

```

int MPI_Ireduce_scatter(const void *sendbuf, void *recvbuf,
                       const int recvcnts[], MPI_Datatype datatype, MPI_Op op,
                       MPI_Comm comm, MPI_Request *request)
int MPI_Ireduce_scatter_c(const void *sendbuf, void *recvbuf,
                          const MPI_Count recvcnts[], MPI_Datatype datatype, MPI_Op op,
                          MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Ireduce_scatter(sendbuf, recvbuf, recvcnts, datatype, op, comm, request,
                   ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcnts(*)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Ireduce_scatter(sendbuf, recvbuf, recvcnts, datatype, op, comm, request,
                   ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcnts(*)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_IREDUCE_SCATTER(SENDBUF, RECVBUF, RECVCOUNTS, DATATYPE, OP, COMM, REQUEST,
                   IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER RECVCOUNTS(*), DATATYPE, OP, COMM, REQUEST, IERROR

```

This call starts a nonblocking variant of MPI_REDUCE_SCATTER (see Section 6.10.2).

6.12.11 Nonblocking Inclusive Scan

```

MPI_ISCAN(sendbuf, recvbuf, count, datatype, op, comm, request)
IN      sendbuf      starting address of send buffer (choice)
OUT     recvbuf      starting address of receive buffer (choice)
IN      count        number of elements in input buffer (non-negative
                    integer)
IN      datatype     datatype of elements of input buffer (handle)

```

```

1      IN      op                operation (handle)
2      IN      comm             communicator (handle)
3
4      OUT     request          communication request (handle)

```

C binding

```

7      int MPI_Iscan(const void *sendbuf, void *recvbuf, int count,
8                  MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
9                  MPI_Request *request)
10
11     int MPI_Iscan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
12                   MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
13                   MPI_Request *request)

```

Fortran 2008 binding

```

14     MPI_Iscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
15     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
16     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
17     INTEGER, INTENT(IN) :: count
18     TYPE(MPI_Datatype), INTENT(IN) :: datatype
19     TYPE(MPI_Op), INTENT(IN) :: op
20     TYPE(MPI_Comm), INTENT(IN) :: comm
21     TYPE(MPI_Request), INTENT(OUT) :: request
22     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24     MPI_Iscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror) !(_c)
25     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
26     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
27     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
28     TYPE(MPI_Datatype), INTENT(IN) :: datatype
29     TYPE(MPI_Op), INTENT(IN) :: op
30     TYPE(MPI_Comm), INTENT(IN) :: comm
31     TYPE(MPI_Request), INTENT(OUT) :: request
32     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

34     MPI_ISCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, REQUEST, IERROR)
35     <type> SENDBUF(*), RECVBUF(*)
36     INTEGER COUNT, DATATYPE, OP, COMM, REQUEST, IERROR

```

This call starts a nonblocking variant of MPI_SCAN (see Section 6.11).

6.12.12 Nonblocking Exclusive Scan

```

43     MPI_IEXSCAN(sendbuf, recvbuf, count, datatype, op, comm, request)
44
45     IN      sendbuf          starting address of send buffer (choice)
46     OUT     recvbuf          starting address of receive buffer (choice)
47     IN      count            number of elements in input buffer (non-negative
48                               integer)

```

IN	datatype	datatype of elements of input buffer (handle)	1
IN	op	operation (handle)	2
IN	comm	intra-communicator (handle)	3
OUT	request	communication request (handle)	4

C binding

```
int MPI_Iexscan(const void *sendbuf, void *recvbuf, int count,
               MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
               MPI_Request *request)

```

```
int MPI_Iexscan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
                 MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
                 MPI_Request *request)

```

Fortran 2008 binding

```
MPI_Iexscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```
MPI_Iexscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror) !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Op), INTENT(IN) :: op
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```
MPI_IEXSCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, REQUEST, IERROR)
  <type> SENDBUF(*), RECVBUF(*)
  INTEGER COUNT, DATATYPE, OP, COMM, REQUEST, IERROR

```

This call starts a nonblocking variant of MPI_EXSCAN (see Section 6.11.2).

6.13 Persistent Collective Operations

Many parallel computation algorithms involve repetitively executing a collective communication operation with the same arguments each time. As with persistent point-to-point operations (see Section 3.9), persistent collective operations allow the MPI programmer to

1 specify operations that will be reused frequently (with fixed arguments). MPI can be de-
2 signed to select a more efficient way to perform the collective operation based on the param-
3 eters specified when the operation is initialized. This “planned-transfer” approach [53, 42]
4 can offer significant performance benefits for programs with repetitive communication pat-
5 terns.

6 In terms of data movement, each persistent collective operation has the same effect as its
7 blocking and nonblocking counterparts for intra-communicators and inter-communicators
8 after completion. Likewise, upon completion, persistent collective reduction operations
9 perform the same operation as their blocking and nonblocking counterparts, and the same
10 restrictions and recommendations on reduction orders apply (see also Section 6.9.1).

11 Initialization calls for MPI persistent collective operations are nonlocal and follow all
12 the existing rules for collective operations, in particular ordering; programs that do not
13 conform to these restrictions are erroneous. After initialization, all arrays associated with
14 input arguments (such as arrays of counts, displacements, and datatypes in the vector
15 versions of the collectives) must not be modified until the corresponding persistent request
16 is freed with `MPI_REQUEST_FREE`.

17 According to the definitions in Section 2.4.2, the persistent collective initialization
18 procedures are incomplete. They are also nonlocal procedures because they may or may
19 not return before they are called in all MPI processes of the MPI process group associated
20 with the specified communicator.

21 *Advice to users.* This is one of the exceptions in which incomplete procedures are
22 nonlocal and therefore blocking. (*End of advice to users.*)

23
24 The `request` argument is an output argument that can be used zero or more times with
25 `MPI_START` or `MPI_STARTALL` in order to start the collective operation. The `request` is
26 initially inactive after the initialization call. Once initialized, persistent collective opera-
27 tions can be started in any order and the order can differ among the MPI processes in the
28 communicator.

29
30 *Rationale.* All ordering requirements that an implementation may need to match
31 up collective operations across the communicator are achieved through the ordering
32 requirements of the initialization functions. This enables out-of-order starts for the
33 persistent operations, and particularly supports their use in `MPI_STARTALL`. (*End of*
34 *rationale.*)

35
36 *Advice to implementors.* An MPI implementation should do no worse than duplicat-
37 ing the communicator during the initialization function, caching the input arguments,
38 and calling the appropriate nonblocking collective function, using the cached argu-
39 ments, during `MPI_START`. High-quality implementations should be able to amortize
40 setup costs and further optimize by taking advantage of early-binding, such as effi-
41 cient and effective pre-allocation of certain resources and algorithm selection. (*End*
42 *of advice to implementors.*)

43
44 A request must be inactive when it is started. Starting the operation makes the request
45 active. Once any MPI process starts a persistent collective operation, it must complete that
46 operation and all other MPI processes in the communicator must eventually start (and
47 complete) the same persistent collective operation. Persistent collective operations cannot
48 be *matched* with blocking or nonblocking collective operations. Completion of a persistent

collective operation makes the corresponding request inactive. After starting a persistent collective operation, all associated send buffers must not be modified and all associated receive buffers must not be accessed until the corresponding persistent request is completed.

Completing a persistent collective request, for example using `MPI_TEST` or `MPI_WAIT`, makes it inactive, but does not free the request. This is the same behavior as for persistent point-to-point requests. Inactive persistent collective requests can be freed using `MPI_REQUEST_FREE`. It is erroneous to free an active persistent collective request. Persistent collective operations cannot be canceled; it is erroneous to use `MPI_CANCEL` on a persistent collective request.

For every nonblocking collective communication operation in MPI, there is a corresponding persistent collective operation with the analogous API signature.

The collective persistent API signatures include an info object in order to support optimization hints and other information that may be nonstandard. Persistent collective operations may be optimized during communicator creation or by the initialization operation of an individual persistent collective. Note that communicator-scoped hints should be provided using `MPI_COMM_SET_INFO` while, for operation-scoped hints, they are supplied to the persistent collective communication initialization functions using the `info` argument.

6.13.1 Persistent Barrier Synchronization

`MPI_BARRIER_INIT(comm, info, request)`

IN	<code>comm</code>	communicator (handle)
IN	<code>info</code>	info argument (handle)
OUT	<code>request</code>	communication request (handle)

C binding

```
int MPI_Barrier_init(MPI_Comm comm, MPI_Info info, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Barrier_init(comm, info, request, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_BARRIER_INIT(COMM, INFO, REQUEST, IERROR)
  INTEGER COMM, INFO, REQUEST, IERROR
```

Creates a persistent collective communication request for the barrier operation.

6.13.2 Persistent Broadcast

```

1 MPI_BCAST_INIT(buffer, count, datatype, root, comm, info, request)
2
3
4

```

```

5     INOUT    buffer                starting address of buffer (choice)
6
7     IN      count                  number of entries in buffer (non-negative integer)
8
9     IN      datatype               datatype of buffer (handle)
10
11     IN      root                   rank of broadcast root (integer)
12
13     IN      comm                   communicator (handle)
14
15     IN      info                    info argument (handle)
16
17     OUT     request                 communication request (handle)

```

C binding

```

18 int MPI_Bcast_init(void *buffer, int count, MPI_Datatype datatype, int root,
19                   MPI_Comm comm, MPI_Info info, MPI_Request *request)

```

```

20 int MPI_Bcast_init_c(void *buffer, MPI_Count count, MPI_Datatype datatype,
21                     int root, MPI_Comm comm, MPI_Info info, MPI_Request *request)

```

Fortran 2008 binding

```

22 MPI_Bcast_init(buffer, count, datatype, root, comm, info, request, ierror)

```

```

23     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
24     INTEGER, INTENT(IN) :: count, root
25     TYPE(MPI_Datatype), INTENT(IN) :: datatype
26     TYPE(MPI_Comm), INTENT(IN) :: comm
27     TYPE(MPI_Info), INTENT(IN) :: info
28     TYPE(MPI_Request), INTENT(OUT) :: request
29     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

30 MPI_Bcast_init(buffer, count, datatype, root, comm, info, request, ierror)

```

```

31     !(_c)
32     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
33     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
34     TYPE(MPI_Datatype), INTENT(IN) :: datatype
35     INTEGER, INTENT(IN) :: root
36     TYPE(MPI_Comm), INTENT(IN) :: comm
37     TYPE(MPI_Info), INTENT(IN) :: info
38     TYPE(MPI_Request), INTENT(OUT) :: request
39     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

40 MPI_BCAST_INIT(BUFFER, COUNT, DATATYPE, ROOT, COMM, INFO, REQUEST, IERROR)

```

```

41     <type> BUFFER(*)
42     INTEGER COUNT, DATATYPE, ROOT, COMM, INFO, REQUEST, IERROR

```

```

43     Creates a persistent collective communication request for the broadcast operation.
44
45
46
47
48

```

6.13.3 Persistent Gather

MPI_GATHER_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root, comm, info, request)

IN	sendbuf	starting address of send buffer (choice)	1
IN	sendcount	number of elements in send buffer (non-negative integer)	2
IN	sendtype	datatype of send buffer elements (handle)	3
OUT	recvbuf	address of receive buffer (choice, significant only at root)	4
IN	recvcount	number of elements for any single receive (non-negative integer, significant only at root)	5
IN	recvtype	datatype of recv buffer elements (handle, significant only at root)	6
IN	root	rank of receiving MPI process (integer)	7
IN	comm	communicator (handle)	8
IN	info	info argument (handle)	9
OUT	request	communication request (handle)	10

C binding

```
int MPI_Gather_init(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
                  void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
                  MPI_Comm comm, MPI_Info info, MPI_Request *request)
```

```
int MPI_Gather_init_c(const void *sendbuf, MPI_Count sendcount,
                    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                    MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info,
                    MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Gather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
               root, comm, info, request, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER, INTENT(IN) :: sendcount, recvcount, root
```

```
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Info), INTENT(IN) :: info
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Gather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
               root, comm, info, request, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
```

```

1     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
2     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
3     INTEGER, INTENT(IN) :: root
4     TYPE(MPI_Comm), INTENT(IN) :: comm
5     TYPE(MPI_Info), INTENT(IN) :: info
6     TYPE(MPI_Request), INTENT(OUT) :: request
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

9     MPI_GATHER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE,
10                    ROOT, COMM, INFO, REQUEST, IERROR)
11     <type> SENDBUF(*), RECVBUF(*)
12     INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, ROOT, COMM, INFO,
13                    REQUEST, IERROR

```

Creates a persistent collective communication request for the gather operation.

```

14
15
16
17
18     MPI_GATHERV_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs, recvtype,
19                    root, comm, info, request)
20     IN     sendbuf          starting address of send buffer (choice)
21     IN     sendcount       number of elements in send buffer (non-negative
22                                     integer)
23
24     IN     sendtype        datatype of send buffer elements (handle)
25     OUT    recvbuf         address of receive buffer (choice, significant only at
26                                     root)
27
28     IN     recvcnts        nonnegative integer array (of length group size)
29                                     containing the number of elements that are received
30                                     from each MPI process (significant only at root)
31
32     IN     displs          integer array (of length group size). Entry i specifies
33                                     the displacement relative to recvbuf at which to place
34                                     the incoming data from MPI process i (significant
35                                     only at root)
36
37     IN     recvtype        datatype of recv buffer elements (handle, significant
38                                     only at root)
39
40     IN     root            rank of receiving MPI process (integer)
41
42     IN     comm            communicator (handle)
43
44     IN     info            info argument (handle)
45
46     OUT    request        communication request (handle)

```

C binding

```

44     int MPI_Gatherv_init(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
45                        void *recvbuf, const int recvcnts[], const int displs[],
46                        MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info,
47                        MPI_Request *request)
48

```



```

int MPI_Gatherv_init_c(const void *sendbuf, MPI_Count sendcount,
                      MPI_Datatype sendtype, void *recvbuf,
                      const MPI_Count recvcnts[], const MPI_Aint displs[],
                      MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info,
                      MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Gatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
                 recvtype, root, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcnts(*), displs(*)
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Gatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
                 recvtype, root, comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcnts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_GATHERV_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
                 RECVTYPE, ROOT, COMM, INFO, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT,
COMM, INFO, REQUEST, IERROR

```

Creates a persistent collective communication request for the gatherv operation.

6.13.4 Persistent Scatter

```

MPI_SCATTER_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcnt, recvtype, root,
                 comm, info, request)

```

IN	sendbuf	address of send buffer (choice, significant only at root)
----	---------	---

1	IN	sendcount	number of elements sent to each MPI process (non-negative integer, significant only at root)
2			
3	IN	sendtype	datatype of send buffer elements (handle, significant only at root)
4			
5	OUT	recvbuf	address of receive buffer (choice)
6			
7	IN	recvcount	number of elements in receive buffer (non-negative integer)
8			
9	IN	recvtype	datatype of receive buffer elements (handle)
10			
11	IN	root	rank of sending MPI process (integer)
12	IN	comm	communicator (handle)
13	IN	info	info argument (handle)
14			
15	OUT	request	communication request (handle)

16

17 **C binding**

```
18 int MPI_Scatter_init(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
19                   void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
20                   MPI_Comm comm, MPI_Info info, MPI_Request *request)
```

```
21 int MPI_Scatter_init_c(const void *sendbuf, MPI_Count sendcount,
22                      MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
23                      MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info,
24                      MPI_Request *request)
```

25

26 **Fortran 2008 binding**

```
27 MPI_Scatter_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
28                root, comm, info, request, ierror)
```

```
29 TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
30 INTEGER, INTENT(IN) :: sendcount, recvcount, root
31 TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
32 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
33 TYPE(MPI_Comm), INTENT(IN) :: comm
34 TYPE(MPI_Info), INTENT(IN) :: info
35 TYPE(MPI_Request), INTENT(OUT) :: request
36 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
37 MPI_Scatter_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
38                root, comm, info, request, ierror) !(_c)
```

```
39 TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
40 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
41 TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
42 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
43 INTEGER, INTENT(IN) :: root
44 TYPE(MPI_Comm), INTENT(IN) :: comm
45 TYPE(MPI_Info), INTENT(IN) :: info
46 TYPE(MPI_Request), INTENT(OUT) :: request
47 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

48

Fortran binding

```
MPI_SCATTER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
                 ROOT, COMM, INFO, REQUEST, IERROR)
```

```
<type> SENDBUF(*), RECVBUF(*)
```

```
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, INFO,
        REQUEST, IERROR
```

Creates a persistent collective communication request for the scatter operation.

```
MPI_SCATTERV_INIT(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcoun, recvtype,
                  root, comm, info, request)
```

IN	sendbuf	address of send buffer (choice, significant only at root)
IN	sendcounts	nonnegative integer array (of length group size) specifying the number of elements to send to each MPI process (significant only at root)
IN	displs	integer array (of length group size). Entry <i>i</i> specifies the displacement (relative to <i>sendbuf</i>) from which to take the outgoing data to MPI process <i>i</i> (significant only at root)
IN	sendtype	datatype of send buffer elements (handle, significant only at root)
OUT	recvbuf	address of receive buffer (choice, significant only at root)
IN	recvcoun	number of elements in receive buffer (non-negative integer)
IN	recvtype	datatype of receive buffer elements (handle)
IN	root	rank of sending MPI process (integer)
IN	comm	communicator (handle)
IN	info	info argument (handle)
OUT	request	communication request (handle)

C binding

```
int MPI_Scatterv_init(const void *sendbuf, const int sendcounts[],
                    const int displs[], MPI_Datatype sendtype, void *recvbuf,
                    int recvcoun, MPI_Datatype recvtype, int root, MPI_Comm comm,
                    MPI_Info info, MPI_Request *request)
```

```
int MPI_Scatterv_init_c(const void *sendbuf, const MPI_Count sendcounts[],
                      const MPI_Aint displs[], MPI_Datatype sendtype, void *recvbuf,
                      MPI_Count recvcoun, MPI_Datatype recvtype, int root,
                      MPI_Comm comm, MPI_Info info, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Scatterv_init(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcoun,
                  recvtype, root, comm, info, request, ierror)
```

```

1     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
2     INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), displs(*)
3     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
4     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
5     INTEGER, INTENT(IN) :: recvcount, root
6     TYPE(MPI_Comm), INTENT(IN) :: comm
7     TYPE(MPI_Info), INTENT(IN) :: info
8     TYPE(MPI_Request), INTENT(OUT) :: request
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10
11 MPI_Scatterv_init(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
12                 recvtype, root, comm, info, request, ierror) !(_c)
13     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
14     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*)
15     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
16     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
17     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
18     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcount
19     INTEGER, INTENT(IN) :: root
20     TYPE(MPI_Comm), INTENT(IN) :: comm
21     TYPE(MPI_Info), INTENT(IN) :: info
22     TYPE(MPI_Request), INTENT(OUT) :: request
23     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

24 MPI_SCATTERV_INIT(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,
25                 RECVTYPE, ROOT, COMM, INFO, REQUEST, IERROR)
26     <type> SENDBUF(*), RECVBUF(*)
27     INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,
28     COMM, INFO, REQUEST, IERROR

```

Creates a persistent collective communication request for the scatterv operation.

6.13.5 Persistent Gather-to-all

```

35 MPI_ALLGATHER_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
36                 info, request)

```

38	IN	sendbuf	starting address of send buffer (choice)
39	IN	sendcount	number of elements in send buffer (non-negative integer)
40			
41			
42	IN	sendtype	datatype of send buffer elements (handle)
43	OUT	recvbuf	address of receive buffer (choice)
44	IN	recvcount	number of elements received from any MPI process (non-negative integer)
45			
46			
47	IN	recvtype	datatype of receive buffer elements (handle)
48	IN	comm	communicator (handle)

IN	info	info argument (handle)	1
OUT	request	communication request (handle)	2

C binding

```

int MPI_Allgather_init(const void *sendbuf, int sendcount,
                      MPI_Datatype sendtype, void *recvbuf, int recvcount,
                      MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                      MPI_Request *request)

```

```

int MPI_Allgather_init_c(const void *sendbuf, MPI_Count sendcount,
                        MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                        MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                        MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
                  comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
                  comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_ALLGATHER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
                  COMM, INFO, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, INFO, REQUEST,
IERROR

```

Creates a persistent collective communication request for the allgather operation.

```

1 MPI_ALLGATHERV_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcoun
2     comm, info, request)
3
4     IN     sendbuf     starting address of send buffer (choice)
5     IN     sendcount   number of elements in send buffer (non-negative
6     integer)
7     IN     sendtype    datatype of send buffer elements (handle)
8     OUT    recvbuf     address of receive buffer (choice)
9
10    IN     recvcoun
11           ts          nonnegative integer array (of length group size)
12           containing the number of elements that are received
13           from each MPI process
14
15    IN     displs      integer array (of length group size). Entry i specifies
16           the displacement (relative to recvbuf) at which to
17           place the incoming data from MPI process i
18
19    IN     recvtype    datatype of receive buffer elements (handle)
20
21    IN     comm         communicator (handle)
22
23    IN     info         info argument (handle)
24
25    OUT    request     communication request (handle)

```

C binding

```

26 int MPI_Allgatherv_init(const void *sendbuf, int sendcount,
27     MPI_Datatype sendtype, void *recvbuf, const int recvcoun
28     ts[],
29     const int displs[], MPI_Datatype recvtype, MPI_Comm comm,
30     MPI_Info info, MPI_Request *request)
31
32 int MPI_Allgatherv_init_c(const void *sendbuf, MPI_Count sendcount,
33     MPI_Datatype sendtype, void *recvbuf,
34     const MPI_Count recvcoun
35     ts[], const MPI_Aint displs[],
36     MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
37     MPI_Request *request)

```

Fortran 2008 binding

```

38 MPI_Allgatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcoun
39     ts,
40     displs, recvtype, comm, info, request, ierror)
41     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
42     INTEGER, INTENT(IN) :: sendcount
43     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
44     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
45     INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcoun
46     ts(*), displs(*)
47     TYPE(MPI_Comm), INTENT(IN) :: comm
48     TYPE(MPI_Info), INTENT(IN) :: info
49     TYPE(MPI_Request), INTENT(OUT) :: request
50     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
51
52 MPI_Allgatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcoun
53     ts,
54     displs, recvtype, comm, info, request, ierror) !(_c)
55     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf

```

```

INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount           1
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype           2
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf               3
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcunts(*) 4
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*) 5
TYPE(MPI_Comm), INTENT(IN) :: comm                             6
TYPE(MPI_Info), INTENT(IN) :: info                             7
TYPE(MPI_Request), INTENT(OUT) :: request                       8
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       9

```

Fortran binding

```

MPI_ALLGATHERV_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
                   RECVTYPE, COMM, INFO, REQUEST, IERROR)      11
<type> SENDBUF(*), RECVBUF(*)                                  12
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
                   INFO, REQUEST, IERROR                       13

```

Creates a persistent collective communication request for the allgather operation.

6.13.6 Persistent All-to-All Scatter/Gather

```

MPI_ALLTOALL_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
                  info, request)                                19

```

IN	sendbuf	starting address of send buffer (choice)	20
IN	sendcount	number of elements sent to each MPI process (non-negative integer)	21
IN	sendtype	datatype of send buffer elements (handle)	22
OUT	recvbuf	address of receive buffer (choice)	23
IN	recvcount	number of elements received from any MPI process (non-negative integer)	24
IN	recvtype	datatype of receive buffer elements (handle)	25
IN	comm	communicator (handle)	26
IN	info	info argument (handle)	27
OUT	request	communication request (handle)	28

C binding

```

int MPI_Alltoall_init(const void *sendbuf, int sendcount,      29
                    MPI_Datatype sendtype, void *recvbuf, int recvcount,
                    MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                    MPI_Request *request)                       30
int MPI_Alltoall_init_c(const void *sendbuf, MPI_Count sendcount,
                    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                    MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                    MPI_Request *request)                       31

```

Fortran 2008 binding

```

1 MPI_Alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
2 comm, info, request, ierror)
3

```

```

4 TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
5

```

```

6 INTEGER, INTENT(IN) :: sendcount, recvcount
7

```

```

8 TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
9

```

```

10 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
11

```

```

12 TYPE(MPI_Comm), INTENT(IN) :: comm
13

```

```

14 TYPE(MPI_Info), INTENT(IN) :: info
15

```

```

16 TYPE(MPI_Request), INTENT(OUT) :: request
17

```

```

18 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19

```

```

20 MPI_Alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
21 comm, info, request, ierror) !(_c)
22

```

```

23 TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
24

```

```

25 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
26

```

```

27 TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
28

```

```

29 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
30

```

```

31 TYPE(MPI_Comm), INTENT(IN) :: comm
32

```

```

33 TYPE(MPI_Info), INTENT(IN) :: info
34

```

```

35 TYPE(MPI_Request), INTENT(OUT) :: request
36

```

```

37 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
38

```

Fortran binding

```

39 MPI_ALLTOALL_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
40 COMM, INFO, REQUEST, IERROR)
41

```

```

42 <type> SENDBUF(*), RECVBUF(*)
43

```

```

44 INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, INFO, REQUEST,
45 IERROR
46

```

47 Creates a persistent collective communication request for the alltoall operation.

```

48 MPI_ALLTOALLV_INIT(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts, rdispls,
49 recvtype, comm, info, request)
50

```

```

51 IN sendbuf starting address of send buffer (choice)
52

```

```

53 IN sendcounts nonnegative integer array (of length group size)
54 specifying the number of elements to send to each
55 MPI process
56

```

```

57 IN sdispls integer array (of length group size). Entry j specifies
58 the displacement (relative to sendbuf) from which to
59 take the outgoing data destined for MPI process j
60

```

```

61 IN sendtype datatype of send buffer elements (handle)
62

```

```

63 OUT recvbuf address of receive buffer (choice)
64

```

```

65 IN recvcounts nonnegative integer array (of length group size)
66 specifying the number of elements that can be
67 received from each MPI process
68

```


IN	rdispls	integer array (of length group size). Entry <i>i</i> specifies the displacement (relative to <i>recvbuf</i>) at which to place the incoming data from MPI process <i>i</i>	1 2 3
IN	recvtype	datatype of receive buffer elements (handle)	4 5
IN	comm	communicator (handle)	6
IN	info	info argument (handle)	7 8
OUT	request	communication request (handle)	9

C binding

```

int MPI_Alltoallv_init(const void *sendbuf, const int sendcounts[],
                      const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
                      const int rcvcounts[], const int rdispls[],
                      MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                      MPI_Request *request)

```

```

int MPI_Alltoallv_init_c(const void *sendbuf, const MPI_Count sendcounts[],
                        const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf,
                        const MPI_Count rcvcounts[], const MPI_Aint rdispls[],
                        MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                        MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf, rcvcounts,
                  rdispls, recvtype, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
rcvcounts(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf, rcvcounts,
                  rdispls, recvtype, comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
rcvcounts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

1 MPI_ALLTOALLV_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, RECVCOUNTS,
2   RDISPLS, RECVTYPE, COMM, INFO, REQUEST, IERROR)
3

```

```

4   <type> SENDBUF(*), RECVBUF(*)
5

```

```

6   INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*),
7   RECVTYPE, COMM, INFO, REQUEST, IERROR
8

```

Creates a persistent collective communication request for the alltoallv operation.

```

10 MPI_ALLTOALLW_INIT(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcoun
11   ts, rdispls,
12   recvtypes, comm, info, request)
13

```

12	IN	sendbuf	starting address of send buffer (choice)
13			
14	IN	sendcounts	integer array (of length group size) specifying the
15			number of elements to send to each MPI process
16			(array of non-negative integers)
17	IN	sdispls	integer array (of length group size). Entry j specifies
18			the displacement in bytes (relative to sendbuf) from
19			which to take the outgoing data destined for MPI
20			process j (array of integers)
21			
22	IN	sendtypes	array of datatypes (of length group size). Entry j
23			specifies the type of data to send to MPI process j
24			(array of handles)
25	OUT	recvbuf	address of receive buffer (choice)
26	IN	recvcoun	integer array (of length group size) specifying the
27		ts	number of elements that can be received from each
28			MPI process (array of non-negative integers)
29	IN	rdispls	integer array (of length group size). Entry i specifies
30			the displacement in bytes (relative to recvbuf) at
31			which to place the incoming data from MPI process i
32			(array of integers)
33			
34	IN	recvtypes	array of datatypes (of length group size). Entry i
35			specifies the type of data received from MPI process i
36			(array of handles)
37	IN	comm	communicator (handle)
38	IN	info	info argument (handle)
39	OUT	request	communication request (handle)
40			

C binding

```

43 int MPI_Alltoallw_init(const void *sendbuf, const int sendcounts[],
44   const int sdispls[], const MPI_Datatype sendtypes[],
45   void *recvbuf, const int recvcoun
46   ts, const int rdispls[],
47   const MPI_Datatype recvtypes[], MPI_Comm comm, MPI_Info info,
48   MPI_Request *request)

```

```

int MPI_Alltoallw_init_c(const void *sendbuf, const MPI_Count sendcounts[],
                        const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
                        void *recvbuf, const MPI_Count recvcounts[],
                        const MPI_Aint rdispls[], const MPI_Datatype recvtypes[],
                        MPI_Comm comm, MPI_Info info, MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
                  recvcounts, rdispls, recvtypes, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
    recvcounts(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
                  recvcounts, rdispls, recvtypes, comm, info, request, ierror)
    !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
    recvcounts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
    rdispls(*)
TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_ALLTOALLW_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF,
                  RECVCOUNTS, RDISPLS, RECVTYPES, COMM, INFO, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPES(*), RECVCOUNTS(*), RDISPLS(*),
    RECVTYPES(*), COMM, INFO, REQUEST, IERROR

```

Creates a persistent collective communication request for the alltoallw operation.

6.13.7 Persistent Reduce

```

MPI_REDUCE_INIT(sendbuf, recvbuf, count, datatype, op, root, comm, info, request)

```

```

IN          sendbuf          address of send buffer (choice)

```

1	OUT	recvbuf	address of receive buffer (choice, significant only at root)
2			
3	IN	count	number of elements in send buffer (non-negative integer)
4			
5	IN	datatype	datatype of elements of send buffer (handle)
6			
7	IN	op	reduce operation (handle)
8			
9	IN	root	rank of the root (integer)
10	IN	comm	communicator (handle)
11	IN	info	info argument (handle)
12	OUT	request	communication request (handle)
13			

C binding

```

15 int MPI_Reduce_init(const void *sendbuf, void *recvbuf, int count,
16                   MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
17                   MPI_Info info, MPI_Request *request)
18
19 int MPI_Reduce_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
20                      MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
21                      MPI_Info info, MPI_Request *request)

```

Fortran 2008 binding

```

23 MPI_Reduce_init(sendbuf, recvbuf, count, datatype, op, root, comm, info,
24               request, ierror)
25
26   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
27   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
28   INTEGER, INTENT(IN) :: count, root
29   TYPE(MPI_Datatype), INTENT(IN) :: datatype
30   TYPE(MPI_Op), INTENT(IN) :: op
31   TYPE(MPI_Comm), INTENT(IN) :: comm
32   TYPE(MPI_Info), INTENT(IN) :: info
33   TYPE(MPI_Request), INTENT(OUT) :: request
34   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
35
36 MPI_Reduce_init(sendbuf, recvbuf, count, datatype, op, root, comm, info,
37               request, ierror) !(_c)
38
39   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
40   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
41   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
42   TYPE(MPI_Datatype), INTENT(IN) :: datatype
43   TYPE(MPI_Op), INTENT(IN) :: op
44   INTEGER, INTENT(IN) :: root
45   TYPE(MPI_Comm), INTENT(IN) :: comm
46   TYPE(MPI_Info), INTENT(IN) :: info
47   TYPE(MPI_Request), INTENT(OUT) :: request
48   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_REDUCE_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, INFO,
               REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER COUNT, DATATYPE, OP, ROOT, COMM, INFO, REQUEST, IERROR

Creates a persistent collective communication request for the reduce operation.

```

6.13.8 Persistent All-Reduce

```

MPI_ALLREDUCE_INIT(sendbuf, recvbuf, count, datatype, op, comm, info, request)

IN      sendbuf      starting address of send buffer (choice)
OUT     recvbuf      starting address of receive buffer (choice)
IN      count        number of elements in send buffer (non-negative
                    integer)
IN      datatype     datatype of elements of send buffer (handle)
IN      op           operation (handle)
IN      comm         communicator (handle)
IN      info         info argument (handle)
OUT     request      communication request (handle)

```

C binding

```

int MPI_Allreduce_init(const void *sendbuf, void *recvbuf, int count,
                      MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
                      MPI_Request *request)

int MPI_Allreduce_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
                        MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
                        MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Allreduce_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
                  ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Allreduce_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
                  ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf

```

```

1     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
2     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
3     TYPE(MPI_Datatype), INTENT(IN) :: datatype
4     TYPE(MPI_Op), INTENT(IN) :: op
5     TYPE(MPI_Comm), INTENT(IN) :: comm
6     TYPE(MPI_Info), INTENT(IN) :: info
7     TYPE(MPI_Request), INTENT(OUT) :: request
8     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

10 MPI_ALLREDUCE_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, INFO, REQUEST,
11                    IERROR)

```

```

12     <type> SENDBUF(*), RECVBUF(*)
13     INTEGER COUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR

```

Creates a persistent collective communication request for the allreduce operation.

6.13.9 Persistent Reduce-Scatter with Equal Blocks

```

20 MPI_REDUCE_SCATTER_BLOCK_INIT(sendbuf, recvbuf, recvcnt, datatype, op, comm,
21                               info, request)

```

23	IN	sendbuf	starting address of send buffer (choice)
24	OUT	recvbuf	starting address of receive buffer (choice)
25	IN	recvcnt	element count per block (non-negative integer)
26	IN	datatype	datatype of elements of send and receive buffers (handle)
27			
28			
29	IN	op	operation (handle)
30	IN	comm	communicator (handle)
31	IN	info	info argument (handle)
32			
33	OUT	request	communication request (handle)

C binding

```

36 int MPI_Reduce_scatter_block_init(const void *sendbuf, void *recvbuf,
37                                 int recvcnt, MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
38                                 MPI_Info info, MPI_Request *request)

```

```

40 int MPI_Reduce_scatter_block_init_c(const void *sendbuf, void *recvbuf,
41                                   MPI_Count recvcnt, MPI_Datatype datatype, MPI_Op op,
42                                   MPI_Comm comm, MPI_Info info, MPI_Request *request)

```

Fortran 2008 binding

```

44 MPI_Reduce_scatter_block_init(sendbuf, recvbuf, recvcnt, datatype, op, comm,
45                               info, request, ierror)

```

```

46     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
47     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
48     INTEGER, INTENT(IN) :: recvcnt

```

```

TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Reduce_scatter_block_init(sendbuf, recvbuf, recvcnt, datatype, op, comm,
                             info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcnt
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_REDUCE_SCATTER_BLOCK_INIT(SENDBUF, RECVBUF, RECVCOUNT, DATATYPE, OP, COMM,
                              INFO, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER REVCOUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR

```

Creates a persistent collective communication request for the reduce-scatter with equal blocks operation.

6.13.10 Persistent Reduce-Scatter

```

MPI_REDUCE_SCATTER_INIT(sendbuf, recvbuf, recvcnts, datatype, op, comm, info,
                        request)

```

IN	sendbuf	starting address of send buffer (choice)	33
OUT	recvbuf	starting address of receive buffer (choice)	34
IN	recvcnts	nonnegative integer array specifying the number of elements in result distributed to each MPI process. This array must be identical on all calling MPI processes.	35 36 37 38 39
IN	datatype	datatype of elements of input buffer (handle)	40
IN	op	operation (handle)	41
IN	comm	communicator (handle)	42
IN	info	info argument (handle)	43 44
OUT	request	communication request (handle)	45 46

C binding

```

int MPI_Reduce_scatter_init(const void *sendbuf, void *recvbuf,

```

```

1         const int recvcnts[], MPI_Datatype datatype, MPI_Op op,
2         MPI_Comm comm, MPI_Info info, MPI_Request *request)
3
4     int MPI_Reduce_scatter_init_c(const void *sendbuf, void *recvbuf,
5         const MPI_Count recvcnts[], MPI_Datatype datatype, MPI_Op op,
6         MPI_Comm comm, MPI_Info info, MPI_Request *request)
7
8     Fortran 2008 binding
9     MPI_Reduce_scatter_init(sendbuf, recvbuf, recvcnts, datatype, op, comm, info,
10        request, ierror)
11     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
12     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
13     INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcnts(*)
14     TYPE(MPI_Datatype), INTENT(IN) :: datatype
15     TYPE(MPI_Op), INTENT(IN) :: op
16     TYPE(MPI_Comm), INTENT(IN) :: comm
17     TYPE(MPI_Info), INTENT(IN) :: info
18     TYPE(MPI_Request), INTENT(OUT) :: request
19     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21     MPI_Reduce_scatter_init(sendbuf, recvbuf, recvcnts, datatype, op, comm, info,
22        request, ierror) !(_c)
23     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
24     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
25     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcnts(*)
26     TYPE(MPI_Datatype), INTENT(IN) :: datatype
27     TYPE(MPI_Op), INTENT(IN) :: op
28     TYPE(MPI_Comm), INTENT(IN) :: comm
29     TYPE(MPI_Info), INTENT(IN) :: info
30     TYPE(MPI_Request), INTENT(OUT) :: request
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33     Fortran binding
34     MPI_REDUCE_SCATTER_INIT(SENDBUF, RECVBUF, REVCOUNTS, DATATYPE, OP, COMM, INFO,
35        REQUEST, IERROR)
36     <type> SENDBUF(*), RECVBUF(*)
37     INTEGER REVCOUNTS(*), DATATYPE, OP, COMM, INFO, REQUEST, IERROR
38
39     Creates a persistent collective communication request for the reduce-scatter operation.
40
41
42
43
44
45
46
47
48

```

6.13.11 Persistent Inclusive Scan

```

41
42     MPI_SCAN_INIT(sendbuf, recvbuf, count, datatype, op, comm, info, request)
43
44     IN         sendbuf           starting address of send buffer (choice)
45     OUT        recvbuf           starting address of receive buffer (choice)
46     IN         count             number of elements in input buffer (non-negative
47                                     integer)
48     IN         datatype           datatype of elements of input buffer (handle)

```


IN	op	operation (handle)	1
IN	comm	communicator (handle)	2
IN	info	info argument (handle)	3
OUT	request	communication request (handle)	4

C binding

```
int MPI_Scan_init(const void *sendbuf, void *recvbuf, int count,
                MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
                MPI_Request *request)
```

```
int MPI_Scan_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
                  MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
                  MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Scan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
              ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
INTEGER, INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Op), INTENT(IN) :: op
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Info), INTENT(IN) :: info
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Scan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
              ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
TYPE(MPI_Op), INTENT(IN) :: op
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Info), INTENT(IN) :: info
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SCAN_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, INFO, REQUEST,
              IERROR)
```

```
<type> SENDBUF(*), RECVBUF(*)
```

```
INTEGER COUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR
```

Creates a persistent collective communication request for the inclusive scan operation.

6.13.12 Persistent Exclusive Scan

```

1 MPI_EXSCAN_INIT(sendbuf, recvbuf, count, datatype, op, comm, info, request)
2
3
4 MPI_EXSCAN_INIT(sendbuf, recvbuf, count, datatype, op, comm, info, request)
5
6     IN      sendbuf      starting address of send buffer (choice)
7     OUT    recvbuf      starting address of receive buffer (choice)
8     IN      count        number of elements in input buffer (non-negative
9                          integer)
10
11     IN      datatype     datatype of elements of input buffer (handle)
12     IN      op           operation (handle)
13     IN      comm         intra-communicator (handle)
14     IN      info         info argument (handle)
15     IN      info         info argument (handle)
16     OUT    request      communication request (handle)
17

```

C binding

```

18
19 int MPI_Exscan_init(const void *sendbuf, void *recvbuf, int count,
20                   MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
21                   MPI_Request *request)
22
23 int MPI_Exscan_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
24                      MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
25                      MPI_Request *request)
26

```

Fortran 2008 binding

```

27 MPI_Exscan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
28                ierror)
29     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
30     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
31     INTEGER, INTENT(IN) :: count
32     TYPE(MPI_Datatype), INTENT(IN) :: datatype
33     TYPE(MPI_Op), INTENT(IN) :: op
34     TYPE(MPI_Comm), INTENT(IN) :: comm
35     TYPE(MPI_Info), INTENT(IN) :: info
36     TYPE(MPI_Request), INTENT(OUT) :: request
37     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
38
39 MPI_Exscan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
40                ierror) !(_c)
41     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
42     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
43     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
44     TYPE(MPI_Datatype), INTENT(IN) :: datatype
45     TYPE(MPI_Op), INTENT(IN) :: op
46     TYPE(MPI_Comm), INTENT(IN) :: comm
47     TYPE(MPI_Info), INTENT(IN) :: info
48     TYPE(MPI_Request), INTENT(OUT) :: request

```

INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

MPI_EXSCAN_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, INFO, REQUEST,
IERROR)

<type> SENDBUF(*), RECVBUF(*)

INTEGER COUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR

Creates a persistent collective communication request for the exclusive scan operation.

6.14 Correctness

A correct, portable program must invoke collective communications so that deadlock will not occur, whether collective communications are synchronizing or not. The following examples illustrate dangerous use of collective routines on intra-communicators.

Example 6.26. The following is erroneous.

```

/* ----- THIS EXAMPLE IS ERRONEOUS ----- */
switch(rank) {
  case 0:
    MPI_Bcast(buf1, count, type, 0, comm);
    MPI_Bcast(buf2, count, type, 1, comm);
    break;
  case 1:
    MPI_Bcast(buf2, count, type, 1, comm);
    MPI_Bcast(buf1, count, type, 0, comm);
    break;
}

```

We assume that the group of `comm` is $\{0,1\}$. Two MPI processes execute two broadcast operations in reverse order. If the operation is synchronizing then a deadlock will occur. Collective operations must be executed in the same order at all members of the communication group.

Example 6.27. The following is erroneous.

```

/* ----- THIS EXAMPLE IS ERRONEOUS ----- */
switch(rank) {
  case 0:
    MPI_Bcast(buf1, count, type, 0, comm0);
    MPI_Bcast(buf2, count, type, 2, comm2);
    break;
  case 1:
    MPI_Bcast(buf1, count, type, 1, comm1);
    MPI_Bcast(buf2, count, type, 0, comm0);
    break;
  case 2:
    MPI_Bcast(buf1, count, type, 2, comm2);
    MPI_Bcast(buf2, count, type, 1, comm1);
    break;
}

```

Assume that the group of `comm0` is $\{0,1\}$, of `comm1` is $\{1, 2\}$ and of `comm2` is $\{2,0\}$. If the broadcast is a synchronizing operation, then there is a cyclic dependency: the broadcast in `comm2` completes only after the broadcast in `comm0`; the broadcast in `comm0` completes only after the broadcast in `comm1`; and the broadcast in `comm1` completes only after the broadcast in `comm2`. Thus, the code will deadlock.

Collective operations must be executed in an order so that no cyclic dependencies occur. Nonblocking collective operations can alleviate this issue.

Example 6.28. The following is erroneous.

```

/* ----- THIS EXAMPLE IS ERRONEOUS ----- */
switch(rank) {
  case 0:
    MPI_Bcast(buf1, count, type, 0, comm);
    MPI_Send(buf2, count, type, 1, tag, comm);
    break;
  case 1:
    MPI_Recv(buf2, count, type, 0, tag, comm, status);
    MPI_Bcast(buf1, count, type, 0, comm);
    break;
}

```

MPI process with rank 0 executes a broadcast, followed by a blocking send operation. MPI process with rank 1 first executes a blocking receive that matches the send, followed by a broadcast call that matches the broadcast of MPI process with rank 0. This program may deadlock. The broadcast call on MPI process with rank 0 *may* block until MPI process with rank 1 executes the matching broadcast call, so that the send is not executed. MPI process with rank 1 will definitely block on the receive and so, in this case, never executes the broadcast.

The relative order of execution of collective operations and point-to-point operations should be such, so that even if the collective operations and the point-to-point operations are synchronizing, no deadlock will occur.

Example 6.29. An unsafe, nondeterministic program.

```

switch(rank) {
  case 0:
    MPI_Bcast(buf1, count, type, 0, comm);
    MPI_Send(buf2, count, type, 1, tag, comm);
    break;
  case 1:
    MPI_Recv(buf2, count, type, MPI_ANY_SOURCE, tag, comm, status);
    MPI_Bcast(buf1, count, type, 0, comm);
    MPI_Recv(buf2, count, type, MPI_ANY_SOURCE, tag, comm, status);
    break;
  case 2:
    MPI_Send(buf2, count, type, 1, tag, comm);
    MPI_Bcast(buf1, count, type, 0, comm);
    break;
}

```

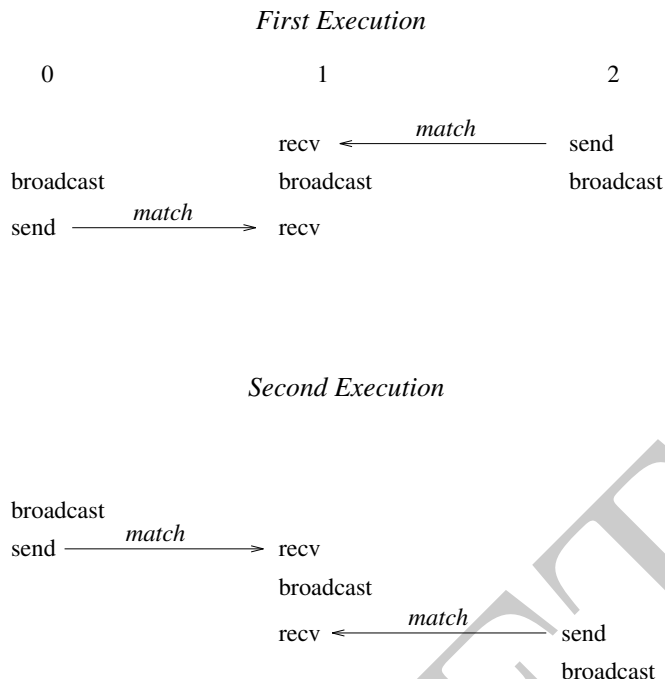


Figure 6.12: A race condition causes nondeterministic matching of sends and receives. One cannot rely on synchronization from a broadcast to make the program deterministic.

All three MPI processes participate in a broadcast. MPI process with rank 0 sends a message to MPI process with rank 1 after the broadcast, and MPI process with rank 2 sends a message to MPI process with rank 1 before the broadcast. MPI process with rank 1 receives before and after the broadcast, with a wildcard source argument. Two possible executions of this program, with different matchings of sends and receives, are illustrated in Figure 6.12. Note that the second execution has the peculiar effect that a send executed after the broadcast is received at another node before the broadcast. This example illustrates the fact that one should not rely on collective communication functions to have particular synchronization effects. A program that works correctly only when the first execution occurs (only when broadcast is synchronizing) is erroneous.

Finally, in multithreaded implementations, one can have more than one, concurrently executing, collective communication initialization call at an MPI process. In these situations, it is the user’s responsibility to ensure that the same communicator is not used concurrently by two different collective communication initialization calls at the same MPI process. Collective communication initialization calls include all calls for blocking collective operations, all initiation calls for nonblocking collective operations, and all initialization calls for persistent collective operations.

Advice to implementors. Assume that broadcast is implemented using point-to-point MPI communication. Suppose the following two rules are followed.

1. All receives specify their source explicitly (no wildcards).
2. Each MPI process sends all messages that pertain to one collective call before sending any message that pertain to a subsequent collective call.

1 Then, messages belonging to successive broadcasts cannot be confused, as the order
2 of point-to-point messages is preserved.

3 It is the implementor's responsibility to ensure that point-to-point messages are not
4 confused with collective messages. One way to accomplish this is, whenever a commu-
5 nicator is created, to also create a "hidden communicator" for collective communica-
6 tion. One could achieve a similar effect more cheaply, for example, by using a hidden
7 tag or context bit to indicate whether the communicator is used for point-to-point or
8 collective communication. (*End of advice to implementors.*)
9

10 **Example 6.30.** Blocking and nonblocking collective operations can be interleaved, i.e., a
11 blocking collective operation can be posted even if there is a nonblocking collective operation
12 outstanding.

```
13 MPI_Request req;  
14  
15 MPI_Ibarrier(comm, &req);  
16 MPI_Bcast(buf1, count, type, 0, comm);  
17 MPI_Wait(&req, MPI_STATUS_IGNORE);  
18
```

19 Each MPI process starts a nonblocking barrier operation, participates in a blocking broadcast
20 and then waits until every other MPI process started the barrier operation. This effectively
21 turns the broadcast into a synchronizing broadcast with possible communication/communi-
22 cation overlap (MPI_Bcast is allowed, but not required to synchronize).
23

24 **Example 6.31.** The starting order of collective operations on a particular communicator
25 defines their matching. The following example shows an erroneous matching of different
26 collective operations on the same communicator.

```
27 /* ----- THIS EXAMPLE IS ERRONEOUS ----- */  
28 MPI_Request req;  
29 switch(rank) {  
30     case 0:  
31         /* erroneous matching */  
32         MPI_Ibarrier(comm, &req);  
33         MPI_Bcast(buf1, count, type, 0, comm);  
34         MPI_Wait(&req, MPI_STATUS_IGNORE);  
35         break;  
36     case 1:  
37         /* erroneous matching */  
38         MPI_Bcast(buf1, count, type, 0, comm);  
39         MPI_Ibarrier(comm, &req);  
40         MPI_Wait(&req, MPI_STATUS_IGNORE);  
41         break;  
42 }  
43
```

44 This ordering would match MPI_Ibarrier on rank 0 with MPI_Bcast on rank 1, which is
45 erroneous and the program behavior is undefined. However, if such an order is required, the
46 user must create different duplicate communicators and perform the operations on them. If
47 started with two MPI processes, the following program would be correct:
48

```
49 MPI_Request req;  
50 MPI_Comm dupcomm;
```

```

MPI_Comm_dup(comm, &dupcomm);
switch(rank) {
  case 0:
    MPI_Ibarrier(comm, &req);
    MPI_Bcast(buf1, count, type, 0, dupcomm);
    MPI_Wait(&req, MPI_STATUS_IGNORE);
    break;
  case 1:
    MPI_Bcast(buf1, count, type, 0, dupcomm);
    MPI_Ibarrier(comm, &req);
    MPI_Wait(&req, MPI_STATUS_IGNORE);
    break;
}

```

Advice to users. The use of different communicators offers some flexibility regarding the matching of nonblocking collective operations. In this sense, communicators could be used as an equivalent to tags. However, communicator construction might induce overheads so that this should be used carefully. (*End of advice to users.*)

Example 6.32. Nonblocking collective operations can rely on the similar progress rules as nonblocking point-to-point operations. Thus, if started with two MPI processes, the following program is a valid MPI program and is guaranteed to terminate:

```

MPI_Request req;

switch(rank) {
  case 0:
    MPI_Ibarrier(comm, &req);
    MPI_Wait(&req, MPI_STATUS_IGNORE);
    MPI_Send(buf, count, dtype, 1, tag, comm);
    break;
  case 1:
    MPI_Ibarrier(comm, &req);
    MPI_Recv(buf, count, dtype, 0, tag, comm, MPI_STATUS_IGNORE);
    MPI_Wait(&req, MPI_STATUS_IGNORE);
    break;
}

```

The MPI library must *progress* the barrier in the MPI_Recv call. Thus, the MPI_Wait call in rank 0 will eventually complete, which enables the matching MPI_Send so all calls eventually return.

Example 6.33. Blocking and nonblocking collective operations do not match. The following example is erroneous.

```

/* ----- THIS EXAMPLE IS ERRONEOUS ----- */
MPI_Request req;

switch(rank) {
  case 0:
    /* erroneous false matching of Alltoall and Ialltoall */
    MPI_Ialltoall(sbuf, scnt, stype, rbuf, rcnt, rtype, comm, &req);

```

```

1   MPI_Wait(&req, MPI_STATUS_IGNORE);
2   break;
3   case 1:
4       /* erroneous false matching of Alltoall and Ialltoall */
5       MPI_Alltoall(sbuf, scnt, stype, rbuf, rcnt, rtype, comm);
6       break;
7   }

```

Example 6.34. Collective and point-to-point requests can be mixed in functions that enable multiple completions. If started with two MPI processes, the following program is valid.

```

13  MPI_Request reqs[2];
14
15  switch(rank) {
16      case 0:
17          MPI_Ibarrier(comm, &reqs[0]);
18          MPI_Send(buf, count, dtype, 1, tag, comm);
19          MPI_Wait(&reqs[0], MPI_STATUS_IGNORE);
20          break;
21      case 1:
22          MPI_Irecv(buf, count, dtype, 0, tag, comm, &reqs[0]);
23          MPI_Ibarrier(comm, &reqs[1]);
24          MPI_Waitall(2, reqs, MPI_STATUSES_IGNORE);
25          break;
26  }

```

The MPI_Waitall call returns only after the barrier and the receive completed.

Example 6.35. Multiple nonblocking collective operations can be outstanding on a single communicator and match in order.

```

31  MPI_Request reqs[3];
32
33  compute(buf1);
34  MPI_Ibcast(buf1, count, type, 0, comm, &reqs[0]);
35  compute(buf2);
36  MPI_Ibcast(buf2, count, type, 0, comm, &reqs[1]);
37  compute(buf3);
38  MPI_Ibcast(buf3, count, type, 0, comm, &reqs[2]);
39  MPI_Waitall(3, reqs, MPI_STATUSES_IGNORE);

```

Advice to users. Pipelining and double-buffering techniques can efficiently be used to overlap computation and communication. However, having too many outstanding requests might have a negative impact on performance. (*End of advice to users.*)

Advice to implementors. The use of pipelining may generate many outstanding requests. A high-quality hardware-supported implementation with limited resources should be able to fall back to a software implementation if its resources are exhausted. In this way, the implementation could limit the number of outstanding requests only by the available memory. (*End of advice to implementors.*)




Figure 6.13: Example with overlapping communicators.

Example 6.36. Nonblocking collective operations can also be used to enable simultaneous collective operations on multiple overlapping communicators (see Figure 6.13). The following example is started with three MPI processes and three communicators. The first communicator `comm1` includes ranks 0 and 1, `comm2` includes ranks 1 and 2, and `comm3` spans ranks 0 and 2. It is not possible to perform a blocking collective operation on all communicators because there exists no deadlock-free order to invoke them. However, nonblocking collective operations can easily be used to achieve this task.

```

MPI_Request reqs[2];

switch(rank) {
  case 0:
    MPI_Iallreduce(sbuf1, rbuf1, count, dtype, MPI_SUM, comm1, &reqs[0]);
    MPI_Iallreduce(sbuf3, rbuf3, count, dtype, MPI_SUM, comm3, &reqs[1]);
    break;
  case 1:
    MPI_Iallreduce(sbuf1, rbuf1, count, dtype, MPI_SUM, comm1, &reqs[0]);
    MPI_Iallreduce(sbuf2, rbuf2, count, dtype, MPI_SUM, comm2, &reqs[1]);
    break;
  case 2:
    MPI_Iallreduce(sbuf2, rbuf2, count, dtype, MPI_SUM, comm2, &reqs[0]);
    MPI_Iallreduce(sbuf3, rbuf3, count, dtype, MPI_SUM, comm3, &reqs[1]);
    break;
}
MPI_Waitall(2, reqs, MPI_STATUSES_IGNORE);

```

Advice to users. This method can be useful if overlapping neighboring regions (halo or ghost zones) are used in collective operations. The sequence of the two calls in each MPI process is irrelevant because the two nonblocking operations are performed on different communicators. (*End of advice to users.*)

Example 6.37. The *progress* of multiple outstanding nonblocking collective operations is completely independent.

```

MPI_Request reqs[2];

compute(buf1);
MPI_Ibcast(buf1, count, type, 0, comm, &reqs[0]);

```

```
1 compute(buf2);  
2 MPI_Ibcast(buf2, count, type, 0, comm, &reqs[1]);  
3 MPI_Wait(&reqs[1], MPI_STATUS_IGNORE);  
4 /* nothing is known about the status of the first bcast here */  
5 MPI_Wait(&reqs[0], MPI_STATUS_IGNORE);
```

6 Finishing the second MPI_IBCAST is completely independent of the first one. This means
7 that it is not guaranteed that the first broadcast operation is finished or even started after
8 the second one is completed via reqs[1].
9

10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

DRAFT

Chapter 7

Groups, Contexts, Communicators, and Caching

7.1 Introduction

This chapter introduces MPI features that support the development of parallel libraries. Parallel libraries are needed to encapsulate the distracting complications inherent in parallel implementations of key algorithms. They help to ensure consistent correctness of such procedures, and provide a “higher level” of portability than MPI itself can provide. As such, libraries prevent each programmer from repeating the work of defining consistent data structures, data layouts, and methods that implement key algorithms (such as matrix operations). Since the best libraries come with several variations on parallel systems (different data layouts, different strategies depending on the size of the system or problem, or type of floating point), this too needs to be hidden from the user.

We refer the reader to [5] and [63] for further information on writing libraries in MPI, using the features described in this chapter.

7.1.1 Features Needed to Support Libraries

The key features needed to support the creation of robust parallel libraries are as follows:

- Safe communication space, that guarantees that libraries can communicate as they need to, without conflicting with communication extraneous to the library,
- Group scope for collective operations, that allow libraries to avoid unnecessarily synchronizing uninvolved MPI processes (potentially running unrelated code),
- Abstract naming of MPI processes to allow libraries to describe their communication in terms suitable to their own data structures and algorithms,
- The ability to “adorn” a set of communicating MPI processes with additional user-defined attributes, such as extra collective operations. This mechanism should provide a means for the user or library writer effectively to extend a message-passing notation.

In addition, a unified mechanism or object is needed for conveniently denoting communication context, the group of communicating MPI processes, to house abstract naming of MPI processes, and to store adornments.

7.1.2 MPI’s Support for Libraries

The corresponding concepts that MPI provides, specifically to support robust libraries, are as follows:

- 1 • **Contexts** of communication,
- 2
- 3 • **Groups** of MPI processes,
- 4
- 5 • **Virtual topologies**,
- 6
- 7 • **Attribute caching**,
- 8
- 9 • **Communicators**.

9 **Communicators** (see [23, 61, 65]) encapsulate all of these ideas in order to provide the
10 appropriate scope for all communication operations in MPI. Communicators are divided
11 into two kinds: intra-communicators for operations within a single group of MPI processes
12 and inter-communicators for operations between two groups of MPI processes.

13
14 **Caching.** Communicators (see below) provide a “caching” mechanism that allows one
15 to associate new attributes with communicators, on par with MPI built-in features. This
16 can be used by advanced users to adorn communicators further, and by MPI to implement
17 some communicator functions. For example, the virtual-topology functions described in
18 Chapter 8 are likely to be supported this way.

19
20 **Groups.** Groups define an ordered collection of MPI processes, each with a rank, and it
21 is this group that defines the low-level names (ranks) for communication. Thus, groups
22 define a scope for MPI process names in point-to-point communication. In addition, groups
23 define the scope of collective operations. Groups may be manipulated separately from
24 communicators in MPI, but only communicators can be used in communication operations.

25
26 **Intra-Communicators.** The most commonly used means for message-passing in MPI is
27 via intra-communicators. Intra-communicators contain an instance of a group, contexts of
28 communication for both point-to-point and collective communication, and the ability to
29 include virtual topology and other attributes. These features work as follows:

- 30 • **Contexts** provide the ability to have separate safe “universes” of message-passing in
31 MPI. A context is akin to an additional tag that differentiates messages. The system
32 manages this differentiation process. The use of separate communication contexts
33 by distinct libraries (or distinct library invocations) insulates communication internal
34 to the library execution from external communication. This allows the invocation
35 of the library even if there are *pending* communication operations or *decoupled MPI*
36 *activities* on “other” communicators, and avoids the need to synchronize entry or exit
37 into library code. *Pending* communication or *decoupled MPI activities* of point-to-
38 point operations are also guaranteed not to interfere with collective communication
39 operations within a single communicator.
- 40
- 41 • **Groups** define the participants in the communication (see above) of a communicator.
- 42
- 43 • A **virtual topology** defines a special mapping of the MPI processes ranks in a group to
44 and from a topology. Special constructors for communicators are defined in Chapter 8
45 to provide this feature. Intra-communicators as described in this chapter do not have
46 topologies.
- 47 • **Attributes** define the local information that the user or library has added to a com-
48 municator for later reference.

Advice to users. The practice in many communication libraries is that there is a unique, predefined communication universe that includes all MPI processes available when the parallel program is initiated; the MPI processes are assigned consecutive ranks. Participants in a point-to-point communication are identified by their rank; a collective communication (such as broadcast) always involves all MPI processes. When using the World Model (Section 11.2), this practice can be followed in MPI by using the predefined communicator `MPI_COMM_WORLD`. (*End of advice to users.*)

Inter-Communicators. The discussion has dealt so far with **intra-communication**: communication within a group. MPI also supports **inter-communication**: communication between two nonoverlapping groups. When an application is built by composing several parallel modules, it is convenient to allow one module to communicate with another using local ranks for addressing within the second module. This is especially convenient in a client-server computing paradigm, where either client or server are parallel. The support of inter-communication also provides a mechanism for the extension of MPI to a dynamic model where not all MPI processes are preallocated at initialization time. In such a situation, it becomes necessary to support communication across “universes.” Inter-communication is supported by objects called **inter-communicators**. These objects bind two groups together with communication contexts shared by both groups. For inter-communicators, these features work as follows:

- Contexts provide the ability to have a separate safe “universe” of message-passing between the two groups. A send operation in the local group is always matched by a receive operation in the remote group, and vice versa. The system manages this differentiation process. The use of separate communication contexts by distinct libraries (or distinct library invocations) insulates communication internal to the library execution from external communication. This allows the invocation of the library even if there are *pending* communication operations or *decoupled MPI activities* on “other” communicators, and avoids the need to synchronize entry or exit into library code.
- A local and remote group specify the recipients and destinations for an inter-communicator.
- Virtual topology is undefined for an inter-communicator.
- As before, attributes cache defines the local information that the user or library has added to a communicator for later reference.

MPI provides mechanisms for creating and manipulating inter-communicators. They are used for point-to-point and collective communication in a related manner to intra-communicators. Users who do not need inter-communication in their applications can safely ignore this extension. Users who require inter-communication between overlapping groups must layer this capability on top of MPI.

7.2 Basic Concepts

In this section, we turn to a more formal definition of the concepts introduced above.

7.2.1 Groups

A **group** is an ordered set of MPI process identifiers (henceforth MPI processes); MPI processes are implementation-dependent objects. Each MPI process in a group is associated with an integer **rank**. Ranks are consecutive and start from zero. Groups are represented by opaque **group objects**, and hence cannot be directly transferred from one MPI process to another. A group is used within a communicator to describe the participants in a communication “universe” and to rank such participants (thus giving them unique names within that “universe” of communication).

There is a special pre-defined group: `MPI_GROUP_EMPTY`, which is a group with no members. The predefined constant `MPI_GROUP_NULL` is the value used for invalid group handles.

Advice to users. `MPI_GROUP_EMPTY`, which is a valid handle to an empty group, should not be confused with `MPI_GROUP_NULL`, which in turn is an invalid handle. The former may be used as an argument to group procedures; the latter is not a valid input value for an input argument. (*End of advice to users.*)

Advice to implementors. Simple implementations of MPI will enumerate groups, such as in a table. However, more advanced data structures make sense in order to improve scalability and memory usage with large numbers of MPI processes. Such implementations are possible with MPI. (*End of advice to implementors.*)

7.2.2 Contexts

A **context** is a property of communicators (defined next) that allows partitioning of the communication space. A message sent in one context cannot be received in another context. Furthermore, where permitted, collective operations are independent of *pending* point-to-point operations and *decoupled MPI activities* of point-to-point operations. Contexts are not explicit MPI objects; they appear only as part of the realization of communicators (below).

Advice to implementors. Distinct communicators in the same MPI process have distinct contexts. A context is essentially a system-managed tag (or tags) needed to make a communicator safe for point-to-point and MPI-defined collective communication. Safety means that collective and point-to-point communication within one communicator do not interfere, and that communication over distinct communicators do not interfere.

A possible implementation for a context is as a supplemental tag attached to messages on send and matched on receive. Each intra-communicator stores the value of its two tags (one for point-to-point and one for collective communication). Communicator-generating functions use a collective communication to agree on a new group-wide unique context.

Analogously, in inter-communication, two context tags are stored per communicator, one used by group A to send and group B to receive, and a second used by group B to send and for group A to receive.

Since contexts are not explicit objects, other implementations are also possible. (*End of advice to implementors.*)

7.2.3 Intra-Communicators

Intra-communicators bring together the concepts of group and context. To support implementation-specific optimizations, and application topologies (defined in the next chapter, Chapter 8), communicators may also “cache” additional information (see Section 7.7). MPI communication operations reference communicators to determine the scope and the “communication universe” in which a point-to-point or collective operation is to operate.

Each communicator contains a group of valid participants; this group always includes the local MPI process. The source and destination of a message are identified by MPI process ranks within that group.

For collective communication, the intra-communicator specifies the set of MPI processes that participate in the collective operation (and their order, when significant). Thus, the communicator restricts the “spatial” scope of communication, and provides machine-independent MPI process addressing through ranks.

Intra-communicators are represented by opaque **intra-communicator objects**, and hence cannot be directly transferred from one MPI process to another.

7.2.4 Predefined Intra-Communicators

When using the World Model (Section 11.2) for MPI initialization, an initial intra-communicator `MPI_COMM_WORLD` of all MPI processes the local MPI process can communicate with after initialization (itself included) is defined once `MPI_INIT` or `MPI_INIT_THREAD` has been called. In addition, the communicator `MPI_COMM_SELF` is provided, which includes only the MPI process itself. When using the Sessions Model (Section 11.3) for initialization of MPI resources, `MPI_COMM_WORLD` and `MPI_COMM_SELF` are not valid for use as a communicator. See the discussion concerning use of MPI named constants in 2.5.4 for valid uses of `MPI_COMM_WORLD` and `MPI_COMM_SELF` prior to initialization of MPI. See also the discussion concerning interoperability of the World Model and Sessions Model in Section 11.1.

The predefined constant `MPI_COMM_NULL` is the value used for invalid communicator handles.

In a static-process-model implementation of MPI, all MPI processes that participate in the computation are available after MPI is initialized. For this case, `MPI_COMM_WORLD` is a communicator of all MPI processes available for the computation; this communicator has the same value in all MPI processes. In an implementation of MPI where MPI processes can dynamically join an MPI execution, it may be the case that an MPI process starts an MPI computation without having access to all other MPI processes. In such situations, `MPI_COMM_WORLD` is a communicator incorporating all MPI processes with which the joining MPI process can immediately communicate. Therefore, `MPI_COMM_WORLD` may simultaneously represent disjoint groups in different MPI processes.

All MPI implementations are required to provide the `MPI_COMM_WORLD` communicator. It cannot be deallocated during the life of an MPI process. The group corresponding to this communicator does not appear as a pre-defined constant, but it may be accessed using `MPI_COMM_GROUP` (see below). MPI does not specify the correspondence between the MPI process rank in `MPI_COMM_WORLD` and its (machine-dependent) absolute address. Other implementation-dependent, predefined communicators may also be provided.

7.3 Group Management

This section describes the manipulation of MPI process groups. These operations are local.

7.3.1 Group Accessors

MPI_GROUP_SIZE(group, size)

IN	group	group (handle)
OUT	size	number of MPI processes in the group (integer)

C binding

```
int MPI_Group_size(MPI_Group group, int *size)
```

Fortran 2008 binding

```
MPI_Group_size(group, size, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group
  INTEGER, INTENT(OUT) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GROUP_SIZE(GROUP, SIZE, IERROR)
  INTEGER GROUP, SIZE, IERROR
```

MPI_GROUP_RANK(group, rank)

IN	group	group (handle)
OUT	rank	rank of the calling MPI process in group, or MPI_UNDEFINED if the MPI process is not a member (integer)

C binding

```
int MPI_Group_rank(MPI_Group group, int *rank)
```

Fortran 2008 binding

```
MPI_Group_rank(group, rank, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group
  INTEGER, INTENT(OUT) :: rank
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GROUP_RANK(GROUP, RANK, IERROR)
  INTEGER GROUP, RANK, IERROR
```


MPI_GROUP_TRANSLATE_RANKS(group1, n, ranks1, group2, ranks2)			1
IN	group1	group1 (handle)	2
IN	n	number of elements in ranks1 and ranks2 arrays (integer)	3
IN	ranks1	array of zero or more valid ranks in group1	4
IN	group2	group2 (handle)	5
OUT	ranks2	array of corresponding ranks in group2, MPI_UNDEFINED when no correspondence exists.	6

C binding

```
int MPI_Group_translate_ranks(MPI_Group group1, int n, const int ranks1[],
                             MPI_Group group2, int ranks2[])
```

Fortran 2008 binding

```
MPI_Group_translate_ranks(group1, n, ranks1, group2, ranks2, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group1, group2
  INTEGER, INTENT(IN) :: n, ranks1(n)
  INTEGER, INTENT(OUT) :: ranks2(n)
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GROUP_TRANSLATE_RANKS(GROUP1, N, RANKS1, GROUP2, RANKS2, IERROR)
  INTEGER GROUP1, N, RANKS1(*), GROUP2, RANKS2(*), IERROR
```

This function is important for determining the relative numbering of the same MPI processes in two different groups. For instance, if one knows the ranks of certain MPI processes in the group of MPI_COMM_WORLD, one might want to know their ranks in a subset of that group.

MPI_PROC_NULL is a valid rank for input to MPI_GROUP_TRANSLATE_RANKS, which returns MPI_PROC_NULL as the translated rank.

```
MPI_GROUP_COMPARE(group1, group2, result)
```

IN	group1	first group (handle)
IN	group2	second group (handle)
OUT	result	result (integer)

C binding

```
int MPI_Group_compare(MPI_Group group1, MPI_Group group2, int *result)
```

Fortran 2008 binding

```
MPI_Group_compare(group1, group2, result, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group1, group2
  INTEGER, INTENT(OUT) :: result
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GROUP_COMPARE(GROUP1, GROUP2, RESULT, IERROR)
```

1 INTEGER GROUP1, GROUP2, RESULT, IERROR

2
3 MPI_IDENT results if the group members and group order are exactly the same in both
4 groups. This happens for instance if `group1` and `group2` are the same handle. MPI_SIMILAR
5 results if the group members are the same but the order is different. MPI_UNEQUAL results
6 otherwise.

7 8 7.3.2 Group Constructors

9 MPI provides two approaches to constructing groups. In the first approach, MPI procedures
10 are provided to subset and superset existing groups. These constructors construct new
11 groups from existing groups. In the second approach, a group is created using a session
12 handle and associated process set. This second approach is available when using the Ses-
13 sions Model. With both approaches, these are local operations, and distinct groups may
14 be defined on different MPI processes; an MPI process may also define a group that does
15 not include itself. Consistent definitions are required when groups are used as arguments
16 in communicator creation functions. When using the World Model (Section 11.2) for MPI
17 initialization, the base group, upon which all other groups are defined, is the group asso-
18 ciated with the initial communicator MPI_COMM_WORLD (accessible through the function
19 MPI_COMM_GROUP).

20
21 *Rationale.* In what follows, there is no group duplication function analogous to
22 MPI_COMM_DUP, defined later in this chapter. There is no need for a group dupli-
23 cator. A group, once created, can have several references to it by making copies of
24 the handle. The following constructors address the need for subsets and supersets of
25 existing groups. (*End of rationale.*)

26
27 *Advice to implementors.* Each group constructor behaves as if it returned a new
28 group object. When this new group is a copy of an existing group, then one can
29 avoid creating such new objects, using a reference-count mechanism. (*End of advice*
30 *to implementors.*)

31
32
33 MPI_COMM_GROUP(comm, group)

34 IN comm communicator (handle)
35
36 OUT group group corresponding to comm (handle)

37 38 **C binding**

39 int MPI_Comm_group(MPI_Comm comm, MPI_Group *group)

40 41 **Fortran 2008 binding**

42 MPI_Comm_group(comm, group, ierror)
43 TYPE(MPI_Comm), INTENT(IN) :: comm
44 TYPE(MPI_Group), INTENT(OUT) :: group
45 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

46 47 **Fortran binding**

48 MPI_COMM_GROUP(COMM, GROUP, IERROR)
 INTEGER COMM, GROUP, IERROR

MPI_COMM_GROUP returns in group a handle to the group of comm.

MPI_GROUP_UNION(group1, group2, newgroup)

IN	group1	first group (handle)
IN	group2	second group (handle)
OUT	newgroup	union group (handle)

C binding

```
int MPI_Group_union(MPI_Group group1, MPI_Group group2, MPI_Group *newgroup)
```

Fortran 2008 binding

```
MPI_Group_union(group1, group2, newgroup, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group1, group2
  TYPE(MPI_Group), INTENT(OUT) :: newgroup
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GROUP_UNION(GROUP1, GROUP2, NEWGROUP, IERROR)
  INTEGER GROUP1, GROUP2, NEWGROUP, IERROR
```

MPI_GROUP_INTERSECTION(group1, group2, newgroup)

IN	group1	first group (handle)
IN	group2	second group (handle)
OUT	newgroup	intersection group (handle)

C binding

```
int MPI_Group_intersection(MPI_Group group1, MPI_Group group2,
  MPI_Group *newgroup)
```

Fortran 2008 binding

```
MPI_Group_intersection(group1, group2, newgroup, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group1, group2
  TYPE(MPI_Group), INTENT(OUT) :: newgroup
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GROUP_INTERSECTION(GROUP1, GROUP2, NEWGROUP, IERROR)
  INTEGER GROUP1, GROUP2, NEWGROUP, IERROR
```

```

1 MPI_GROUP_DIFFERENCE(group1, group2, newgroup)
2     IN     group1           first group (handle)
3
4     IN     group2           second group (handle)
5
6     OUT    newgroup         difference group (handle)

```

C binding

```

8 int MPI_Group_difference(MPI_Group group1, MPI_Group group2,
9                          MPI_Group *newgroup)

```

Fortran 2008 binding

```

12 MPI_Group_difference(group1, group2, newgroup, ierror)
13     TYPE(MPI_Group), INTENT(IN) :: group1, group2
14     TYPE(MPI_Group), INTENT(OUT) :: newgroup
15     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

17 MPI_GROUP_DIFFERENCE(GROUP1, GROUP2, NEWGROUP, IERROR)
18     INTEGER GROUP1, GROUP2, NEWGROUP, IERROR

```

The set-like operations are defined as follows:

union: All elements of the first group (*group1*), followed by all elements of second group (*group2*) not in the first group.

intersect: All elements of the first group that are also in the second group, ordered as in the first group.

difference All elements of the first group that are not in the second group, ordered as in the first group.

Note that for these operations the order of MPI processes in the output group is determined primarily by order in the first group (if possible) and then, if necessary, by order in the second group. Neither union nor intersection are commutative, but both are associative. The new group can be empty, that is, equal to `MPI_GROUP_EMPTY`.

```

35 MPI_GROUP_INCL(group, n, ranks, newgroup)

```

```

36     IN     group           group (handle)
37
38     IN     n               number of elements in array ranks (and size of
39                          newgroup) (integer)
40
41     IN     ranks           ranks of processes in group to appear in newgroup
42                          (array of integers)
43
44     OUT    newgroup         new group derived from above, in the order defined
45                          by ranks (handle)

```

C binding

```

46 int MPI_Group_incl(MPI_Group group, int n, const int ranks[],
47                   MPI_Group *newgroup)

```

Fortran 2008 binding

```

MPI_Group_incl(group, n, ranks, newgroup, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group
  INTEGER, INTENT(IN) :: n, ranks(n)
  TYPE(MPI_Group), INTENT(OUT) :: newgroup
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_GROUP_INCL(GROUP, N, RANKS, NEWGROUP, IERROR)
  INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR

```

The function `MPI_GROUP_INCL` creates a group `newgroup` that consists of the `n` MPI processes in `group` with ranks `ranks[0], ..., ranks[n-1]`; the MPI process with rank `i` in `newgroup` is the MPI process with rank `ranks[i]` in `group`. Each of the `n` elements of `ranks` must be a valid rank in `group` and all elements must be distinct, or else the program is erroneous. If `n = 0`, then `newgroup` is `MPI_GROUP_EMPTY`. This function can, for instance, be used to reorder the elements of a group. See also `MPI_GROUP_COMPARE`.

```

MPI_GROUP_EXCL(group, n, ranks, newgroup)

```

IN	group	group (handle)
IN	n	number of elements in array <code>ranks</code> (integer)
IN	ranks	array of integer ranks of MPI processes in <code>group</code> not to appear in <code>newgroup</code>
OUT	newgroup	new group derived from above, preserving the order defined by <code>group</code> (handle)

C binding

```

int MPI_Group_excl(MPI_Group group, int n, const int ranks[],
  MPI_Group *newgroup)

```

Fortran 2008 binding

```

MPI_Group_excl(group, n, ranks, newgroup, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group
  INTEGER, INTENT(IN) :: n, ranks(n)
  TYPE(MPI_Group), INTENT(OUT) :: newgroup
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_GROUP_EXCL(GROUP, N, RANKS, NEWGROUP, IERROR)
  INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR

```

The function `MPI_GROUP_EXCL` creates a group of MPI processes `newgroup` that is obtained by deleting from `group` those MPI processes with ranks `ranks[0], ..., ranks[n-1]`. The ordering of MPI processes in `newgroup` is identical to the ordering in `group`. Each of the `n` elements of `ranks` must be a valid rank in `group` and all elements must be distinct; otherwise, the program is erroneous. If `n = 0`, then `newgroup` is identical to `group`.

```

1 MPI_GROUP_RANGE_INCL(group, n, ranges, newgroup)
2   IN      group          group (handle)
3
4   IN      n              number of triplets in array ranges (integer)
5
6   IN      ranges         a one-dimensional array of integer triplets, of the
7                          form (first rank, last rank, stride) indicating ranks in
8                          group of MPI processes to be included in newgroup
9
10  OUT     newgroup       new group derived from above, in the order defined
11                          by ranges (handle)

```

C binding

```

12 int MPI_Group_range_incl(MPI_Group group, int n, int ranges[][3],
13                          MPI_Group *newgroup)

```

Fortran 2008 binding

```

15 MPI_Group_range_incl(group, n, ranges, newgroup, ierror)
16   TYPE(MPI_Group), INTENT(IN) :: group
17   INTEGER, INTENT(IN) :: n, ranges(3, n)
18   TYPE(MPI_Group), INTENT(OUT) :: newgroup
19   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

21 MPI_GROUP_RANGE_INCL(GROUP, N, RANGES, NEWGROUP, IERROR)
22   INTEGER GROUP, N, RANGES(3, *), NEWGROUP, IERROR

```

If ranges consists of the triplets

$$(first_1, last_1, stride_1), \dots, (first_n, last_n, stride_n)$$

then newgroup consists of the sequence of MPI processes in group with ranks

$$first_1, first_1 + stride_1, \dots, first_1 + \left\lfloor \frac{last_1 - first_1}{stride_1} \right\rfloor stride_1, \dots,$$

$$first_n, first_n + stride_n, \dots, first_n + \left\lfloor \frac{last_n - first_n}{stride_n} \right\rfloor stride_n.$$

Each computed rank must be a valid rank in group and all computed ranks must be distinct, or else the program is erroneous. Note that we may have $first_i > last_i$, and $stride_i$ may be negative, but cannot be zero.

The functionality of this routine is specified to be equivalent to expanding the array of ranges to an array of the included ranks and passing the resulting array of ranks and other arguments to MPI_GROUP_INCL. A call to MPI_GROUP_INCL is equivalent to a call to MPI_GROUP_RANGE_INCL with each rank i in ranks replaced by the triplet $(i,i,1)$ in the argument ranges.

MPI_GROUP_RANGE_EXCL(group, n, ranges, newgroup)			1
IN	group	group (handle)	2
IN	n	number of triplets in array <code>ranges</code> (integer)	3
IN	ranges	a one-dimensional array of integer triplets, of the form (first rank, last rank, stride) indicating ranks in group of MPI processes to be excluded from the output group <code>newgroup</code> (array of integers)	4
OUT	newgroup	new group derived from above, preserving the order in group (handle)	5

C binding

```
int MPI_Group_range_excl(MPI_Group group, int n, int ranges[][3],
                        MPI_Group *newgroup)
```

Fortran 2008 binding

```
MPI_Group_range_excl(group, n, ranges, newgroup, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group
  INTEGER, INTENT(IN) :: n, ranges(3, n)
  TYPE(MPI_Group), INTENT(OUT) :: newgroup
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GROUP_RANGE_EXCL(GROUP, N, RANGES, NEWGROUP, IERROR)
  INTEGER GROUP, N, RANGES(3, *), NEWGROUP, IERROR
```

Each computed rank must be a valid rank in `group` and all computed ranks must be distinct, or else the program is erroneous.

The functionality of this routine is specified to be equivalent to expanding the array of `ranges` to an array of the excluded ranks and passing the resulting array of ranks and other arguments to `MPI_GROUP_EXCL`. A call to `MPI_GROUP_EXCL` is equivalent to a call to `MPI_GROUP_RANGE_EXCL` with each rank `i` in `ranks` replaced by the triplet `(i,i,1)` in the argument `ranges`.

Advice to users. The range operations do not explicitly enumerate ranks, and therefore are more scalable if implemented efficiently. Hence, we recommend MPI programmers to use them whenever possible, as high-quality implementations will take advantage of this fact. (*End of advice to users.*)

Advice to implementors. The range operations should be implemented, if possible, without enumerating the group members, in order to obtain better scalability (time and space). (*End of advice to implementors.*)

```

1 MPI_GROUP_FROM_SESSION_PSET(session, pset_name, newgroup)
2     IN      session          session (handle)
3
4     IN      pset_name       name of process set to use to create the new group
5                               (string)
6
7     OUT     newgroup        new group derived from supplied session and process
8                               set (handle)

```

C binding

```

10 int MPI_Group_from_session_pset(MPI_Session session, const char *pset_name,
11                               MPI_Group *newgroup)

```

Fortran 2008 binding

```

13 MPI_Group_from_session_pset(session, pset_name, newgroup, ierror)
14     TYPE(MPI_Session), INTENT(IN) :: session
15     CHARACTER(LEN=*), INTENT(IN) :: pset_name
16     TYPE(MPI_Group), INTENT(OUT) :: newgroup
17     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

19 MPI_GROUP_FROM_SESSION_PSET(SESSION, PSET_NAME, NEWGROUP, IERROR)
20     INTEGER SESSION, NEWGROUP, IERROR
21     CHARACTER*(*) PSET_NAME

```

The function `MPI_GROUP_FROM_SESSION_PSET` creates a group `newgroup` using the provided session handle and process set. The process set name must be one returned from an invocation of `MPI_SESSION_GET_NTH_PSET` using the supplied session handle. If the `pset_name` does not exist, `MPI_GROUP_NULL` will be returned in the `newgroup` argument. As with other group constructors, `MPI_GROUP_FROM_SESSION_PSET` is a local function. See Section 11.3 for more information on sessions and process sets.

7.3.3 Group Destructors

```

34 MPI_GROUP_FREE(group)
35     INOUT  group          group (handle)

```

C binding

```

38 int MPI_Group_free(MPI_Group *group)

```

Fortran 2008 binding

```

40 MPI_Group_free(group, ierror)
41     TYPE(MPI_Group), INTENT(INOUT) :: group
42     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

44 MPI_GROUP_FREE(GROUP, IERROR)
45     INTEGER GROUP, IERROR

```


This operation marks a group object for deallocation. The handle `group` is set to `MPI_GROUP_NULL` by the call. Any on-going operation using this group will complete normally.

Advice to implementors. One can keep a reference count that is incremented for each call to `MPI_COMM_GROUP`, `MPI_COMM_CREATE`, `MPI_COMM_DUP`, `MPI_COMM_IDUP`, `MPI_COMM_DUP_WITH_INFO`, `MPI_COMM_IDUP_WITH_INFO`, `MPI_COMM_SPLIT`, `MPI_COMM_SPLIT_TYPE`, `MPI_COMM_CREATE_GROUP`, `MPI_COMM_CREATE_FROM_GROUP`, `MPI_INTERCOMM_CREATE`, and `MPI_INTERCOMM_CREATE_FROM_GROUPS`, and decremented for each call to `MPI_GROUP_FREE` or `MPI_COMM_FREE`; the group object is ultimately deallocated when the reference count drops to zero. (*End of advice to implementors.*)

7.4 Communicator Management

This section describes the manipulation of communicators in MPI. Operations that access communicators are local. Operations that create communicators are collective.

Advice to implementors. High-quality implementations should amortize the overheads associated with the creation of communicators (for the same group, or subsets thereof) over several calls, by allocating multiple contexts with one collective communication. (*End of advice to implementors.*)

7.4.1 Communicator Accessors

The following are all local operations.

`MPI_COMM_SIZE(comm, size)`

IN	<code>comm</code>	communicator (handle)
OUT	<code>size</code>	number of MPI processes in the group of <code>comm</code> (integer)

C binding

```
int MPI_Comm_size(MPI_Comm comm, int *size)
```

Fortran 2008 binding

```
MPI_Comm_size(comm, size, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(OUT) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_SIZE(COMM, SIZE, IERROR)
  INTEGER COMM, SIZE, IERROR
```

Rationale. This function is equivalent to accessing the communicator's group with `MPI_COMM_GROUP` (see above), computing the size using `MPI_GROUP_SIZE`, and then freeing the temporary group via `MPI_GROUP_FREE`. However, this functionality is so commonly used that this shortcut was introduced. (*End of rationale.*)

Advice to users. This function indicates the number of MPI processes involved in a communicator. For `MPI_COMM_WORLD`, it indicates the total number of MPI processes available unless the number of MPI processes has been changed by using the functions described in Chapter 11; note that the number of MPI processes in `MPI_COMM_WORLD` does not change during the life of an MPI program.

This call is often used with the next call to determine the amount of concurrency available for a specific library or program. The following call, `MPI_COMM_RANK` indicates the rank of the MPI process that calls it in the range from 0, ..., `size-1`, where `size` is the return value of `MPI_COMM_SIZE`. (*End of advice to users.*)

`MPI_COMM_RANK(comm, rank)`

IN	<code>comm</code>	communicator (handle)
OUT	<code>rank</code>	rank of the calling MPI process in group of <code>comm</code> (integer)

C binding

```
int MPI_Comm_rank(MPI_Comm comm, int *rank)
```

Fortran 2008 binding

```
MPI_Comm_rank(comm, rank, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(OUT) :: rank
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_RANK(COMM, RANK, IERROR)
  INTEGER COMM, RANK, IERROR
```

Rationale. This function is equivalent to accessing the communicator's group with `MPI_COMM_GROUP` (see above), computing the rank using `MPI_GROUP_RANK`, and then freeing the temporary group via `MPI_GROUP_FREE`. However, this functionality is so commonly used that this shortcut was introduced. (*End of rationale.*)

Advice to users. This function gives the rank of the MPI process in the particular communicator's group. It is useful, as noted above, in conjunction with `MPI_COMM_SIZE`.

Many programs will follow the supervisor/executor or manager/worker model, where one MPI process will play a supervisory role while the other MPI processes will play an executory role. In this framework, the two preceding calls are useful for determining the roles of the various MPI processes of a communicator. (*End of advice to users.*)

MPI_COMM_COMPARE(comm1, comm2, result)			1
IN	comm1	first communicator (handle)	2
IN	comm2	second communicator (handle)	3
OUT	result	result (integer)	4
			5
			6

C binding

```
int MPI_Comm_compare(MPI_Comm comm1, MPI_Comm comm2, int *result)
```

Fortran 2008 binding

```
MPI_Comm_compare(comm1, comm2, result, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm1, comm2
  INTEGER, INTENT(OUT) :: result
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_COMPARE(COMM1, COMM2, RESULT, IERROR)
  INTEGER COMM1, COMM2, RESULT, IERROR
```

MPI_IDENT results if and only if `comm1` and `comm2` are handles for the same object (identical groups and same contexts). MPI_CONGRUENT results if the underlying groups are identical in constituents and rank order; these communicators differ only by context. MPI_SIMILAR results if the group members of both communicators are the same but the rank order differs. MPI_UNEQUAL results otherwise.

7.4.2 Communicator Constructors

The following are collective functions that are invoked by all MPI processes in the group or groups associated with `comm`, with the exception of `MPI_COMM_CREATE_GROUP`, `MPI_COMM_CREATE_FROM_GROUP`, and `MPI_INTERCOMM_CREATE_FROM_GROUPS`. `MPI_COMM_CREATE_GROUP` and `MPI_COMM_CREATE_FROM_GROUP` are invoked only by the MPI processes in the group of the new communicator being constructed. `MPI_INTERCOMM_CREATE_FROM_GROUPS` is invoked by all the MPI processes in the local and remote groups of the new communicator being constructed. See the discussion below for the definition of local and remote groups.

Rationale. Note that, when using the World Model, there is a chicken-and-egg aspect to MPI in that a communicator is needed to create a new communicator. In the World Model, the base communicator for all MPI communicators is predefined outside of MPI, and is `MPI_COMM_WORLD`. The World Model was arrived at after considerable debate, and was chosen to increase “safety” of programs written in MPI. (*End of rationale.*)

This chapter presents the following communicator construction routines:

```
MPI_COMM_CREATE, MPI_COMM_DUP, MPI_COMM_IDUP,
MPI_COMM_DUP_WITH_INFO, MPI_COMM_IDUP_WITH_INFO, MPI_COMM_SPLIT
and MPI_COMM_SPLIT_TYPE can be used to create both intra-communicators and inter-communicators; MPI_COMM_CREATE_GROUP, MPI_COMM_CREATE_FROM_GROUP and MPI_INTERCOMM_MERGE (see Section 7.6.2) can be used to create intra-communicators;
```

1 MPI_INTERCOMM_CREATE and MPI_INTERCOMM_CREATE_FROM_GROUPS (see Section
2 7.6.2) can be used to create inter-communicators.

3 An intra-communicator involves a single group while an inter-communicator involves
4 two groups. Where the following discussions address inter-communicator semantics, the
5 two groups in an inter-communicator are called the *left* and *right* groups. An MPI process
6 in an inter-communicator is a member of either the left or the right group. From the point
7 of view of that MPI process, the group that the MPI process is a member of is called the
8 *local group*; the other group (relative to that MPI process) is the *remote group*. The left
9 and right group labels give us a way to describe the two groups in an inter-communicator
10 that is not relative to any particular MPI process (as the local and remote groups are).
11

12 MPI_COMM_DUP(comm, newcomm)

14	IN	comm	communicator (handle)
15	OUT	newcomm	copy of comm (handle)

17 C binding

18 int MPI_Comm_dup(MPI_Comm comm, MPI_Comm *newcomm)

20 Fortran 2008 binding

21 MPI_Comm_dup(comm, newcomm, ierror)
22 TYPE(MPI_Comm), INTENT(IN) :: comm
23 TYPE(MPI_Comm), INTENT(OUT) :: newcomm
24 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

25 Fortran binding

26 MPI_COMM_DUP(COMM, NEWCOMM, IERROR)
27 INTEGER COMM, NEWCOMM, IERROR

28 MPI_COMM_DUP duplicates the existing communicator `comm` with associated key
29 values, topology information and error handlers. For each key value, the respective copy
30 callback function determines the attribute value associated with this key in the new com-
31 municator; one particular action that a copy callback may take is to delete the attribute
32 from the new communicator. MPI_COMM_DUP returns in `newcomm` a new communicator
33 with the same group or groups, same topology, same error handlers and any copied cached
34 information, but a new context (see Section 7.7.1). The newly created communicator will
35 have no buffer attached (see Section 3.6).
36

37
38 *Advice to users.* This operation is used to provide a parallel library with a duplicate
39 communication space that has the same properties as the original communicator.
40 This includes any attributes (see below) and topologies (see Chapter 8). This call is
41 valid even if there are *pending* point-to-point communication operations or *decoupled*
42 *MPI activities* involving the communicator `comm`. A typical call might involve a
43 MPI_COMM_DUP at the beginning of the parallel call, and an MPI_COMM_FREE of
44 that duplicated communicator at the end of the call. Other models of communicator
45 management are also possible.

46 This call applies to both intra- and inter-communicators. (*End of advice to users.*)
47
48

Advice to implementors. One need not actually copy the group information, but only add a new reference and increment the reference count. Copy on write can be used for the cached information. (*End of advice to implementors.*)

MPI_COMM_DUP_WITH_INFO(comm, info, newcomm)

IN	comm	communicator (handle)
IN	info	info object (handle)
OUT	newcomm	copy of comm (handle)

C binding

int MPI_Comm_dup_with_info(MPI_Comm comm, MPI_Info info, MPI_Comm *newcomm)

Fortran 2008 binding

```
MPI_Comm_dup_with_info(comm, info, newcomm, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(MPI_Comm), INTENT(OUT) :: newcomm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_DUP_WITH_INFO(COMM, INFO, NEWCOMM, IERROR)
  INTEGER COMM, INFO, NEWCOMM, IERROR
```

MPI_COMM_DUP_WITH_INFO behaves exactly as MPI_COMM_DUP except that the hints provided by the argument `info` are associated with the output communicator `newcomm`.

Rationale. It is expected that some hints will only be valid at communicator creation time. However, for legacy reasons, most communicator creation calls do not provide an `info` argument. One may associate `info` hints with a duplicate of any communicator at creation time through a call to MPI_COMM_DUP_WITH_INFO. (*End of rationale.*)

MPI_COMM_IDUP(comm, newcomm, request)

IN	comm	communicator (handle)
OUT	newcomm	copy of comm (handle)
OUT	request	communication request (handle)

C binding

int MPI_Comm_idup(MPI_Comm comm, MPI_Comm *newcomm, MPI_Request *request)

Fortran 2008 binding

```
MPI_Comm_idup(comm, newcomm, request, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Comm), INTENT(OUT), ASYNCHRONOUS :: newcomm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```

MPI_COMM_IDUP(COMM, NEWCOMM, REQUEST, IERROR)
  INTEGER COMM, NEWCOMM, REQUEST, IERROR

```

MPI_COMM_IDUP is a nonblocking variant of MPI_COMM_DUP. With the exception of its nonblocking behavior, the semantics of MPI_COMM_IDUP are as if MPI_COMM_DUP was executed at the time that MPI_COMM_IDUP is called. For example, attributes changed after MPI_COMM_IDUP will not be copied to the new communicator. All restrictions and assumptions for nonblocking collective operations (see Section 6.12) apply to MPI_COMM_IDUP and the returned request.

It is erroneous to use the communicator `newcomm` as an input argument to other MPI functions before the MPI_COMM_IDUP operation completes.

```

MPI_COMM_IDUP_WITH_INFO(comm, info, newcomm, request)

```

IN	comm	communicator (handle)
IN	info	info object (handle)
OUT	newcomm	copy of comm (handle)
OUT	request	communication request (handle)

C binding

```

int MPI_Comm_idup_with_info(MPI_Comm comm, MPI_Info info, MPI_Comm *newcomm,
  MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Comm_idup_with_info(comm, info, newcomm, request, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(MPI_Comm), INTENT(OUT), ASYNCHRONOUS :: newcomm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_COMM_IDUP_WITH_INFO(COMM, INFO, NEWCOMM, REQUEST, IERROR)
  INTEGER COMM, INFO, NEWCOMM, REQUEST, IERROR

```

MPI_COMM_IDUP_WITH_INFO is a nonblocking variant of MPI_COMM_DUP_WITH_INFO. With the exception of its nonblocking behavior, the semantics of MPI_COMM_IDUP_WITH_INFO are as if MPI_COMM_DUP_WITH_INFO was executed at the time that MPI_COMM_IDUP_WITH_INFO is called. For example, attributes or info hints changed after MPI_COMM_IDUP_WITH_INFO will not be copied to the new communicator. All restrictions and assumptions for nonblocking collective operations (see Section 6.12) apply to MPI_COMM_IDUP_WITH_INFO and the returned request.

It is erroneous to use the communicator `newcomm` as an input argument to other MPI functions before the MPI_COMM_IDUP_WITH_INFO operation completes.

Rationale. The MPI_COMM_IDUP and MPI_COMM_IDUP_WITH_INFO functions are crucial for the development of purely nonblocking libraries (see [41]). (*End of rationale.*)

MPI_COMM_CREATE(comm, group, newcomm)			1
IN	comm	communicator (handle)	2
IN	group	group, which is a subset of the group of comm (handle)	3
OUT	newcomm	new communicator (handle)	4
			5
			6
			7

C binding

```
int MPI_Comm_create(MPI_Comm comm, MPI_Group group, MPI_Comm *newcomm)
```

Fortran 2008 binding

```
MPI_Comm_create(comm, group, newcomm, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Group), INTENT(IN) :: group
  TYPE(MPI_Comm), INTENT(OUT) :: newcomm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

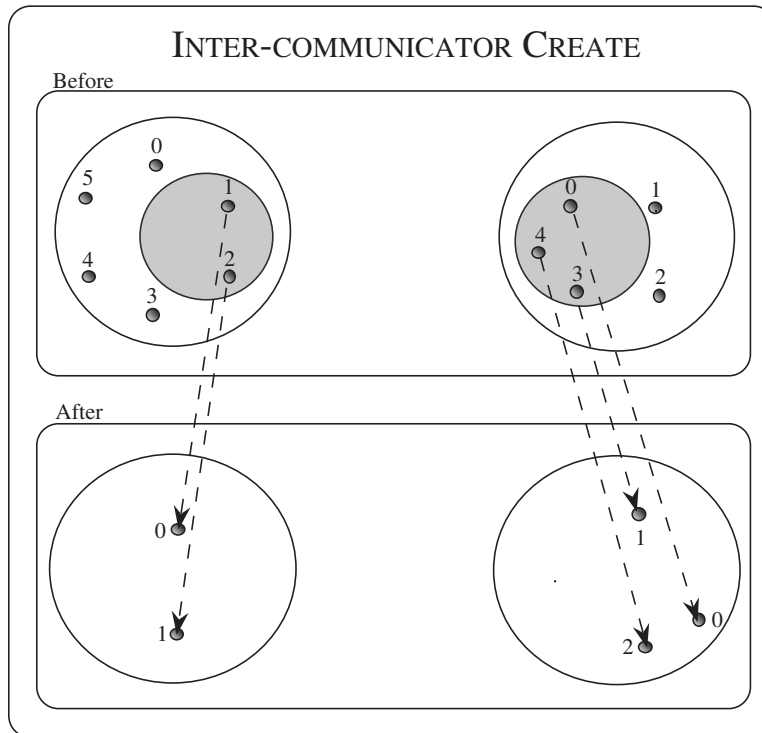
```
MPI_COMM_CREATE(COMM, GROUP, NEWCOMM, IERROR)
  INTEGER COMM, GROUP, NEWCOMM, IERROR
```

If `comm` is an intra-communicator, this function returns a new communicator `newcomm` with communication group defined by the `group` argument. No cached information propagates from `comm` to `newcomm` and no virtual topology information is added to the created communicator. Each MPI process must call `MPI_COMM_CREATE` with a `group` argument that is a subgroup of the group associated with `comm`; this could be `MPI_GROUP_EMPTY`. The MPI processes may specify different values for the `group` argument. If an MPI process calls with a nonempty `group` then all MPI processes in that `group` must call the function with the same `group` as argument, that is the same MPI processes in the same order. Otherwise, the call is erroneous. This implies that the set of groups specified across the MPI processes must be disjoint. If the calling MPI process is a member of the group given as `group` argument, then `newcomm` is a communicator with `group` as its associated group. In the case that an MPI process calls with a `group` to which it does not belong, e.g., `MPI_GROUP_EMPTY`, then `MPI_COMM_NULL` is returned as `newcomm`. The function is collective and must be called by all MPI processes in the group of `comm`.

Rationale. The interface supports the original mechanism from MPI-1.1, which required the same `group` in all MPI processes of `comm`. It was extended in MPI-2.2 to allow the use of disjoint subgroups in order to allow implementations to eliminate unnecessary communication that `MPI_COMM_SPLIT` would incur when the user already knows the membership of the disjoint subgroups. (*End of rationale.*)

Rationale. The requirement that the entire group of `comm` participate in the call stems from the following considerations:

- It allows the implementation to layer `MPI_COMM_CREATE` on top of regular collective communications.
- It provides additional safety, in particular in the case where partially overlapping groups are used to create new communicators.



22 Figure 7.1: Inter-communicator creation using `MPI_COMM_CREATE` extended to inter-communicators. The input groups are those in the grey circle.

- 26 • It permits implementations to sometimes avoid communication related to context creation.

28 (*End of rationale.*)

30 *Advice to users.* `MPI_COMM_CREATE` provides a means to subset a group of MPI processes for the purpose of separate MIMD computation, with separate communication space. `newcomm`, which emerges from `MPI_COMM_CREATE`, can be used in subsequent calls to `MPI_COMM_CREATE` (or other communicator constructors) to further subdivide a computation into parallel sub-computations. A more general service is provided by `MPI_COMM_SPLIT`, below. (*End of advice to users.*)

37 *Advice to implementors.* When calling `MPI_COMM_DUP`, all MPI processes call with the same `group` (the `group` associated with the communicator). When calling `MPI_COMM_CREATE`, the MPI processes provide the same `group` or disjoint subgroups. For both calls, it is theoretically possible to agree on a group-wide unique context with no communication. However, local execution of these functions requires use of a larger context name space and reduces error checking. Implementations may strike various compromises between these conflicting goals, such as bulk allocation of multiple contexts in one collective operation.

45 Important: If new communicators are created without synchronizing the MPI processes involved then the communication system must be able to cope with messages arriving in a context that has not yet been allocated at the receiving MPI process. (*End of advice to implementors.*)

If `comm` is an inter-communicator, then the output communicator is also an inter-communicator where the local group consists only of those MPI processes contained in `group` (see Figure 7.1). The `group` argument should only contain those MPI processes in the local group of the input inter-communicator that are to be a part of `newcomm`. All MPI processes in the same local group of `comm` must specify the same value for `group`, i.e., the same members in the same order. If either `group` does not specify at least one MPI process in the local group of the inter-communicator, or if the calling MPI process is not included in the `group`, `MPI_COMM_NULL` is returned.

Rationale. In the case where either the left or right group is empty, a null communicator is returned instead of an inter-communicator with `MPI_GROUP_EMPTY` because the side with the empty group must return `MPI_COMM_NULL`. (*End of rationale.*)

Example 7.1. Inter-communicator creation.

The following example illustrates how the first node in the left side of an inter-communicator could be joined with all members on the right side of an inter-communicator to form a new inter-communicator.

```
MPI_Comm  inter_comm, new_inter_comm;
MPI_Group local_group, group;
int       rank = 0; /* rank on left side to include in
                    new inter-comm */

/* Construct the original inter-communicator: "inter_comm" */
...

/* Construct the group of MPI processes to be in new
inter-communicator */
if (/* I'm on the left side of the inter-communicator */) {
    MPI_Comm_group(inter_comm, &local_group);
    MPI_Group_incl(local_group, 1, &rank, &group);
    MPI_Group_free(&local_group);
}
else
    MPI_Comm_group(inter_comm, &group);

MPI_Comm_create(inter_comm, group, &new_inter_comm);
MPI_Group_free(&group);
```

`MPI_COMM_CREATE_GROUP(comm, group, tag, newcomm)`

IN	<code>comm</code>	intra-communicator (handle)
IN	<code>group</code>	group, which is a subset of the group of <code>comm</code> (handle)
IN	<code>tag</code>	tag (integer)
OUT	<code>newcomm</code>	new communicator (handle)

C binding

```
int MPI_Comm_create_group(MPI_Comm comm, MPI_Group group, int tag,
```

```

1           MPI_Comm *newcomm)
2
3 Fortran 2008 binding
4 MPI_Comm_create_group(comm, group, tag, newcomm, ierror)
5     TYPE(MPI_Comm), INTENT(IN) :: comm
6     TYPE(MPI_Group), INTENT(IN) :: group
7     INTEGER, INTENT(IN) :: tag
8     TYPE(MPI_Comm), INTENT(OUT) :: newcomm
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

10 MPI_COMM_CREATE_GROUP(COMM, GROUP, TAG, NEWCOMM, IERROR)
11     INTEGER COMM, GROUP, TAG, NEWCOMM, IERROR

```

13 MPI_COMM_CREATE_GROUP is similar to MPI_COMM_CREATE; however,
14 MPI_COMM_CREATE must be called by all MPI processes in the group of `comm`, whereas
15 MPI_COMM_CREATE_GROUP must be called by all MPI processes in `group`, which is
16 a subgroup of the group of `comm`. In addition, MPI_COMM_CREATE_GROUP requires
17 that `comm` is an intra-communicator. MPI_COMM_CREATE_GROUP returns a new intra-
18 communicator, `newcomm`, for which the `group` argument defines the communication group.
19 No cached information propagates from `comm` to `newcomm` and no virtual topology infor-
20 mation is added to the created communicator. Each MPI process must provide a `group`
21 argument that is a subgroup of the group associated with `comm`; this could be
22 MPI_GROUP_EMPTY. If a nonempty group is specified, then all MPI processes in that group
23 must call the function, and each of these MPI processes must provide the same arguments,
24 including a group that contains the same members with the same ordering. Otherwise the
25 call is erroneous. If the calling MPI process is a member of the group given as the `group`
26 argument, then `newcomm` is a communicator with `group` as its associated group. If the
27 calling MPI process is not a member of `group`, e.g., `group` is MPI_GROUP_EMPTY, then the
28 call is a local operation and MPI_COMM_NULL is returned as `newcomm`.

30 *Rationale.* Functionality similar to MPI_COMM_CREATE_GROUP can be imple-
31 mented through repeated MPI_INTERCOMM_CREATE and
32 MPI_INTERCOMM_MERGE calls that start with the MPI_COMM_SELF communica-
33 tors at each MPI process in `group` and build up an intra-communicator with `group`
34 `group` [18]. Such an algorithm requires the creation of many intermediate communica-
35 tors; MPI_COMM_CREATE_GROUP can provide a more efficient implementation that
36 avoids this overhead. (*End of rationale.*)

37 *Advice to users.* An inter-communicator can be created collectively over MPI pro-
38 cesses in the union of the local and remote groups by creating the local communicator
39 using MPI_COMM_CREATE_GROUP and using that communicator as the local com-
40 municator argument to MPI_INTERCOMM_CREATE. (*End of advice to users.*)

42 The `tag` argument does not conflict with tags used in point-to-point communication
43 and is not permitted to be a wildcard. If multiple threads at a given MPI process per-
44 form concurrent MPI_COMM_CREATE_GROUP operations, the user must distinguish these
45 operations by providing different `tag` or `comm` arguments.

46 *Advice to users.* MPI_COMM_CREATE may provide lower overhead than
47 MPI_COMM_CREATE_GROUP because it can take advantage of collective communi-
48 cation on `comm` when constructing `newcomm`. (*End of advice to users.*)

MPI_COMM_SPLIT(comm, color, key, newcomm)			1
IN	comm	communicator (handle)	2
IN	color	control of subset assignment (integer)	3
IN	key	control of rank assignment (integer)	4
OUT	newcomm	new communicator (handle)	5
			6
			7
			8

C binding

```
int MPI_Comm_split(MPI_Comm comm, int color, int key, MPI_Comm *newcomm)
```

Fortran 2008 binding

```
MPI_Comm_split(comm, color, key, newcomm, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: color, key
  TYPE(MPI_Comm), INTENT(OUT) :: newcomm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_SPLIT(COMM, COLOR, KEY, NEWCOMM, IERROR)
  INTEGER COMM, COLOR, KEY, NEWCOMM, IERROR
```

This function partitions the group associated with `comm` into disjoint subgroups, one for each value of `color`. Each subgroup contains all MPI processes of the same color. Within each subgroup, the MPI processes are ranked in the order defined by the value of the argument `key`, with ties broken according to their rank in the old group. A new communicator is created for each subgroup and returned in `newcomm`. An MPI process may supply the color value `MPI_UNDEFINED`, in which case `newcomm` returns `MPI_COMM_NULL`. This is a collective call, but each MPI process is permitted to provide different values for `color` and `key`. No cached information propagates from `comm` to `newcomm` and no virtual topology information is added to the created communicators.

With an intra-communicator `comm`, a call to `MPI_COMM_CREATE(comm, group, newcomm)` is equivalent to a call to `MPI_COMM_SPLIT(comm, color, key, newcomm)`, where MPI processes that are members of their `group` argument provide a `color` argument equal to the number of the `group` (based on a unique numbering of all disjoint groups) and a `key` argument equal to their rank in `group`, and all MPI processes that are not members of their `group` argument provide a `color` argument equal to `MPI_UNDEFINED`. The value of `color` must be nonnegative or `MPI_UNDEFINED`.

Advice to users. This is an extremely powerful mechanism for dividing a single communicating group of MPI processes into k subgroups, with k chosen implicitly by the user (by the number of colors asserted over all the MPI processes). Each resulting communicator will be nonoverlapping. Such a division could be useful for defining a hierarchy of computations, such as for multigrid, or linear algebra. For intra-communicators, `MPI_COMM_SPLIT` provides similar capability as `MPI_COMM_CREATE` to split a communicating group into disjoint subgroups. `MPI_COMM_SPLIT` is useful when some MPI processes do not have complete information of the other members in their group, but all MPI processes know (the color of) the group to which they belong. In this case, the MPI implementation discovers the other group members via communication. `MPI_COMM_CREATE` is useful when all MPI processes have complete

1 information of the members of their group. In this case, MPI can avoid the extra com-
 2 munication required to discover group membership. `MPI_COMM_CREATE_GROUP`
 3 is useful when all MPI processes in a given group have complete information of the
 4 members of their group and synchronization with MPI processes outside the group
 5 can be avoided.

6 Multiple calls to `MPI_COMM_SPLIT` can be used to overcome the requirement that
 7 any call have no overlap of the resulting communicators (each MPI process is of only
 8 one color per call). In this way, multiple overlapping communication structures can be
 9 created. Creative use of the `color` and `key` in such splitting operations is encouraged.
 10

11 Note that, for a fixed color, the keys need not be unique. It is `MPI_COMM_SPLIT`'s
 12 responsibility to sort MPI processes in ascending order according to this key, and to
 13 break ties in a consistent way. If all the keys are specified in the same way, then all
 14 the MPI processes in a given color will have the relative rank order as they did in their
 15 parent group.

16 (*End of advice to users.*)

17
 18 *Rationale.* `color` is restricted to be nonnegative, so as not to conflict with the value
 19 assigned to `MPI_UNDEFINED`. (*End of rationale.*)

20 The result of `MPI_COMM_SPLIT` on an inter-communicator is that those MPI processes on
 21 the left with the same color as those MPI processes on the right combine to create a new
 22 inter-communicator. The `key` argument describes the relative rank of MPI processes on each
 23 side of the inter-communicator (see Figure 7.2). For those colors that are specified only on
 24 one side of the inter-communicator, `MPI_COMM_NULL` is returned. `MPI_COMM_NULL` is also
 25 returned to those MPI processes that specify `MPI_UNDEFINED` as the color.
 26

27 *Advice to users.* For inter-communicators, `MPI_COMM_SPLIT` is more general than
 28 `MPI_COMM_CREATE`. A single call to `MPI_COMM_SPLIT` can create a set of disjoint
 29 inter-communicators, while a call to `MPI_COMM_CREATE` creates only one. (*End of*
 30 *advice to users.*)
 31

32 **Example 7.2.** Parallel client-server model.

33 The following client code illustrates how clients on the left side of an inter-communicator
 34 could be assigned to a single server from a pool of servers on the right side of an inter-
 35 communicator.
 36

```

37 /* Client code */
38 MPI_Comm multiple_server_comm;
39 MPI_Comm single_server_comm;
40 int color, rank, num_servers;
41
42 /* Create inter-communicator with clients and servers:
43 multiple_server_comm */
44 ...
45
46 /* Find out the number of servers available */
47 MPI_Comm_remote_size(multiple_server_comm, &num_servers);
48
49 /* Determine my color */

```

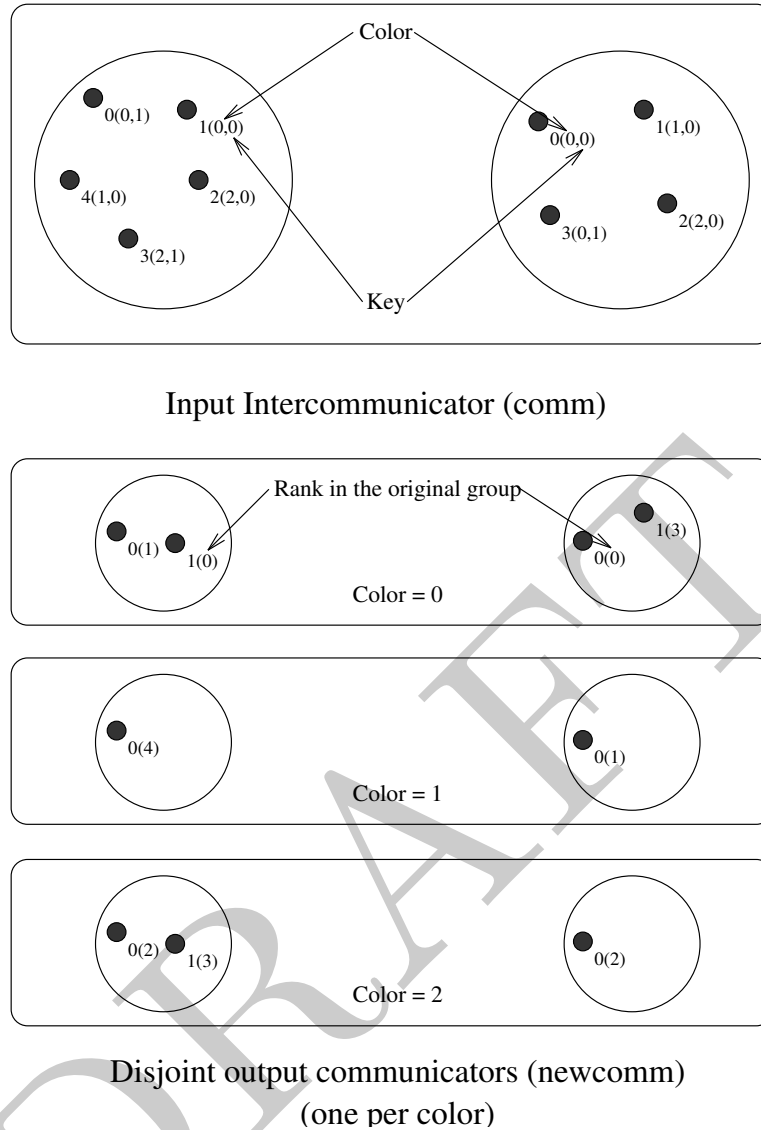


Figure 7.2: Inter-communicator construction achieved by splitting an existing inter-communicator with `MPI_COMM_SPLIT` extended to inter-communicators.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

1 MPI_Comm_rank(multiple_server_comm, &rank);
2 color = rank % num_servers;
3
4 /* Split the inter-communicator */
5 MPI_Comm_split(multiple_server_comm, color, rank,
6                &single_server_comm);

```

The following is the corresponding server code:

```

8 /* Server code */
9 MPI_Comm multiple_client_comm;
10 MPI_Comm single_server_comm;
11 int rank;
12
13 /* Create inter-communicator with clients and servers:
14 multiple_client_comm */
15 ...
16
17 /* Split the inter-communicator for a single server per group
18 of clients */
19 MPI_Comm_rank(multiple_client_comm, &rank);
20 MPI_Comm_split(multiple_client_comm, rank, 0,
21                &single_server_comm);

```

MPI_COMM_SPLIT_TYPE(comm, split_type, key, info, newcomm)

IN	comm	communicator (handle)
IN	split_type	type of processes to be grouped together (integer)
IN	key	control of rank assignment (integer)
INOUT	info	info argument (handle)
OUT	newcomm	new communicator (handle)

C binding

```

34 int MPI_Comm_split_type(MPI_Comm comm, int split_type, int key, MPI_Info info,
35                        MPI_Comm *newcomm)

```

Fortran 2008 binding

```

37 MPI_Comm_split_type(comm, split_type, key, info, newcomm, ierror)
38     TYPE(MPI_Comm), INTENT(IN) :: comm
39     INTEGER, INTENT(IN) :: split_type, key
40     TYPE(MPI_Info), INTENT(IN) :: info
41     TYPE(MPI_Comm), INTENT(OUT) :: newcomm
42     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

44 MPI_COMM_SPLIT_TYPE(COMM, SPLIT_TYPE, KEY, INFO, NEWCOMM, IERROR)
45     INTEGER COMM, SPLIT_TYPE, KEY, INFO, NEWCOMM, IERROR

```

This function partitions the group associated with `comm` into disjoint subgroups such that each subgroup contains all MPI processes in the same grouping referred to by `split_type`.

Within each subgroup, the MPI processes are ranked in the order defined by the value of the argument `key`, with ties broken according to their rank in the old group. A new communicator is created for each subgroup and returned in `newcomm`. This is a collective call. All MPI processes in the group associated with `comm` must provide the same `split_type`, but each MPI process is permitted to provide different values for `key`. An exception to this rule is that an MPI process may supply the type value `MPI_UNDEFINED`, in which case `MPI_COMM_NULL` is returned in `newcomm` for such MPI process. No cached information propagates from `comm` to `newcomm` and no virtual topology information is added to the created communicators.

For `split_type`, the following values are defined by MPI:

MPI_COMM_TYPE_SHARED: all MPI processes in the group of `newcomm` are part of the same *shared memory domain* and can create a *shared memory segment* (e.g., with a successful call to `MPI_WIN_ALLOCATE_SHARED`). This segment can subsequently be used for load/store accesses by all MPI processes in `newcomm`.

Advice to users. Since the location of some of the MPI processes may change during the application execution, the communicators created with the value `MPI_COMM_TYPE_SHARED` before this change may not reflect an actual ability to share memory between MPI processes after this change. (*End of advice to users.*)

MPI_COMM_TYPE_HW_GUIDED: this value specifies that the communicator `comm` is split according to a **hardware resource type** (for example a computing core or an L3 cache) specified by the `"mpi_hw_resource_type"` info key. Each output communicator `newcomm` corresponds to a single instance of the specified hardware resource type. The MPI processes in the group associated with the output communicator `newcomm` utilize that specific hardware resource type instance, and no other instance of the same hardware resource type.

If an MPI process does not meet the above criteria, then `MPI_COMM_NULL` is returned in `newcomm` for such MPI process.

`MPI_COMM_NULL` is also returned in `newcomm` in the following cases:

- `MPI_INFO_NULL` is provided.
- The info handle does not include the key `"mpi_hw_resource_type"`.
- The MPI implementation neither recognizes nor supports the info key `"mpi_hw_resource_type"`.
- The MPI implementation does not recognize the value associated with the info key `"mpi_hw_resource_type"`.

The MPI implementation will return in the group of the output communicator `newcomm` the largest subset of MPI processes that match the splitting criterion.

The MPI processes in the group associated with `newcomm` are ranked in the order defined by the value of the argument `key` with ties broken according to their rank in the group associated with `comm`.

Advice to users. The set of hardware resources that an MPI process is able to utilize may change during the application execution (e.g., because of the relocation of an MPI process), in which case the communicators created with the value

MPI_COMM_TYPE_HW_GUIDED before this change may not reflect the utilization of hardware resources of such MPI process at any time after the communicator creation. (*End of advice to users.*)

The user explicitly constrains with the info argument the splitting of the input communicator `comm`. To this end, the info key `"mpi_hw_resource_type"` is reserved and its associated value is an implementation-defined string designating the type of the requested hardware resource (e.g., `"NUMANode"`, `"Package"` or `"L3Cache"`).

The value `"mpi_shared_memory"` is reserved and its use is equivalent to using `MPI_COMM_TYPE_SHARED` for the `split_type` parameter.

Rationale. The value `"mpi_shared_memory"` is defined in order to ensure consistency between the use of `MPI_COMM_TYPE_SHARED` and the use of `MPI_COMM_TYPE_HW_GUIDED`. (*End of rationale.*)

All MPI processes must provide the same value for the info key `"mpi_hw_resource_type"`.

Example 7.3. Splitting `MPI_COMM_WORLD` into `NUMANode` subcommunicators.

```
MPI_Info info;
MPI_Comm hwcomm;
int rank;

MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Info_create(&info);
MPI_Info_set(info, "mpi_hw_resource_type", "NUMANode");
MPI_Comm_split_type(MPI_COMM_WORLD,
                   MPI_COMM_TYPE_HW_GUIDED,
                   rank, info, &hwcomm);
```

MPI_COMM_TYPE_RESOURCE_GUIDED: this value specifies that the communicator `comm` is split according to a **hardware resource type** (for example a computing core or an L3 cache) specified by the `"mpi_hw_resource_type"` info key or to a **logical resource type** (for example a process set name, see Section 11.3.2) specified by the `"mpi_pset_name"` info key.

Each output communicator `newcomm` corresponds to a single instance of the specified resource type. The MPI processes in the group associated with the output communicator `newcomm` utilize that specific resource type instance, and no other instance of the same resource type.

If an MPI process does not meet the above criteria, then `MPI_COMM_NULL` is returned in `newcomm` for such process.

`MPI_COMM_NULL` is also returned in `newcomm` in the following cases:

- `MPI_INFO_NULL` is provided.
- The info handle includes neither the key `"mpi_hw_resource_type"` nor the key `"mpi_pset_name"`.
- The MPI implementation neither recognizes nor supports the info keys `"mpi_hw_resource_type"` and `"mpi_pset_name"`.
- The MPI implementation does not recognize the value associated with the info key `"mpi_hw_resource_type"` or `"mpi_pset_name"`.

The MPI implementation will return in the group of the output communicator `newcomm` the largest subset of MPI processes that match the splitting criterion.

Advice to users. The set of resources that an MPI process is able to utilize may change during the application execution (e.g., because of the relocation of an MPI process), in which case the communicators created with the value `MPI_COMM_TYPE_RESOURCE_GUIDED` before this change may not reflect the utilization of resources of such process at any time after the communicator creation. (*End of advice to users.*)

The user explicitly constrains with the `info` argument the splitting of the input communicator `comm`. To this end, the following info keys are reserved and their associated values are implementation-defined strings designating the type of the requested resource. Only one of these info keys can be used in `info` at a time in a call to `MPI_COMM_SPLIT_TYPE`; use of more than one info key is erroneous.

"`mpi_hw_resource_type`" is used to specify the type of a requested hardware resource (e.g., "NUMANode", "Package" or "L3Cache"). The value "`mpi_shared_memory`" is reserved and its use is equivalent to using `MPI_COMM_TYPE_SHARED` for the `split_type` parameter.

Rationale. The value "`mpi_shared_memory`" is defined in order to ensure consistency between the use of `MPI_COMM_TYPE_SHARED` and the use of `MPI_COMM_TYPE_RESOURCE_GUIDED`. (*End of rationale.*)

All MPI processes in the group of the input communicator `comm` must provide the same info key to perform the splitting action. All MPI processes in the group of the input communicator `comm` must provide the same value for the info key "`mpi_hw_resource_type`".

"`mpi_pset_name`" is used to specify the type of a requested logical resource through the utilization of a process set name (e.g., "app://ocean" or "app://atmos"). This process set name must be valid in the session from which the input communicator `comm` is derived. If this input communicator is not derived from a session, then `MPI_COMM_NULL` is returned in `newcomm`.

All MPI processes that are both in the group of the input communicator `comm` and in the process set identified by the given process set name must provide the same info key to perform the splitting action. All MPI processes that are both in the group of the input communicator `comm` and in the process set identified by the given process set name must provide the same value for the info key "`mpi_pset_name`".

Example 7.4. Splitting `MPI_COMM_WORLD` into NUMANode subcommunicators.

```
MPI_Info info;
MPI_Comm hwcomm;
int rank;

MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Info_create(&info);
MPI_Info_set(info, "mpi_hw_resource_type", "NUMANode");
MPI_Comm_split(MPI_COMM_WORLD,
```

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

```

MPI_COMM_TYPE_RESOURCE_GUIDED ,
rank, info, &hwcomm);

```

MPI_COMM_TYPE_HW_UNGUIDED: the group of MPI processes associated with `newcomm` must be a *strict* subset of the group associated with `comm` and each `newcomm` corresponds to a single instance of a **hardware resource type** (for example a computing core or an L3 cache).

All MPI processes in the group associated with `comm` that utilize that specific hardware resource type instance—and no other instance of the same hardware resource type—are included in the group of `newcomm`.

If a given MPI process cannot be a member of a communicator that forms such a strict subset, or does not meet the above criteria, then `MPI_COMM_NULL` is returned in `newcomm` for this process.

Advice to implementors. In a high-quality MPI implementation, the number of different new valid communicators `newcomm` produced by this splitting operation should be minimal unless the user provides a key/value pair that modifies this behavior. The sets of hardware resource types used for the splitting operation are implementation-dependent, but should reflect the hardware of the actual system on which the application is currently executing. (*End of advice to implementors.*)

Rationale. If the hardware resources are hierarchically organized, calling this routine several times using as its input communicator `comm` the output communicator `newcomm` of the previous call creates a sequence of `newcomm` communicators in each MPI process, which exposes a hierarchical view of the hardware platform, as shown in Example 7.5. This sequence of returned `newcomm` communicators may differ from the sets of hardware resource types, as shown in the second splitting operation in Figure 7.3. (*End of rationale.*)

Advice to users. Each output communicator `newcomm` can represent a different hardware resource type (see Figure 7.3 for an example). The set of hardware resources an MPI process utilizes may change during the application execution (e.g., because of MPI process relocation), in which case the communicators created with the value `MPI_COMM_TYPE_HW_UNGUIDED` before this change may not reflect the utilization of hardware resources for such MPI process at any time after the communicator creation. (*End of advice to users.*)

If a valid `info` handle is provided as an argument, the MPI implementation sets the `info` key `"mpi_hw_resource_type"` for each MPI process in the group associated with a returned `newcomm` communicator and the `info` key value is an implementation-defined string that indicates the hardware resource type represented by `newcomm`. The same hardware resource type must be set in all MPI processes in the group associated with `newcomm`.

Example 7.5. Recursive splitting of `MPI_COMM_WORLD`.

```

#define MAX_NUM_LEVELS 32

MPI_Comm hwcomm[MAX_NUM_LEVELS];

```

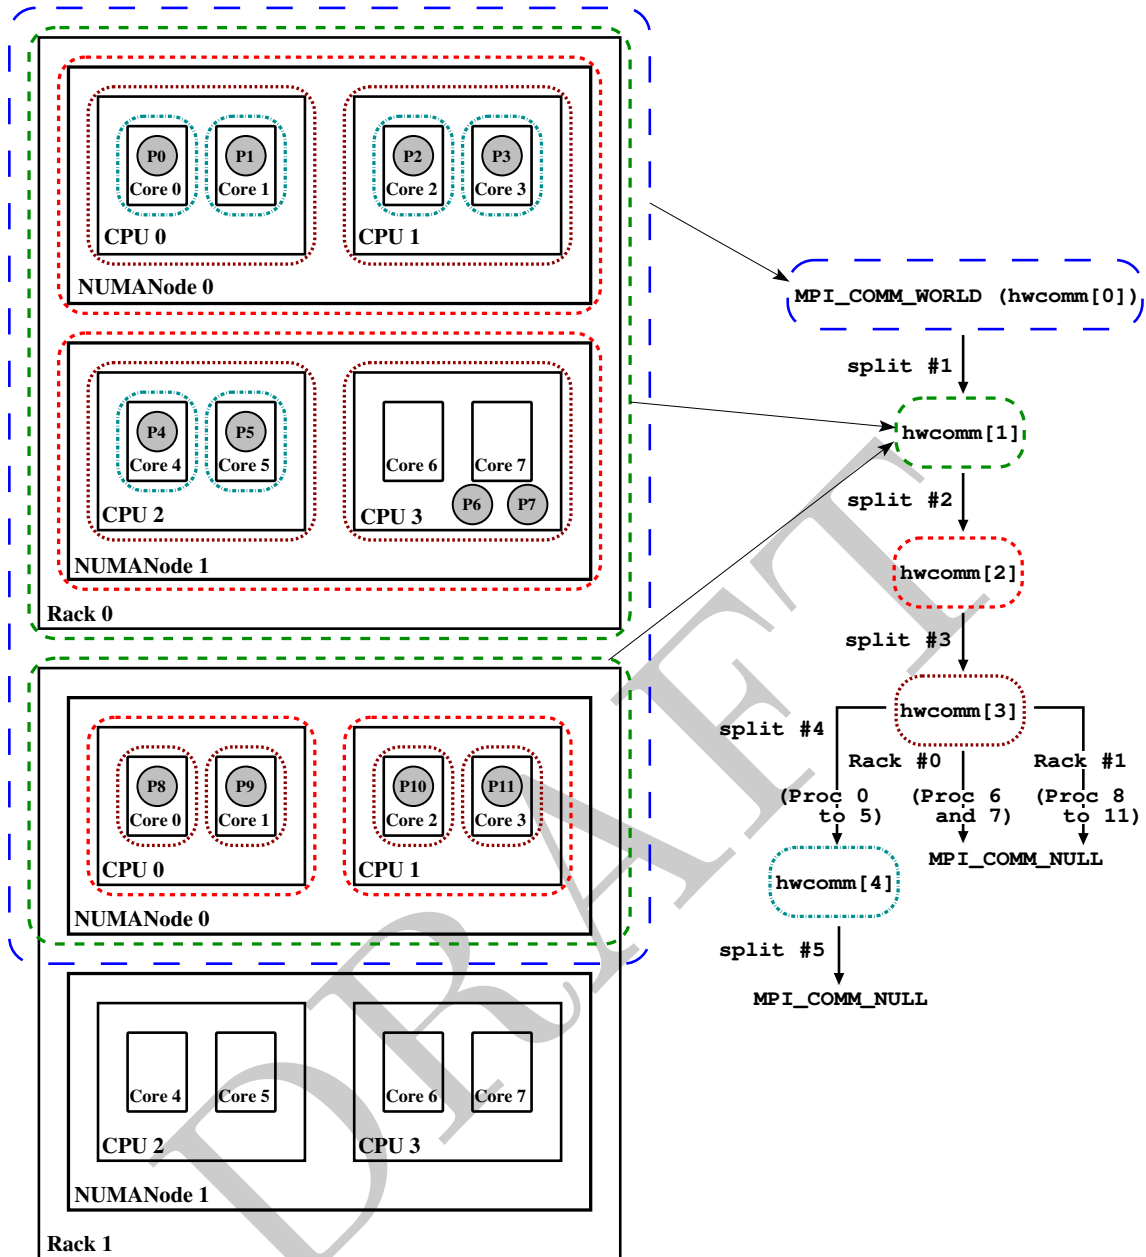


Figure 7.3: Recursive splitting of `MPI_COMM_WORLD` with `MPI_COMM_SPLIT_TYPE` and `MPI_COMM_TYPE_HW_UNGUIDED`. Dashed lines represent communicators whilst solid lines represent hardware resources. MPI processes (P0 to P11) utilize exclusively their respective core, except for P6 and P7, which utilize CPU #3 of Rack #0 and can therefore use Cores #6 and #7 indifferently. The second splitting operation yields two subcommunicators corresponding to NUMANodes in Rack #0 and to CPUs in Rack #1 because Rack #1 features only one NUMANode, which corresponds to the whole portion of the Rack that is included in `MPI_COMM_WORLD` and `hwcomm[1]`. For the first splitting operation, the hardware resource type returned in the info argument is “Rack” on the MPI processes on Rack #0, whereas on Rack #1, it can be either “Rack” or “NUMANode”.

```

1  int      rank, level_num = 0;
2
3  hwcomm[level_num] = MPI_COMM_WORLD;
4
5  while((hwcomm[level_num] != MPI_COMM_NULL)
6        && (level_num < MAX_NUM_LEVELS-1)){
7      MPI_Comm_rank(hwcomm[level_num], &rank);
8      MPI_Comm_split_type(hwcomm[level_num],
9                          MPI_COMM_TYPE_HW_UNGUIDED,
10                         rank,
11                         MPI_INFO_NULL,
12                         &hwcomm[level_num+1]);
13
14     level_num++;
15 }

```

Advice to implementors. Implementations can define their own `split_type` values, or use the `info` argument, to assist in creating communicators that help expose platform-specific information to the application. The concept of hardware-based communicators was first described by Träff [68] for SMP systems. Guided and unguided modes description as well as an implementation path are introduced by Goglin et al. [28]. (*End of advice to implementors.*)

`MPI_COMM_CREATE_FROM_GROUP(group, stringtag, info, errhandler, newcomm)`

IN	group	group (handle)
IN	stringtag	unique identifier for this operation (string)
IN	info	info object (handle)
IN	errhandler	error handler to be attached to new intra-communicator (handle)
OUT	newcomm	new communicator (handle)

C binding

```

int MPI_Comm_create_from_group(MPI_Group group, const char *stringtag,
                              MPI_Info info, MPI_Errhandler errhandler, MPI_Comm *newcomm)

```

Fortran 2008 binding

```

MPI_Comm_create_from_group(group, stringtag, info, errhandler, newcomm, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group
  CHARACTER(LEN=*), INTENT(IN) :: stringtag
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
  TYPE(MPI_Comm), INTENT(OUT) :: newcomm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_COMM_CREATE_FROM_GROUP(GROUP, STRINGTAG, INFO, ERRHANDLER, NEWCOMM, IERROR)
  INTEGER GROUP, INFO, ERRHANDLER, NEWCOMM, IERROR

```

CHARACTER*(*) STRINGTAG

MPI_COMM_CREATE_FROM_GROUP is similar to MPI_COMM_CREATE_GROUP, except that the set of MPI processes involved in the creation of the new intra-communicator is specified by a `group` argument, rather than the group associated with a pre-existing communicator. If a nonempty `group` is specified, then all MPI processes in that group must call the function and each of these MPI processes must provide the same arguments, including a `group` that contains the same members with the same ordering, and identical `stringtag` value. In the event that MPI_GROUP_EMPTY is supplied as the `group` argument, then the call is a local operation and MPI_COMM_NULL is returned as `newcomm`. The `stringtag` argument is analogous to the `tag` used for MPI_COMM_CREATE_GROUP. If multiple threads at a given MPI process perform concurrent MPI_COMM_CREATE_FROM_GROUP operations, the user must distinguish these operations by providing different `stringtag` arguments. The `stringtag` shall not exceed MPI_MAX_STRINGTAG_LEN characters in length. For C, this includes space for a null terminating character. MPI_MAX_STRINGTAG_LEN shall have a value of at least 63.

The `errhandler` argument specifies an error handler to be attached to the new intra-communicator. Section 9.3 specifies the error handler to be invoked if an error is encountered during the invocation of MPI_COMM_CREATE_FROM_GROUP.

The `info` argument provides hints and assertions, possibly MPI implementation dependent, which indicate desired characteristics and guide communicator creation.

Advice to users. The `stringtag` argument is used to distinguish concurrent communicator construction operations issued by different entities. As such, it is important to ensure that this argument is unique for each concurrent call to MPI_COMM_CREATE_FROM_GROUP. Reverse domain name notation convention [2] is one approach to constructing unique `stringtag` arguments. See also example 11.10. (*End of advice to users.*)

7.4.3 Communicator Destructors

MPI_COMM_FREE(comm)

INOUT comm communicator to be destroyed (handle)

C binding

int MPI_Comm_free(MPI_Comm *comm)

Fortran 2008 binding

MPI_Comm_free(comm, ierror)

TYPE(MPI_Comm), INTENT(INOUT) :: comm

INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

MPI_COMM_FREE(COMM, IERROR)

INTEGER COMM, IERROR

This collective operation marks the communication object for deallocation. Any operations that use the communicator `comm` (whether active or inactive at the time of this

1 procedure call) will continue to work; the object is actually deallocated only if there are
2 no other active references to it. The handle is set to `MPI_COMM_NULL` in the calling MPI
3 process. This call applies to intra- and inter-communicators. The delete callback functions
4 for all cached attributes (see Section 7.7) are called in arbitrary order.

5
6 *Advice to implementors.* Though collective, it is anticipated that this operation will
7 normally be implemented to be local, though a debugging version of an MPI library
8 might choose to synchronize. (*End of advice to implementors.*)
9

10 7.4.4 Communicator Info

11 Hints specified via info (see Chapter 10) allow a user to provide information to direct opti-
12 mization. Providing hints may enable an implementation to deliver increased performance
13 or minimize use of system resources. As described in Section 10, an implementation is free
14 to ignore all hints; however, applications must comply with any info hints they provide that
15 are used by the MPI implementation (i.e., are returned by a call to `MPI_COMM_GET_INFO`)
16 and that place a restriction on the behavior of the application. Hints are specified on a per
17 communicator basis, in `MPI_COMM_DUP_WITH_INFO`, `MPI_COMM_IDUP_WITH_INFO`,
18 `MPI_COMM_SET_INFO`, `MPI_COMM_SPLIT_TYPE`, `MPI_DIST_GRAPH_CREATE`, and
19 `MPI_DIST_GRAPH_CREATE_ADJACENT`, via the opaque info object. When an info object
20 that specifies a subset of valid hints is passed to `MPI_COMM_SET_INFO`, there will be no
21 effect on previously set or defaulted hints that the info does not specify.
22

23 *Advice to implementors.* It may happen that a program is coded with hints for one
24 system, and later executes on another system that does not support these hints. In
25 general, unsupported hints should simply be ignored. Needless to say, no hint can be
26 mandatory. However, for each hint used by a specific implementation, a default value
27 must be provided when the user does not specify a value for this hint. (*End of advice*
28 *to implementors.*)
29

30 *Advice to users.* Some optimizations may only be possible when all processes in the
31 group of the communicator provide a given info key with the same value. (*End of*
32 *advice to users.*)
33

34 Info hints are not propagated by MPI from one communicator to another. The following
35 info keys are valid for all communicators.
36

37 *Advice to users.* Some optimizations may only be possible when all MPI processes in
38 the group of the communicator provide a given info key with the same value. (*End of*
39 *advice to users.*)
40

41 **"mpi_assert_no_any_tag" (boolean, default: "false"):** If set to "true", then the implemen-
42 tation may assume that the MPI process will not use the `MPI_ANY_TAG` wildcard on
43 the given communicator.
44

45 **"mpi_assert_no_any_source" (boolean, default: "false"):** If set to "true", then the imple-
46 mentation may assume that the MPI process will not use the `MPI_ANY_SOURCE` wild-
47 card on the given communicator.
48

"mpi_assert_exact_length" (boolean, default: "false"): If set to "true", then the implementation may assume that the lengths of messages received by the MPI process are equal to the lengths of the corresponding receive buffers, for point-to-point communication operations on the given communicator.

"mpi_assert_allow_overtaking" (boolean, default: "false"): If set to "true", then the implementation may assume that point-to-point communications on the given communicator do not rely on the nonovertaking rule specified in Section 3.5. In other words, the application asserts that send operations are not required to be matched at the receiver in the order in which the send operations were posted by the sender, and receive operations are not required to be matched in the order in which they were posted by the receiver.

Advice to users. Use of the "mpi_assert_allow_overtaking" info key can result in nondeterminism in the message matching order. (*End of advice to users.*)

"mpi_assert_strict_persistent_collective_ordering" (boolean, default: "false"): If set to "true", then the implementation may assume that all the persistent collective operations are started in the same order across all MPI processes in the group of the communicator. It is required that if this assertion is made on one member of the communicator's group, then it must be made on all members of that communicator's group with the same value.

Advice to users. Use of the "mpi_assert_strict_persistent_collective_ordering" may be needed because some optimizations may only be possible on certain systems when strict collective ordering is asserted for the underlying communicator of a persistent collective operation. (*End of advice to users.*)

"mpi_assert_memory_alloc_kinds" (string, not set by default): If set, the implementation may assume that the memory for all communication buffers passed to MPI operations performed by the calling MPI process on the given communicator will use only the memory allocation kinds listed in the value string. See Section 11.4.3.

MPI_COMM_SET_INFO(comm, info)

INOUT	comm	communicator (handle)
IN	info	info object (handle)

C binding

```
int MPI_Comm_set_info(MPI_Comm comm, MPI_Info info)
```

Fortran 2008 binding

```
MPI_Comm_set_info(comm, info, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Info), INTENT(IN) :: info
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_SET_INFO(COMM, INFO, IERROR)
```

1 INTEGER COMM, INFO, IERROR

2 MPI_COMM_SET_INFO updates the hints of the communicator associated with `comm`
 3 using the hints provided in `info`. This operation has no effect on previously set or defaulted
 4 hints that are not specified by `info`. It also has no effect on previously set or defaulted
 5 hints that are specified by `info`, but are ignored by the MPI implementation in this call to
 6 MPI_COMM_SET_INFO. MPI_COMM_SET_INFO is a collective routine. The `info` object
 7 may be different on each MPI process, but any `info` entries that an implementation requires
 8 to be the same on all MPI processes must appear with the same value in each MPI process's
 9 `info` object.
 10

11 *Advice to users.* Some `info` items that an implementation can use when it creates
 12 a communicator cannot easily be changed once the communicator has been created.
 13 Thus, an implementation may ignore hints issued in this call that it would have
 14 accepted in a creation call. An implementation may also be unable to update certain
 15 `info` hints in a call to MPI_COMM_SET_INFO. MPI_COMM_GET_INFO can be used to
 16 determine whether updates to existing `info` hints were ignored by the implementation.
 17 (*End of advice to users.*)
 18

19 *Advice to users.* Setting `info` hints on the predefined communicators
 20 MPI_COMM_WORLD and MPI_COMM_SELF may have unintended effects, as changes to
 21 these global objects may affect all components of the application, including libraries
 22 and tools. Users must ensure that all components of the application that use a given
 23 communicator, including libraries and tools, can comply with any `info` hints associated
 24 with that communicator. (*End of advice to users.*)
 25

26
 27
 28 MPI_COMM_GET_INFO(`comm`, `info_used`)

29 IN `comm` communicator object (handle)
 30 OUT `info_used` new `info` object (handle)
 31

32 33 **C binding**

34 `int MPI_Comm_get_info(MPI_Comm comm, MPI_Info *info_used)`

35 **Fortran 2008 binding**

36 `MPI_Comm_get_info(comm, info_used, ierror)`
 37 TYPE(MPI_Comm), INTENT(IN) :: `comm`
 38 TYPE(MPI_Info), INTENT(OUT) :: `info_used`
 39 INTEGER, OPTIONAL, INTENT(OUT) :: `ierror`
 40

41 **Fortran binding**

42 `MPI_COMM_GET_INFO(COMM, INFO_USED, IERROR)`
 43 INTEGER COMM, INFO_USED, IERROR

44 MPI_COMM_GET_INFO returns a new `info` object containing the hints of the commu-
 45 nicator associated with `comm`. The current setting of all hints related to this communicator
 46 is returned in `info_used`. An MPI implementation is required to return all hints that are
 47 supported by the implementation and have default values specified; any user-supplied hints
 48 that were not ignored by the implementation; and any additional hints that were set by

the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no key/value pair. The user is responsible for freeing `info_used` via `MPI_INFO_FREE`.

7.5 Motivating Examples

7.5.1 Current Practice #1

Example 7.6. Parallel output of a message

```
int main(int argc, char *argv[])
{
    int me, size;
    ...
    MPI_Init(&argc, &argv);
    MPI_Comm_rank(MPI_COMM_WORLD, &me);
    MPI_Comm_size(MPI_COMM_WORLD, &size);

    (void)printf("MPI process %d size %d\n", me, size);
    ...
    MPI_Finalize();
    return 0;
}
```

Example 7.6 is a do-nothing program that initializes itself, and refers to the “all” communicator, and prints a message. It terminates itself too. This example does not imply that MPI supports `printf`-like communication itself.

Example 7.7. Message exchange (supposing that size is even)

```
int main(int argc, char *argv[])
{
    int me, size;
    int SOME_TAG = 0;
    ...
    MPI_Init(&argc, &argv);

    MPI_Comm_rank(MPI_COMM_WORLD, &me); /* local */
    MPI_Comm_size(MPI_COMM_WORLD, &size); /* local */

    if((me % 2) == 0)
    {
        /* send unless highest-numbered MPI process */
        if((me + 1) < size)
            MPI_Send(..., me + 1, SOME_TAG, MPI_COMM_WORLD);
    }
    else
        MPI_Recv(..., me - 1, SOME_TAG, MPI_COMM_WORLD, &status);

    ...
    MPI_Finalize();
    return 0;
}
```

1
}

2
3 Example 7.7 schematically illustrates message exchanges between “even” and “odd” MPI
4 processes in the “all” communicator.

5
6 7.5.2 Current Practice #27
8 **Example 7.8.**

```
9
10 int main(int argc, char *argv[])
11 {
12     int me, count;
13     void *data;
14     ...
15     MPI_Init(&argc, &argv);
16     MPI_Comm_rank(MPI_COMM_WORLD, &me);
17
18     if(me == 0)
19     {
20         /* get input, create buffer “data” */
21         ...
22     }
23     MPI_Bcast(data, count, MPI_BYTE, 0, MPI_COMM_WORLD);
24
25     ...
26     MPI_Finalize();
27     return 0;
28 }
```

29 Example 7.8 illustrates the use of a collective communication.

30
31 7.5.3 (Approximate) Current Practice #332
33 **Example 7.9.**

```
34
35 int main(int argc, char *argv[])
36 {
37     int me, count, count2;
38     void *send_buf, *recv_buf, *send_buf2, *recv_buf2;
39     MPI_Group group_world, grpem;
40     MPI_Comm commWorker;
41     static int ranks[] = {0};
42     ...
43     MPI_Init(&argc, &argv);
44     MPI_Comm_group(MPI_COMM_WORLD, &group_world);
45     MPI_Comm_rank(MPI_COMM_WORLD, &me); /* local */
46
47     MPI_Group_excl(group_world, 1, ranks, &grpem); /* local */
48     MPI_Comm_create(MPI_COMM_WORLD, grpem, &commWorker);
49
50     if(me != 0)
51     {
```

```

1  /* compute on worker */
2  ...
3  MPI_Reduce(send_buf,recv_buf,count, MPI_INT, MPI_SUM, 1, commWorker);
4  ...
5  MPI_Comm_free(&commWorker);
6  }
7  /* zero falls through immediately to this reduce, others do later... */
8  MPI_Reduce(send_buf2, recv_buf2, count2,
9             MPI_INT, MPI_SUM, 0, MPI_COMM_WORLD);
10
11 MPI_Group_free(&group_world);
12 MPI_Group_free(&grpem);
13 MPI_Finalize();
14 return 0;
15 }

```

Example 7.9 illustrates how a group consisting of all but the zeroth MPI process of the “all” group is created, and then how a communicator is formed (`commWorker`) for that new group. The new communicator is used in a collective call, and all MPI processes execute a collective call in the `MPI_COMM_WORLD` context. This example illustrates how the two communicators (that inherently possess distinct contexts) protect communication. That is, communication in `MPI_COMM_WORLD` is insulated from communication in `commWorker`, and vice versa.

In summary, “group safety” is achieved via communicators because distinct contexts within communicators are enforced to be unique on any MPI process.

7.5.4 Communication Safety Example

The following example (7.10) is meant to illustrate “safety” between point-to-point and collective communication. MPI guarantees that a single communicator can do safe point-to-point and collective communication.

Example 7.10.

```

31 #define TAG_ARBITRARY 12345
32 #define SOME_COUNT    50
33
34 int main(int argc, char *argv[])
35 {
36     int me;
37     MPI_Request request[2];
38     MPI_Status status[2];
39     MPI_Group group_world, subgroup;
40     int ranks[] = {2, 4, 6, 8};
41     MPI_Comm the_comm;
42     ...
43     MPI_Init(&argc, &argv);
44     MPI_Comm_group(MPI_COMM_WORLD, &group_world);
45
46     MPI_Group_incl(group_world, 4, ranks, &subgroup); /* local */
47     MPI_Group_rank(subgroup, &me); /* local */
48
49     MPI_Comm_create(MPI_COMM_WORLD, subgroup, &the_comm);

```

```

1
2  if(me != MPI_UNDEFINED)
3  {
4      MPI_Irecv(buff1, count, MPI_DOUBLE, MPI_ANY_SOURCE,
5                TAG_ARBITRARY, the_comm, request);
6      MPI_Isend(buff2, count, MPI_DOUBLE, (me+1)%4, TAG_ARBITRARY,
7                the_comm, request+1);
8      for(i = 0; i < SOME_COUNT; i++)
9          MPI_Reduce(..., the_comm);
10     MPI_Waitall(2, request, status);
11
12     MPI_Comm_free(&the_comm);
13 }
14
15 MPI_Group_free(&group_world);
16 MPI_Group_free(&subgroup);
17 MPI_Finalize();
18 return 0;
19 }

```

7.5.5 Library Example #1

Example 7.11. First library example
The main program:

```

22 int main(int argc, char *argv[])
23 {
24     int done = 0;
25     user_lib_t *libh_a, *libh_b;
26     void *dataset1, *dataset2;
27     ...
28     MPI_Init(&argc, &argv);
29     ...
30     init_user_lib(MPI_COMM_WORLD, &libh_a);
31     init_user_lib(MPI_COMM_WORLD, &libh_b);
32     ...
33     user_start_op(libh_a, dataset1);
34     user_start_op(libh_b, dataset2);
35     ...
36     while(!done)
37     {
38         /* work */
39         ...
40         MPI_Reduce(..., MPI_COMM_WORLD);
41         ...
42         /* see if done */
43         ...
44     }
45     user_end_op(libh_a);
46     user_end_op(libh_b);
47
48

```

```

uninit_user_lib(libh_a);
uninit_user_lib(libh_b);
MPI_Finalize();
return 0;
}

```

The user library initialization code:

```

void init_user_lib(MPI_Comm comm, user_lib_t **handle)
{
    user_lib_t *save;

    user_lib_initsave(&save); /* local */
    MPI_Comm_dup(comm, &(save->comm));

    /* other inits */
    ...

    *handle = save;
}

```

User start-up code:

```

void user_start_op(user_lib_t *handle, void *data)
{
    MPI_Irecv( ..., handle->comm, &(handle->irecv_handle) );
    MPI_Isend( ..., handle->comm, &(handle->isend_handle) );
}

```

User communication clean-up code:

```

void user_end_op(user_lib_t *handle)
{
    MPI_Status status;
    MPI_Wait(&handle->isend_handle, &status);
    MPI_Wait(&handle->irecv_handle, &status);
}

```

User object clean-up code:

```

void uninit_user_lib(user_lib_t *handle)
{
    MPI_Comm_free(&(handle->comm));
    free(handle);
}

```

7.5.6 Library Example #2

Example 7.12. Second library example

The main program:

```

int main(int argc, char *argv[])
{
    int ma, mb;
    MPI_Group group_world, group_a, group_b;
}

```

```

1  MPI_Comm comm_a, comm_b;
2
3  static int list_a[] = {0, 1};
4  #if defined(EXAMPLE_2B) || defined(EXAMPLE_2C)
5  static int list_b[] = {0, 2 ,3};
6  #else /* EXAMPLE_2A */
7  static int list_b[] = {0, 2};
8  #endif
9  int size_list_a = sizeof(list_a)/sizeof(int);
10 int size_list_b = sizeof(list_b)/sizeof(int);
11
12 ...
13 MPI_Init(&argc, &argv);
14 MPI_Comm_group(MPI_COMM_WORLD, &group_world);
15
16 MPI_Group_incl(group_world, size_list_a, list_a, &group_a);
17 MPI_Group_incl(group_world, size_list_b, list_b, &group_b);
18
19 MPI_Comm_create(MPI_COMM_WORLD, group_a, &comm_a);
20 MPI_Comm_create(MPI_COMM_WORLD, group_b, &comm_b);
21
22 if(comm_a != MPI_COMM_NULL)
23     MPI_Comm_rank(comm_a, &ma);
24 if(comm_b != MPI_COMM_NULL)
25     MPI_Comm_rank(comm_b, &mb);
26
27 if(comm_a != MPI_COMM_NULL)
28     lib_call(comm_a);
29
30 if(comm_b != MPI_COMM_NULL)
31 {
32     lib_call(comm_b);
33     lib_call(comm_b);
34 }
35
36 if(comm_a != MPI_COMM_NULL)
37     MPI_Comm_free(&comm_a);
38 if(comm_b != MPI_COMM_NULL)
39     MPI_Comm_free(&comm_b);
40 MPI_Group_free(&group_a);
41 MPI_Group_free(&group_b);
42 MPI_Group_free(&group_world);
43 MPI_Finalize();
44 return 0;
45 }

```

The library:

```

43 void lib_call(MPI_Comm comm)
44 {
45     int me, done = 0;
46     MPI_Status status;
47     MPI_Comm_rank(comm, &me);
48     if(me == 0)

```

```

1   while(!done)
2   {
3       MPI_Recv(..., MPI_ANY_SOURCE, MPI_ANY_TAG, comm, &status);
4       ...
5   }
6   else
7   {
8       /* work */
9       MPI_Send(..., 0, ARBITRARY_TAG, comm);
10      ...
11  }
12  #ifdef EXAMPLE_2C
13      /* include (resp, exclude) for safety (resp, no safety): */
14      MPI_Barrier(comm);
15  #endif
16  }

```

The above example is three examples, depending on whether or not one includes rank 3 in `list_b`, and whether or not a synchronizing operation is included in `lib_call`. This example illustrates that, despite contexts, subsequent calls to `lib_call` with the same context need not be safe from one another (colloquially, “back-masking”). Safety is realized if a call to `MPI_Barrier` is added. What this demonstrates is that libraries have to be written carefully, even with contexts. When rank 3 is excluded, then the synchronizing operation is not needed to get safety from back-masking.

Algorithms like “reduce” and “allreduce” have strong enough source selectivity properties so that they are inherently okay (no back-masking), provided that MPI provides basic guarantees. So are multiple calls to a typical tree-broadcast algorithm with the same root or different roots (see [65]). Here we rely on two guarantees of MPI: pairwise ordering of messages between MPI processes in the same context, and source selectivity—deleting either feature removes the guarantee that back-masking cannot be required.

Algorithms that try to do nondeterministic broadcasts or other calls that include wildcard operations will not generally have the good properties of the deterministic implementations of “reduce,” “allreduce,” and “broadcast.” Such algorithms would have to utilize the monotonically increasing tags (within a communicator scope) to keep things straight.

All of the foregoing is a supposition of “collective calls” implemented with point-to-point operations. MPI implementations may or may not implement collective calls using point-to-point operations. These algorithms are used to illustrate the issues of correctness and safety, independent of how MPI implements its collective calls. See also Section 7.9.

7.6 Inter-Communication

This section introduces the concept of inter-communication and describes the portions of MPI that support it. It describes support for writing programs that contain user-level servers.

All communication described thus far has involved communication between MPI processes that are members of the same group. This type of communication is called “**intra-communication**” and the communicator used is called an “**intra-communicator**,” as we have noted earlier in the chapter.

In modular and multi-disciplinary applications, different MPI process groups execute

1 distinct modules and MPI processes within different modules communicate with one another
2 in a pipeline or a more general module graph. In these applications, the most natural way for
3 a MPI process to specify a target MPI process is by the rank of the target MPI process within
4 the target group. In applications that contain internal user-level servers, each server may be
5 a MPI process group that provides services to one or more clients, and each client may be a
6 MPI process group that uses the services of one or more servers. It is again most natural to
7 specify the target MPI process by rank within the target group in these applications. This
8 type of communication is called “**inter-communication**” and the communicator used is
9 called an “**inter-communicator**,” as introduced earlier.

10 An **inter-communication** is a point-to-point communication between MPI processes
11 in different groups. The group containing an MPI process that initiates an inter-communi-
12 cation operation is called the “local group,” that is, the sender in a send and the receiver in
13 a receive. The group containing the target MPI process is called the “remote group,” that
14 is, the receiver in a send and the sender in a receive. As in intra-communication, the target
15 MPI process is specified using a (communicator, rank) pair. Unlike intra-communication, the
16 rank is relative to a second, remote group.

17 All inter-communicator constructors are blocking except for `MPI_COMM_IDUP` and
18 require that the local and remote groups be disjoint.

19 *Advice to users.* The groups must be disjoint for several reasons. First, the intent
20 of the inter-communicators is to provide a communicator for communication between
21 disjoint groups. This is reflected in the definition of `MPI_INTERCOMM_MERGE`,
22 which allows the user to control the ranking of the MPI processes in the created
23 intra-communicator; this ranking makes little sense if the groups are not disjoint. In
24 addition, the natural extension of collective operations to inter-communicators makes
25 the most sense when the groups are disjoint. (*End of advice to users.*)

26 Here is a summary of the properties of inter-communication and inter-communicators:

- 27 • The syntax of point-to-point and collective communication is the same for both inter-
28 and intra-communication. The same communicator can be used both for send and for
29 receive operations.
- 30 • A target MPI process is addressed by its rank in the remote group, both for sends and
31 for receives.
- 32 • Communications using an inter-communicator are guaranteed not to conflict with any
33 communications that use a different communicator.
- 34 • A communicator will provide either intra- or inter-communication, never both.

35 The routine `MPI_COMM_TEST_INTER` may be used to determine if a communicator is an
36 inter- or intra-communicator. Inter-communicators can be used as arguments to some of the
37 other communicator access routines. Inter-communicators cannot be used as input to some
38 of the constructor routines for intra-communicators (for instance, `MPI_CART_CREATE`).

39 *Advice to implementors.* For the purpose of point-to-point communication, commu-
40 nicators can be represented in each process by a tuple consisting of:

41 **group**
42 **send_context**

receive_context**source**

For inter-communicators, *group* describes the remote group, and *source* is the rank of the MPI process in the local group. For intra-communicators, *group* is the communicator group (remote=local), *source* is the rank of the MPI process in this group, and *send context* and *receive context* are identical. A group can be represented by a rank-to-absolute-address translation table.

The inter-communicator cannot be discussed sensibly without considering MPI processes in both the local and remote groups. Imagine an MPI process **P** in group \mathcal{P} , which has an inter-communicator $C_{\mathcal{P}}$, and an MPI process **Q** in group \mathcal{Q} , which has an inter-communicator $C_{\mathcal{Q}}$. Then

- $C_{\mathcal{P}}.\mathbf{group}$ describes the group \mathcal{Q} and $C_{\mathcal{Q}}.\mathbf{group}$ describes the group \mathcal{P} .
- $C_{\mathcal{P}}.\mathbf{send_context} = C_{\mathcal{Q}}.\mathbf{receive_context}$ and the context is unique in \mathcal{Q} ;
 $C_{\mathcal{P}}.\mathbf{receive_context} = C_{\mathcal{Q}}.\mathbf{send_context}$ and this context is unique in \mathcal{P} .
- $C_{\mathcal{P}}.\mathbf{source}$ is rank of **P** in \mathcal{P} and $C_{\mathcal{Q}}.\mathbf{source}$ is rank of **Q** in \mathcal{Q} .

Assume that **P** sends a message to **Q** using the inter-communicator. Then **P** uses the **group** table to find the absolute address of **Q**; **source** and **send_context** are appended to the message.

Assume that **Q** posts a receive with an explicit source argument using the inter-communicator. Then **Q** matches **receive_context** to the message context and source argument to the message source.

The same algorithm is appropriate for intra-communicators as well.

In order to support inter-communicator accessors and constructors, it is necessary to supplement this model with additional structures, that store information about the local communication group, and additional safe contexts. (*End of advice to implementors.*)

7.6.1 Inter-Communicator Accessors

MPI_COMM_TEST_INTER(comm, flag)

IN	comm	communicator (handle)
OUT	flag	true if comm is an inter-communicator (logical)

C binding

```
int MPI_Comm_test_inter(MPI_Comm comm, int *flag)
```

Fortran 2008 binding

```
MPI_Comm_test_inter(comm, flag, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_TEST_INTER(COMM, FLAG, IERROR)
```

Table 7.1: MPI_COMM_* function behavior (in inter-communication mode)

MPI_COMM_SIZE	returns the size of the local group.
MPI_COMM_GROUP	returns the local group.
MPI_COMM_RANK	returns the rank in the local group

INTEGER COMM, IERROR
LOGICAL FLAG

This local routine allows the calling MPI process to determine if a communicator is an inter-communicator or an intra-communicator. It returns true if it is an inter-communicator, otherwise false.

Table 7.1 describes the behavior when an inter-communicator is used as an input argument to the communicator accessors described above under intra-communication. Furthermore, the operation MPI_COMM_COMPARE is valid for inter-communicators. Both communicators must be either intra- or inter-communicators, or else MPI_UNEQUAL results. Both corresponding local and remote groups must compare correctly to get the results MPI_CONGRUENT or MPI_SIMILAR. In particular, it is possible for MPI_SIMILAR to result because either the local or remote groups were similar but not identical.

The following accessors provide consistent access to the remote group of an inter-communicator. The following are all local operations.

MPI_COMM_REMOTE_SIZE(comm, size)

IN	comm	inter-communicator (handle)
OUT	size	number of MPI processes in the remote group of comm (integer)

C binding

```
int MPI_Comm_remote_size(MPI_Comm comm, int *size)
```

Fortran 2008 binding

```
MPI_Comm_remote_size(comm, size, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(OUT) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_REMOTE_SIZE(COMM, SIZE, IERROR)
  INTEGER COMM, SIZE, IERROR
```

MPI_COMM_REMOTE_GROUP(comm, group)

IN	comm	inter-communicator (handle)
OUT	group	remote group corresponding to comm (handle)

C binding

```
int MPI_Comm_remote_group(MPI_Comm comm, MPI_Group *group)
```

Fortran 2008 binding

```
MPI_Comm_remote_group(comm, group, ierror)
    TYPE(MPI_Comm), INTENT(IN) :: comm
    TYPE(MPI_Group), INTENT(OUT) :: group
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_REMOTE_GROUP(COMM, GROUP, IERROR)
    INTEGER COMM, GROUP, IERROR
```

Rationale. Symmetric access to both the local and remote groups of an inter-communicator is important, so this function, as well as `MPI_COMM_REMOTE_SIZE` have been provided. (*End of rationale.*)

7.6.2 Inter-Communicator Operations

This section introduces five blocking inter-communicator operations.

`MPI_INTERCOMM_CREATE` is used to bind two intra-communicators into an inter-communicator; the function `MPI_INTERCOMM_CREATE_FROM_GROUPS` constructs an inter-communicator from two previously defined disjoint groups; the function `MPI_INTERCOMM_MERGE` creates an intra-communicator by merging the local and remote groups of an inter-communicator. The functions `MPI_COMM_DUP` and `MPI_COMM_FREE`, introduced previously, duplicate and free an inter-communicator, respectively.

Overlap of local and remote groups that are bound into an inter-communicator is prohibited. If there is overlap, then the program is erroneous and is likely to deadlock.

The function `MPI_INTERCOMM_CREATE` can be used to create an inter-communicator from two existing intra-communicators, in the following situation: At least one selected member from each group (the “group leader”) has the ability to communicate with the selected member from the other group; that is, a “peer” communicator exists to which both leaders belong, and each leader knows the rank of the other leader in this peer communicator. Furthermore, members of each group know the rank of their leader.

Construction of an inter-communicator from two intra-communicators requires separate collective operations in the local group and in the remote group, as well as a point-to-point communication between an MPI process in the local group and an MPI process in the remote group.

When using the World Model (Section 11.2), the `MPI_COMM_WORLD` communicator (or preferably a dedicated duplicate thereof) can be this peer communicator. For applications that use the Sessions Model, or the spawn or join operations, it may be necessary to first create an intra-communicator to be used as the peer communicator.

The application topology functions described in Chapter 8 do not apply to inter-communicators. Users that require this capability should utilize `MPI_INTERCOMM_MERGE` to build an intra-communicator, then apply the graph or cartesian topology capabilities to that intra-communicator, creating an appropriate topology-oriented intra-communicator. Alternatively, it may be reasonable to devise one’s own application topology mechanisms for this case, without loss of generality.

```

1 MPI_INTERCOMM_CREATE(local_comm, local_leader, peer_comm, remote_leader, tag,
2     newintercomm)
3
4     IN     local_comm     local intra-communicator (handle)
5     IN     local_leader   rank of local group leader in local_comm (integer)
6     IN     peer_comm     “peer” communicator; significant only at the
7     local_leader (handle)
8     IN     remote_leader  rank of remote group leader in peer_comm;
9     significant only at the local_leader (integer)
10
11    IN     tag             tag (integer)
12
13    OUT    newintercomm   new inter-communicator (handle)

```

C binding

```

14
15 int MPI_Intercomm_create(MPI_Comm local_comm, int local_leader,
16     MPI_Comm peer_comm, int remote_leader, int tag,
17     MPI_Comm *newintercomm)
18

```

Fortran 2008 binding

```

19 MPI_Intercomm_create(local_comm, local_leader, peer_comm, remote_leader, tag,
20     newintercomm, ierror)
21
22     TYPE(MPI_Comm), INTENT(IN) :: local_comm, peer_comm
23     INTEGER, INTENT(IN) :: local_leader, remote_leader, tag
24     TYPE(MPI_Comm), INTENT(OUT) :: newintercomm
25     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

26
27 MPI_INTERCOMM_CREATE(LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG,
28     NEWINTERCOMM, IERROR)
29
30     INTEGER LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG,
31     NEWINTERCOMM, IERROR

```

32 This call creates an inter-communicator. It is collective over the union of the local and
33 remote groups. MPI processes should provide identical `local_comm` and `local_leader` argu-
34 ments within each group. Wildcards are not permitted for `remote_leader`, `local_leader`, and
35 `tag`.

```

36
37 MPI_INTERCOMM_CREATE_FROM_GROUPS(local_group, local_leader, remote_group,
38     remote_leader, stringtag, info, errhandler, newintercomm)
39
40     IN     local_group     local group (handle)
41     IN     local_leader   rank of local group leader in local_group (integer)
42     IN     remote_group   remote group, significant only at local_leader (handle)
43     IN     remote_leader  rank of remote group leader in remote_group,
44     significant only at local_leader (integer)
45
46     IN     stringtag     unique identifier for this operation (string)
47     IN     info           info object (handle)
48

```

IN	errhandler	error handler to be attached to new inter-communicator (handle)	1 2 3
OUT	newintercomm	new inter-communicator (handle)	4 5

C binding

```
int MPI_Intercomm_create_from_groups(MPI_Group local_group, int local_leader,
    MPI_Group remote_group, int remote_leader, const char *stringtag,
    MPI_Info info, MPI_Errhandler errhandler, MPI_Comm *newintercomm)
```

Fortran 2008 binding

```
MPI_Intercomm_create_from_groups(local_group, local_leader, remote_group,
    remote_leader, stringtag, info, errhandler, newintercomm, ierror)
    TYPE(MPI_Group), INTENT(IN) :: local_group, remote_group
    INTEGER, INTENT(IN) :: local_leader, remote_leader
    CHARACTER(LEN=*), INTENT(IN) :: stringtag
    TYPE(MPI_Info), INTENT(IN) :: info
    TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
    TYPE(MPI_Comm), INTENT(OUT) :: newintercomm
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_INTERCOMM_CREATE_FROM_GROUPS(LOCAL_GROUP, LOCAL_LEADER, REMOTE_GROUP,
    REMOTE_LEADER, STRINGTAG, INFO, ERRHANDLER, NEWINTERCOMM, IERROR)
    INTEGER LOCAL_GROUP, LOCAL_LEADER, REMOTE_GROUP, REMOTE_LEADER, INFO,
    ERRHANDLER, NEWINTERCOMM, IERROR
    CHARACTER*(*) STRINGTAG
```

This call creates an inter-communicator. Unlike `MPI_INTERCOMM_CREATE`, this function uses as input previously defined, disjoint local and remote groups. The calling MPI process must be a member of the local group. The call is collective over the union of the local and remote groups. All involved MPI processes shall provide an identical value for the `stringtag` argument. Within each group, all MPI processes shall provide identical `local_group`, `local_leader` arguments. Wildcards are not permitted for the `remote_leader` or `local_leader` arguments. The `stringtag` argument serves the same purpose as the `stringtag` used in the `MPI_COMM_CREATE_FROM_GROUP` function; it differentiates concurrent calls in a multithreaded environment. The `stringtag` shall not exceed `MPI_MAX_STRINGTAG_LEN` characters in length. For C, this includes space for a null terminating character.

`MPI_MAX_STRINGTAG_LEN` shall have a value of at least 63. In the event that `MPI_GROUP_EMPTY` is supplied as the `local_group` or `remote_group` or both, then the call is a local operation and `MPI_COMM_NULL` is returned as the `newintercomm`.

The `errhandler` argument specifies an error handler to be attached to the new inter-communicator. Section 9.3 specifies the error handler to be invoked if an error is encountered during the invocation of `MPI_INTERCOMM_CREATE_FROM_GROUPS`.

The `info` argument provides hints and assertions, possibly MPI implementation dependent, which indicate desired characteristics and guide communicator creation.

```

1 MPI_INTERCOMM_MERGE(intercomm, high, newintracomm)
2   IN      intercomm      inter-communicator (handle)
3
4   IN      high           ordering of the local and remote groups in the new
5                          intra-communicator (logical)
6
7   OUT     newintracomm   new intra-communicator (handle)

```

C binding

```

9 int MPI_Intercomm_merge(MPI_Comm intercomm, int high, MPI_Comm *newintracomm)

```

Fortran 2008 binding

```

11 MPI_Intercomm_merge(intercomm, high, newintracomm, ierror)
12   TYPE(MPI_Comm), INTENT(IN) :: intercomm
13   LOGICAL, INTENT(IN) :: high
14   TYPE(MPI_Comm), INTENT(OUT) :: newintracomm
15   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

17 MPI_INTERCOMM_MERGE(INTERCOMM, HIGH, NEWINTRACOMM, IERROR)
18   INTEGER INTERCOMM, NEWINTRACOMM, IERROR
19   LOGICAL HIGH

```

This function creates an intra-communicator from the union of the two groups that are associated with `intercomm`. All MPI processes should provide the same `high` value within each of the two groups. If MPI processes in one group provided the value `high = false` and MPI processes in the other group provided the value `high = true` then the union orders the “low” group before the “high” group. If all MPI processes provided the same `high` argument then the order of the union is arbitrary. This call is blocking and collective within the union of the two groups.

The error handler on the new inter-communicator in each MPI process is inherited from the communicator that contributes the local group. Note that this can result in different MPI processes in the same communicator having different error handlers.

Advice to implementors. The implementation of `MPI_INTERCOMM_MERGE`, `MPI_COMM_FREE`, and `MPI_COMM_DUP` are similar to the implementation of `MPI_INTERCOMM_CREATE`, except that contexts private to the input inter-communicator are used for communication between group leaders rather than contexts inside a bridge communicator. (*End of advice to implementors.*)

7.6.3 Inter-Communication Examples*Example 1: Three-Group “Pipeline”*

As shown in Figure 7.4, groups 0 and 1 communicate. Groups 1 and 2 communicate. Therefore, group 0 requires one inter-communicator, group 1 requires two inter-communicators, and group 2 requires 1 inter-communicator.

Example 7.13.

```

46 int main(int argc, char *argv[])
47 {
48   MPI_Comm myComm;      /* intra-communicator of local sub-group */

```




Figure 7.4: Three-group pipeline

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
MPI_Comm myFirstComm; /* inter-communicator */
MPI_Comm mySecondComm; /* second inter-communicator (group 1 only) */
int membershipKey;
int rank;

MPI_Init(&argc, &argv);
MPI_Comm_rank(MPI_COMM_WORLD, &rank);

/* User code must generate membershipKey in the range [0, 1, 2] */
membershipKey = rank % 3;

/* Build intra-communicator for local sub-group */
MPI_Comm_split(MPI_COMM_WORLD, membershipKey, rank, &myComm);

/* Build inter-communicators. Tags are hard-coded. */
if (membershipKey == 0)
{
    /* Group 0 communicates with group 1. */
    MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 1,
                        1, &myFirstComm);
}
else if (membershipKey == 1)
{
    /* Group 1 communicates with groups 0 and 2. */
    MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 0,
                        1, &myFirstComm);
    MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 2,
                        12, &mySecondComm);
}
else if (membershipKey == 2)
{
    /* Group 2 communicates with group 1. */
    MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 1,
                        12, &myFirstComm);
}

/* Do work ... */

switch(membershipKey) /* free communicators appropriately */
{
case 1:
    MPI_Comm_free(&mySecondComm);
case 0:
case 2:
    MPI_Comm_free(&myFirstComm);
    break;
}

MPI_Finalize();
return 0;

```




Figure 7.5: Three-group ring

```

11 }

```

Example 2: Three-Group “Ring”

As shown in Figure 7.5, groups 0 and 1 communicate. Groups 1 and 2 communicate. Groups 0 and 2 communicate. Therefore, each requires two inter-communicators.

Example 7.14.

```

19 int main(int argc, char *argv[])
20 {
21     MPI_Comm    myComm;        /* intra-communicator of local sub-group */
22     MPI_Comm    myFirstComm; /* inter-communicators */
23     MPI_Comm    mySecondComm;
24     int         membershipKey;
25     int         rank;
26
27     MPI_Init(&argc, &argv);
28     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
29     ...
30
31     /* User code must generate membershipKey in the range [0, 1, 2] */
32     membershipKey = rank % 3;
33
34     /* Build intra-communicator for local sub-group */
35     MPI_Comm_split(MPI_COMM_WORLD, membershipKey, rank, &myComm);
36
37     /* Build inter-communicators. Tags are hard-coded. */
38     if (membershipKey == 0)
39     {
40         /* Group 0 communicates with groups 1 and 2. */
41         MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 1,
42                             1, &myFirstComm);
43         MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 2,
44                             2, &mySecondComm);
45     }
46     else if (membershipKey == 1)
47     {
48         /* Group 1 communicates with groups 0 and 2. */
49         MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 0,
50                             1, &myFirstComm);
51         MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 2,
52                             12, &mySecondComm);
53     }
54 }

```



```

else if (membershipKey == 2)
{
    /* Group 2 communicates with groups 0 and 1. */
    MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 0,
                        2, &myFirstComm);
    MPI_Intercomm_create(myComm, 0, MPI_COMM_WORLD, 1,
                        12, &mySecondComm);
}

/* Do some work ... */

/* Then free communicators before terminating... */
MPI_Comm_free(&myFirstComm);
MPI_Comm_free(&mySecondComm);
MPI_Comm_free(&myComm);
MPI_Finalize();
return 0;
}

```

7.7 Caching

MPI provides a “caching” facility that allows an application to attach arbitrary pieces of information, called **attributes**, to three kinds of MPI objects: communicators, windows, and datatypes. More precisely, the caching facility allows a portable library to do the following:

- pass information between calls by associating it with an MPI intra- or inter-communicator, window, or datatype,
- quickly retrieve that information, and
- be guaranteed that out-of-date information is never retrieved, even if the object is freed and its handle subsequently reused by MPI.

The caching capabilities, in some form, are required by built-in MPI routines such as collective communication and application topology. Defining an interface to these capabilities as part of the MPI standard is valuable because it permits routines like collective communication and application topologies to be implemented as portable code, and also because it makes MPI more extensible by allowing user-written routines to use standard MPI calling sequences.

Advice to users. The communicator `MPI_COMM_SELF` is a suitable choice for posting MPI process-local attributes, via this attribute-caching mechanism. (*End of advice to users.*)

Rationale. In one extreme one can allow caching on all opaque handles. The other extreme is to only allow it on communicators. Caching has a cost associated with it and should only be allowed when it is clearly needed and the increased cost is modest. This is the reason that windows and datatypes were added but not other handles. (*End of rationale.*)

1 One difficulty is the potential for size differences between Fortran integers and C
2 pointers. For this reason, the Fortran versions of these routines use integers of kind
3 MPI_ADDRESS_KIND.

4
5 *Advice to implementors.* High-quality implementations should raise an error when
6 a keyval that was created by a call to MPI_XXX_CREATE_KEYVAL is used with an
7 object of the wrong type with a call to MPI_YYY_GET_ATTR, MPI_YYY_SET_ATTR,
8 MPI_YYY_DELETE_ATTR, or MPI_YYY_FREE_KEYVAL. To do so, it is necessary to
9 maintain, with each keyval, information on the type of the associated user function.
10 (*End of advice to implementors.*)

11 7.7.1 Functionality

12 Attributes can be attached to communicators, windows, and datatypes. Attributes are
13 local to the MPI process and specific to the communicator to which they are attached.
14 Attributes are not propagated by MPI from one communicator to another except when the
15 communicator is duplicated using MPI_COMM_DUP, MPI_COMM_IDUP,
16 MPI_COMM_DUP_WITH_INFO, and MPI_COMM_IDUP_WITH_INFO (and even then the
17 application must give specific permission through callback functions for the attribute to be
18 copied. Please refer to Section 7.4.2 and Section 7.7.2 for attributes propagation rules).
19

20
21 *Advice to users.* Attributes in C are of type void*. Typically, such an attribute will
22 be a pointer to a structure that contains further information, or a handle to an MPI
23 object. In Fortran, attributes are of type INTEGER. Such attribute can be a handle to
24 an MPI object, or just an integer-valued attribute. (*End of advice to users.*)

25
26 *Advice to implementors.* Attributes are scalar values, equal in size to, or larger than
27 a C-language pointer. Attributes can always hold an MPI handle. (*End of advice to
28 implementors.*)

29
30 The caching interface defined here requires that attributes be stored by MPI opaquely
31 within a communicator, window, or datatype. Accessor functions include the following:

- 32 • obtain a key value (used to identify an attribute); the user specifies “callback” func-
33 tions by which MPI informs the application when the communicator is destroyed or
34 copied.
- 35 • store and retrieve the value of an attribute;

36
37
38 *Advice to implementors.* Caching and callback functions are only called synchronously,
39 in response to explicit application requests. This avoids problems that result from re-
40 peated crossings between user and system space. (This synchronous calling rule is a
41 general property of MPI.)

42 The choice of key values is under control of MPI. This allows MPI to optimize its
43 implementation of attribute sets. It also avoids conflict between independent modules
44 caching information on the same communicators.

45
46 A much smaller interface, consisting of just a callback facility, would allow the entire
47 caching facility to be implemented by portable code. However, with the minimal call-
48 back interface, some form of table searching is implied by the need to handle arbitrary

communicators. In contrast, the more complete interface defined here permits rapid access to attributes through the use of pointers in communicators (to find the attribute table) and cleverly chosen key values (to retrieve individual attributes). In light of the efficiency “hit” inherent in the minimal interface, the more complete interface defined here is seen to be superior. (*End of advice to implementors.*)

MPI provides the following services related to caching. They are all MPI process local.

7.7.2 Communicators

Functions for caching on communicators are:

`MPI_COMM_CREATE_KEYVAL(comm_copy_attr_fn, comm_delete_attr_fn, comm_keyval, extra_state)`

IN	<code>comm_copy_attr_fn</code>	copy callback function for <code>comm_keyval</code> (function)
IN	<code>comm_delete_attr_fn</code>	delete callback function for <code>comm_keyval</code> (function)
OUT	<code>comm_keyval</code>	key value for future access (integer)
IN	<code>extra_state</code>	extra state for callback function

C binding

```
int MPI_Comm_create_keyval(MPI_Comm_copy_attr_function *comm_copy_attr_fn,
                           MPI_Comm_delete_attr_function *comm_delete_attr_fn,
                           int *comm_keyval, void *extra_state)
```

Fortran 2008 binding

```
MPI_Comm_create_keyval(comm_copy_attr_fn, comm_delete_attr_fn, comm_keyval,
                       extra_state, ierror)
PROCEDURE(MPI_Comm_copy_attr_function) :: comm_copy_attr_fn
PROCEDURE(MPI_Comm_delete_attr_function) :: comm_delete_attr_fn
INTEGER, INTENT(OUT) :: comm_keyval
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_CREATE_KEYVAL(COMM_COPY_ATTR_FN, COMM_DELETE_ATTR_FN, COMM_KEYVAL,
                       EXTRA_STATE, IERROR)
EXTERNAL COMM_COPY_ATTR_FN, COMM_DELETE_ATTR_FN
INTEGER COMM_KEYVAL, IERROR
INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
```

Generates a new attribute key. Keys are locally unique in an MPI process, and opaque to user, though they are explicitly stored in integers. Once allocated, the key value can be used to associate attributes and access them on any locally defined communicator.

The C callback functions are:

```
typedef int MPI_Comm_copy_attr_function(MPI_Comm oldcomm, int comm_keyval,
                                       void *extra_state, void *attribute_val_in,
                                       void *attribute_val_out, int *flag);
```

1 and
 2 typedef int MPI_Comm_delete_attr_function(MPI_Comm comm, int comm_keyval,
 3 void *attribute_val, void *extra_state);
 4
 5 which are the same as the MPI-1.1 calls but with a new name. The old names are deprecated.
 6 With the `mpi_f08` module, the Fortran callback functions are:

```
7 ABSTRACT INTERFACE
8   SUBROUTINE MPI_Comm_copy_attr_function(oldcomm, comm_keyval, extra_state,
9     attribute_val_in, attribute_val_out, flag, ierror)
10  TYPE(MPI_Comm) :: oldcomm
11  INTEGER :: comm_keyval, ierror
12  INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
13    attribute_val_out
14  LOGICAL :: flag
```

```
15 and
16 ABSTRACT INTERFACE
17   SUBROUTINE MPI_Comm_delete_attr_function(comm, comm_keyval, attribute_val,
18     extra_state, ierror)
19   TYPE(MPI_Comm) :: comm
20   INTEGER :: comm_keyval, ierror
21   INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
```

22 With the `mpi` module and (deprecated) `mpif.h` include file, the Fortran callback functions
 23 are:

```
24 SUBROUTINE COMM_COPY_ATTR_FUNCTION(OLDCOMM, COMM_KEYVAL, EXTRA_STATE,
25   ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)
26   INTEGER OLDCOMM, COMM_KEYVAL, IERROR
27   INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
28     ATTRIBUTE_VAL_OUT
29   LOGICAL FLAG
```

```
30 and
31 SUBROUTINE COMM_DELETE_ATTR_FUNCTION(COMM, COMM_KEYVAL, ATTRIBUTE_VAL,
32   EXTRA_STATE, IERROR)
33   INTEGER COMM, COMM_KEYVAL, IERROR
34   INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE
```

36 The `comm_copy_attr_fn` function is invoked when a communicator is duplicated by
 37 `MPI_COMM_DUP`, `MPI_COMM_IDUP`, `MPI_COMM_DUP_WITH_INFO` or
 38 `MPI_COMM_IDUP_WITH_INFO`. `comm_copy_attr_fn` should be of type
 39 `MPI_Comm_copy_attr_function`. The copy callback function is invoked for each key value in
 40 `oldcomm` in arbitrary order. Each call to the copy callback is made with a key value and its
 41 corresponding attribute. If it returns `flag = 0` or `.FALSE.`, then the attribute is deleted in
 42 the duplicated communicator. Otherwise (`flag = 1` or `.TRUE.`), the new attribute value is set
 43 to the value returned in `attribute_val_out`. The function returns `MPI_SUCCESS` on success
 44 and an error code on failure (in which case `MPI_COMM_DUP` or `MPI_COMM_IDUP` will
 45 fail).

46 The argument `comm_copy_attr_fn` may be specified as `MPI_COMM_NULL_COPY_FN`
 47 or `MPI_COMM_DUP_FN` from either C or Fortran. `MPI_COMM_NULL_COPY_FN` is a
 48 function that does nothing other than returning `flag = 0` or `.FALSE.` (depending on whether

the keyval was created with a C or Fortran binding to `MPI_COMM_CREATE_KEYVAL`) and `MPI_SUCCESS`. `MPI_COMM_DUP_FN` is a simple copy function that sets `flag = 1` or `.TRUE.`, returns the value of `attribute_val_in` in `attribute_val_out`, and returns `MPI_SUCCESS`. These replace the MPI-1 predefined callbacks `MPI_NULL_COPY_FN` and `MPI_DUP_FN`, whose use is deprecated.

Advice to users. Even though both formal arguments `attribute_val_in` and `attribute_val_out` are of type `void*`, their usage differs. The C copy function is passed by MPI in `attribute_val_in` the *value* of the attribute, and in `attribute_val_out` the *address* of the attribute, so as to allow the function to return the (new) attribute value. The use of type `void*` for both is to avoid messy type casts.

A valid copy function is one that completely duplicates the information by making a full duplicate copy of the data structures implied by an attribute; another might just make another reference to that data structure, while using a reference-count mechanism. Other types of attributes might not copy at all (they might be specific to `oldcomm` only). (*End of advice to users.*)

Advice to implementors. A C interface should be assumed for copy and delete functions associated with key values created in C; a Fortran calling interface should be assumed for key values created in Fortran. (*End of advice to implementors.*)

Analogous to `comm_copy_attr_fn` is a callback deletion function, defined as follows. The `comm_delete_attr_fn` function is invoked when a communicator is deleted by `MPI_COMM_FREE`, `MPI_COMM_DISCONNECT` or when a call is made explicitly to `MPI_COMM_DELETE_ATTR`. `comm_delete_attr_fn` should be of type `MPI_Comm_delete_attr_function`.

This function is called by `MPI_COMM_FREE`, `MPI_COMM_DISCONNECT`, `MPI_COMM_DELETE_ATTR`, and `MPI_COMM_SET_ATTR` to do whatever is needed to remove an attribute. The function returns `MPI_SUCCESS` on success and an error code on failure (in which case `MPI_COMM_FREE` will fail).

The argument `comm_delete_attr_fn` may be specified as `MPI_COMM_NULL_DELETE_FN` from either C or Fortran. `MPI_COMM_NULL_DELETE_FN` is a function that does nothing, other than returning `MPI_SUCCESS`. `MPI_COMM_NULL_DELETE_FN` replaces `MPI_NULL_DELETE_FN`, whose use is deprecated.

If an attribute copy function or attribute delete function returns other than `MPI_SUCCESS`, then the call that caused it to be invoked (for example, `MPI_COMM_FREE`) is erroneous.

The special key value `MPI_KEYVAL_INVALID` is never returned by `MPI_COMM_CREATE_KEYVAL`. Therefore, it can be used for static initialization of key values.

Advice to implementors. The predefined Fortran functions `MPI_COMM_NULL_COPY_FN`, `MPI_COMM_DUP_FN`, and `MPI_COMM_NULL_DELETE_FN` are defined in the `mpi` module (and deprecated `mpif.h`) and the `mpi_f08` module with the same name, but with different interfaces. Each function can coexist twice with the same name in the same MPI library, one routine as an implicit interface outside of the `mpi` module, i.e., declared as `EXTERNAL`,

1 and the other routine within `mpi_f08` declared with `CONTAINS`. These routines have
 2 different link names, which are also different to the link names used for the routines
 3 used in C. (*End of advice to implementors.*)
 4

5 *Advice to users.* Callbacks, including the predefined Fortran functions
 6 `MPI_COMM_NULL_COPY_FN`, `MPI_COMM_DUP_FN`, and
 7 `MPI_COMM_NULL_DELETE_FN` should not be passed from one application routine
 8 that uses the `mpi_f08` module to another application routine that uses the `mpi` module
 9 or (deprecated) `mpif.h` include file, and vice versa; see also the advice to users on
 10 page 842. (*End of advice to users.*)
 11

12
 13 `MPI_COMM_FREE_KEYVAL(comm_keyval)`

14 `INOUT comm_keyval` key value (integer)
 15
 16

17 **C binding**

18 `int MPI_Comm_free_keyval(int *comm_keyval)`

19 **Fortran 2008 binding**

20 `MPI_Comm_free_keyval(comm_keyval, ierror)`

21 `INTEGER, INTENT(INOUT) :: comm_keyval`

22 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`
 23

24 **Fortran binding**

25 `MPI_COMM_FREE_KEYVAL(COMM_KEYVAL, IERROR)`

26 `INTEGER COMM_KEYVAL, IERROR`
 27

28 Frees an extant attribute key. This function sets the value of `keyval` to
 29 `MPI_KEYVAL_INVALID`. Note that it is not erroneous to free an attribute key that is in use,
 30 because the actual free does not transpire until after all references (in other communicators
 31 on the MPI process) to the key have been freed. These references need to be explicitly freed
 32 by the program, either via calls to `MPI_COMM_DELETE_ATTR` that free one attribute
 33 instance, or by calls to `MPI_COMM_FREE` that free all attribute instances associated with
 34 the freed communicator.
 35

36 `MPI_COMM_SET_ATTR(comm, comm_keyval, attribute_val)`

37 `INOUT comm` communicator to which attribute will be attached
 38 (handle)
 39

40 `IN comm_keyval` key value (integer)

41 `IN attribute_val` attribute value
 42

43 **C binding**

44 `int MPI_Comm_set_attr(MPI_Comm comm, int comm_keyval, void *attribute_val)`

45 **Fortran 2008 binding**

46 `MPI_Comm_set_attr(comm, comm_keyval, attribute_val, ierror)`

47 `TYPE(MPI_Comm), INTENT(IN) :: comm`
 48

```

INTEGER, INTENT(IN) :: comm_keyval
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: attribute_val
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_COMM_SET_ATTR(COMM, COMM_KEYVAL, ATTRIBUTE_VAL, IERROR)
  INTEGER COMM, COMM_KEYVAL, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL

```

This function stores the stipulated attribute value `attribute_val` for subsequent retrieval by `MPI_COMM_GET_ATTR`. If the value is already present, then the outcome is as if `MPI_COMM_DELETE_ATTR` was first called to delete the previous value (and the callback function `comm_delete_attr_fn` was executed), and a new value was next stored. The call is erroneous if there is no key with value `keyval`; in particular `MPI_KEYVAL_INVALID` is an erroneous key value. The call will fail if the `comm_delete_attr_fn` function returned an error code other than `MPI_SUCCESS`.

```

MPI_COMM_GET_ATTR(comm, comm_keyval, attribute_val, flag)

```

IN	comm	communicator to which the attribute is attached (handle)
IN	comm_keyval	key value (integer)
OUT	attribute_val	attribute value, unless flag = false
OUT	flag	false if no attribute is associated with the key (logical)

C binding

```

int MPI_Comm_get_attr(MPI_Comm comm, int comm_keyval, void *attribute_val,
  int *flag)

```

Fortran 2008 binding

```

MPI_Comm_get_attr(comm, comm_keyval, attribute_val, flag, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: comm_keyval
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: attribute_val
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_COMM_GET_ATTR(COMM, COMM_KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)
  INTEGER COMM, COMM_KEYVAL, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL
  LOGICAL FLAG

```

Retrieves attribute value by key. The call is erroneous if there is no key with value `keyval`. On the other hand, the call is correct if the key value exists, but no attribute is attached on `comm` for that key; in such case, the call returns `flag = false`. In particular `MPI_KEYVAL_INVALID` is an erroneous key value.

Advice to users. The call to `MPI_Comm_set_attr` passes in `attribute_val` the *value* of the attribute; the call to `MPI_Comm_get_attr` passes in `attribute_val` the *address* of the location where the attribute value is to be returned. Thus, if the attribute value itself is a pointer of type `void*`, then the actual `attribute_val` parameter to `MPI_Comm_set_attr` will be of type `void*` and the actual `attribute_val` parameter to `MPI_Comm_get_attr` will be of type `void**`. (*End of advice to users.*)

Rationale. The use of a formal parameter `attribute_val` of type `void*` (rather than `void**`) avoids the messy type casting that would be needed if the attribute value is declared with a type other than `void*`. (*End of rationale.*)

`MPI_COMM_DELETE_ATTR(comm, comm_keyval)`

INOUT	comm	communicator from which the attribute is deleted (handle)
IN	comm_keyval	key value (integer)

C binding

```
int MPI_Comm_delete_attr(MPI_Comm comm, int comm_keyval)
```

Fortran 2008 binding

```
MPI_Comm_delete_attr(comm, comm_keyval, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: comm_keyval
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_DELETE_ATTR(COMM, COMM_KEYVAL, IERROR)
  INTEGER COMM, COMM_KEYVAL, IERROR
```

Delete attribute from cache by key. This function invokes the attribute delete function `comm_delete_attr_fn` specified when the `keyval` was created. The call will fail if the `comm_delete_attr_fn` function returns an error code other than `MPI_SUCCESS`.

Whenever a communicator is replicated using the function `MPI_COMM_DUP`, `MPI_COMM_IDUP`, `MPI_COMM_DUP_WITH_INFO` or `MPI_COMM_IDUP_WITH_INFO`, all call-back copy functions for attributes that are currently set are invoked (in arbitrary order). Whenever a communicator is deleted using the function `MPI_COMM_FREE` all callback delete functions for attributes that are currently set are invoked.

7.7.3 Windows

The functions for caching on windows are:

MPI_WIN_CREATE_KEYVAL(win_copy_attr_fn, win_delete_attr_fn, win_keyval, extra_state) 1

IN	win_copy_attr_fn	copy callback function for win_keyval (function)	3
IN	win_delete_attr_fn	delete callback function for win_keyval (function)	4
OUT	win_keyval	key value for future access (integer)	5
IN	extra_state	extra state for callback function	6

C binding 7

```
int MPI_Win_create_keyval(MPI_Win_copy_attr_function *win_copy_attr_fn,
                          MPI_Win_delete_attr_function *win_delete_attr_fn,
                          int *win_keyval, void *extra_state) 8
```

Fortran 2008 binding 9

```
MPI_Win_create_keyval(win_copy_attr_fn, win_delete_attr_fn, win_keyval,
                      extra_state, ierror) 10
PROCEDURE(MPI_Win_copy_attr_function) :: win_copy_attr_fn 11
PROCEDURE(MPI_Win_delete_attr_function) :: win_delete_attr_fn 12
INTEGER, INTENT(OUT) :: win_keyval 13
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state 14
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 15
```

Fortran binding 16

```
MPI_WIN_CREATE_KEYVAL(WIN_COPY_ATTR_FN, WIN_DELETE_ATTR_FN, WIN_KEYVAL,
                       EXTRA_STATE, IERROR) 17
EXTERNAL WIN_COPY_ATTR_FN, WIN_DELETE_ATTR_FN 18
INTEGER WIN_KEYVAL, IERROR 19
INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE 20
```

The argument `win_copy_attr_fn` may be specified as `MPI_WIN_NULL_COPY_FN` or `MPI_WIN_DUP_FN` from either C or Fortran. `MPI_WIN_NULL_COPY_FN` is a function that does nothing other than returning `flag = 0` and `MPI_SUCCESS`. `MPI_WIN_DUP_FN` is a simple copy function that sets `flag = 1`, returns the value of `attribute_val_in` in `attribute_val_out`, and returns `MPI_SUCCESS`.

The argument `win_delete_attr_fn` may be specified as `MPI_WIN_NULL_DELETE_FN` from either C or Fortran. `MPI_WIN_NULL_DELETE_FN` is a function that does nothing, other than returning `MPI_SUCCESS`.

The C callback functions are:

```
typedef int MPI_Win_copy_attr_function(MPI_Win oldwin, int win_keyval,
                                       void *extra_state, void *attribute_val_in,
                                       void *attribute_val_out, int *flag); 21
```

and 22

```
typedef int MPI_Win_delete_attr_function(MPI_Win win, int win_keyval,
                                       void *attribute_val, void *extra_state); 23
```

With the `mpi_f08` module, the Fortran callback functions are: 24

```
ABSTRACT INTERFACE 25
  SUBROUTINE MPI_Win_copy_attr_function(oldwin, win_keyval, extra_state,
                                       attribute_val_in, attribute_val_out, flag, ierror) 26
```

```

1     TYPE(MPI_Win) :: oldwin
2     INTEGER :: win_keyval, ierror
3     INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
4         attribute_val_out
5     LOGICAL :: flag
6
7     and
8     ABSTRACT INTERFACE
9         SUBROUTINE MPI_Win_delete_attr_function(win, win_keyval, attribute_val,
10             extra_state, ierror)
11         TYPE(MPI_Win) :: win
12         INTEGER :: win_keyval, ierror
13         INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state

```

With the `mpi` module and (deprecated) `mpif.h` include file, the Fortran callback functions are:

```

14     SUBROUTINE WIN_COPY_ATTR_FUNCTION(OLDWIN, WIN_KEYVAL, EXTRA_STATE,
15         ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)
16         INTEGER OLDWIN, WIN_KEYVAL, IERROR
17         INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
18             ATTRIBUTE_VAL_OUT
19         LOGICAL FLAG

```

```

20     and
21     SUBROUTINE WIN_DELETE_ATTR_FUNCTION(WIN, WIN_KEYVAL, ATTRIBUTE_VAL,
22         EXTRA_STATE, IERROR)
23         INTEGER WIN, WIN_KEYVAL, IERROR
24         INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE

```

If an attribute copy function or attribute delete function returns other than `MPI_SUCCESS`, then the call that caused it to be invoked (for example, `MPI_WIN_FREE`), is erroneous.

```

25     MPI_WIN_FREE_KEYVAL(win_keyval)
26         INOUT    win_keyval                key value (integer)

```

C binding

```

27     int MPI_Win_free_keyval(int *win_keyval)

```

Fortran 2008 binding

```

28     MPI_Win_free_keyval(win_keyval, ierror)
29         INTEGER, INTENT(INOUT) :: win_keyval
30         INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

31     MPI_WIN_FREE_KEYVAL(WIN_KEYVAL, IERROR)
32         INTEGER WIN_KEYVAL, IERROR

```

MPI_WIN_SET_ATTR(win, win_keyval, attribute_val)	1
INOUT win window to which attribute will be attached (handle)	2
IN win_keyval key value (integer)	3
IN attribute_val attribute value	4
	5
	6
	7
C binding	8
int MPI_Win_set_attr(MPI_Win win, int win_keyval, void *attribute_val)	9
	10
Fortran 2008 binding	11
MPI_Win_set_attr(win, win_keyval, attribute_val, ierror)	12
TYPE(MPI_Win), INTENT(IN) :: win	13
INTEGER, INTENT(IN) :: win_keyval	14
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: attribute_val	15
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	16
	17
Fortran binding	18
MPI_WIN_SET_ATTR(WIN, WIN_KEYVAL, ATTRIBUTE_VAL, IERROR)	19
INTEGER WIN, WIN_KEYVAL, IERROR	20
INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL	21
	22
MPI_WIN_GET_ATTR(win, win_keyval, attribute_val, flag)	23
IN win window to which the attribute is attached (handle)	24
IN win_keyval key value (integer)	25
OUT attribute_val attribute value, unless flag = false	26
OUT flag false if no attribute is associated with the key	27
(logical)	28
	29
	30
	31
C binding	32
int MPI_Win_get_attr(MPI_Win win, int win_keyval, void *attribute_val,	33
int *flag)	34
	35
Fortran 2008 binding	36
MPI_Win_get_attr(win, win_keyval, attribute_val, flag, ierror)	37
TYPE(MPI_Win), INTENT(IN) :: win	38
INTEGER, INTENT(IN) :: win_keyval	39
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: attribute_val	40
LOGICAL, INTENT(OUT) :: flag	41
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	42
	43
Fortran binding	44
MPI_WIN_GET_ATTR(WIN, WIN_KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)	45
INTEGER WIN, WIN_KEYVAL, IERROR	46
INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL	47
LOGICAL FLAG	48

```

1 MPI_WIN_DELETE_ATTR(win, win_keyval)
2     INOUT    win                window from which the attribute is deleted (handle)
3
4     IN      win_keyval          key value (integer)

```

C binding

```
int MPI_Win_delete_attr(MPI_Win win, int win_keyval)
```

Fortran 2008 binding

```

MPI_Win_delete_attr(win, win_keyval, ierror)
    TYPE(MPI_Win), INTENT(IN) :: win
    INTEGER, INTENT(IN) :: win_keyval
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_WIN_DELETE_ATTR(WIN, WIN_KEYVAL, IERROR)
    INTEGER WIN, WIN_KEYVAL, IERROR

```

7.7.4 Datatypes

The new functions for caching on datatypes are:

```

MPI_TYPE_CREATE_KEYVAL(type_copy_attr_fn, type_delete_attr_fn, type_keyval,
    extra_state)
    IN      type_copy_attr_fn      copy callback function for type_keyval (function)
    IN      type_delete_attr_fn    delete callback function for type_keyval (function)
    OUT     type_keyval            key value for future access (integer)
    IN      extra_state            extra state for callback function

```

C binding

```

int MPI_Type_create_keyval(MPI_Type_copy_attr_function *type_copy_attr_fn,
    MPI_Type_delete_attr_function *type_delete_attr_fn,
    int *type_keyval, void *extra_state)

```

Fortran 2008 binding

```

MPI_Type_create_keyval(type_copy_attr_fn, type_delete_attr_fn, type_keyval,
    extra_state, ierror)
    PROCEDURE(MPI_Type_copy_attr_function) :: type_copy_attr_fn
    PROCEDURE(MPI_Type_delete_attr_function) :: type_delete_attr_fn
    INTEGER, INTENT(OUT) :: type_keyval
    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TYPE_CREATE_KEYVAL(TYPE_COPY_ATTR_FN, TYPE_DELETE_ATTR_FN, TYPE_KEYVAL,
    EXTRA_STATE, IERROR)
    EXTERNAL TYPE_COPY_ATTR_FN, TYPE_DELETE_ATTR_FN
    INTEGER TYPE_KEYVAL, IERROR

```

```
INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
```

The argument `type_copy_attr_fn` may be specified as `MPI_TYPE_NULL_COPY_FN` or `MPI_TYPE_DUP_FN` from either C or Fortran. `MPI_TYPE_NULL_COPY_FN` is a function that does nothing other than returning `flag = 0` and `MPI_SUCCESS`. `MPI_TYPE_DUP_FN` is a simple copy function that sets `flag = 1`, returns the value of `attribute_val_in` in `attribute_val_out`, and returns `MPI_SUCCESS`.

The argument `type_delete_attr_fn` may be specified as `MPI_TYPE_NULL_DELETE_FN` from either C or Fortran. `MPI_TYPE_NULL_DELETE_FN` is a function that does nothing, other than returning `MPI_SUCCESS`.

The C callback functions are:

```
typedef int MPI_Type_copy_attr_function(MPI_Datatype oldtype, int type_keyval,
                                       void *extra_state, void *attribute_val_in,
                                       void *attribute_val_out, int *flag);
```

and

```
typedef int MPI_Type_delete_attr_function(MPI_Datatype datatype,
                                         int type_keyval, void *attribute_val, void *extra_state);
```

With the `mpi_f08` module, the Fortran callback functions are:

```
ABSTRACT INTERFACE
  SUBROUTINE MPI_Type_copy_attr_function(oldtype, type_keyval, extra_state,
                                         attribute_val_in, attribute_val_out, flag, ierror)
    TYPE(MPI_Datatype) :: oldtype
    INTEGER :: type_keyval, ierror
    INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
                                     attribute_val_out
    LOGICAL :: flag
```

and

```
ABSTRACT INTERFACE
  SUBROUTINE MPI_Type_delete_attr_function(datatype, type_keyval,
                                         attribute_val, extra_state, ierror)
    TYPE(MPI_Datatype) :: datatype
    INTEGER :: type_keyval, ierror
    INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
```

With the `mpi` module and (deprecated) `mpif.h` include file, the Fortran callback functions are:

```
SUBROUTINE TYPE_COPY_ATTR_FUNCTION(OLDTYPE, TYPE_KEYVAL, EXTRA_STATE,
                                   ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)
  INTEGER OLDTYPE, TYPE_KEYVAL, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
                                   ATTRIBUTE_VAL_OUT
  LOGICAL FLAG
```

and

```
SUBROUTINE TYPE_DELETE_ATTR_FUNCTION(DATATYPE, TYPE_KEYVAL, ATTRIBUTE_VAL,
                                     EXTRA_STATE, IERROR)
  INTEGER DATATYPE, TYPE_KEYVAL, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE
```

1 If an attribute copy function or attribute delete function returns other than
 2 MPI_SUCCESS, then the call that caused it to be invoked (for example, MPI_TYPE_FREE),
 3 is erroneous.

4
 5
 6 MPI_TYPE_FREE_KEYVAL(type_keyval)

7 INOUT type_keyval key value (integer)

8
 9 **C binding**

10 int MPI_Type_free_keyval(int *type_keyval)

11
 12 **Fortran 2008 binding**

13 MPI_Type_free_keyval(type_keyval, ierror)

14 INTEGER, INTENT(INOUT) :: type_keyval

15 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

16 **Fortran binding**

17 MPI_TYPE_FREE_KEYVAL(TYPE_KEYVAL, IERROR)

18 INTEGER TYPE_KEYVAL, IERROR

19
 20
 21 MPI_TYPE_SET_ATTR(datatype, type_keyval, attribute_val)

22 INOUT datatype datatype to which attribute will be attached (handle)

23 IN type_keyval key value (integer)

24 IN attribute_val attribute value

25
 26
 27 **C binding**

28 int MPI_Type_set_attr(MPI_Datatype datatype, int type_keyval,
 29 void *attribute_val)

30
 31 **Fortran 2008 binding**

32 MPI_Type_set_attr(datatype, type_keyval, attribute_val, ierror)

33 TYPE(MPI_Datatype), INTENT(IN) :: datatype

34 INTEGER, INTENT(IN) :: type_keyval

35 INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: attribute_val

36 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

37 **Fortran binding**

38 MPI_TYPE_SET_ATTR(DATATYPE, TYPE_KEYVAL, ATTRIBUTE_VAL, IERROR)

39 INTEGER DATATYPE, TYPE_KEYVAL, IERROR

40 INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL

MPI_TYPE_GET_ATTR(datatype, type_keyval, attribute_val, flag)				1
IN	datatype	datatype to which the attribute is attached (handle)		2
IN	type_keyval	key value (integer)		3
OUT	attribute_val	attribute value, unless flag = false		4
OUT	flag	false if no attribute is associated with the key (logical)		5
				6
				7
				8
				9

C binding

```
int MPI_Type_get_attr(MPI_Datatype datatype, int type_keyval,
                     void *attribute_val, int *flag)
```

Fortran 2008 binding

```
MPI_Type_get_attr(datatype, type_keyval, attribute_val, flag, ierror)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER, INTENT(IN) :: type_keyval
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: attribute_val
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_GET_ATTR(DATATYPE, TYPE_KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)
  INTEGER DATATYPE, TYPE_KEYVAL, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL
  LOGICAL FLAG
```

```
MPI_TYPE_DELETE_ATTR(datatype, type_keyval)
```

INOUT	datatype	datatype from which the attribute is deleted (handle)	29
IN	type_keyval	key value (integer)	30

C binding

```
int MPI_Type_delete_attr(MPI_Datatype datatype, int type_keyval)
```

Fortran 2008 binding

```
MPI_Type_delete_attr(datatype, type_keyval, ierror)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER, INTENT(IN) :: type_keyval
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_DELETE_ATTR(DATATYPE, TYPE_KEYVAL, IERROR)
  INTEGER DATATYPE, TYPE_KEYVAL, IERROR
```

7.7.5 Error Class for Invalid Keyval

Key values for attributes are system-allocated, by MPI_{XXX}_CREATE_KEYVAL. Only such values can be passed to the functions that use

key values as input arguments. In order to signal that an erroneous key value has been passed to one of these functions, there is a new MPI error class: `MPI_ERR_KEYVAL`. It can be returned by `MPI_ATTR_PUT`, `MPI_ATTR_GET`, `MPI_ATTR_DELETE`, `MPI_KEYVAL_FREE`, `MPI_{XXX}_DELETE_ATTR`, `MPI_{XXX}_SET_ATTR`, `MPI_{XXX}_GET_ATTR`, `MPI_{XXX}_FREE_KEYVAL`, `MPI_COMM_DUP`, `MPI_COMM_IDUP`, `MPI_COMM_DUP_WITH_INFO`, `MPI_COMM_IDUP_WITH_INFO`, `MPI_COMM_DISCONNECT`, and `MPI_COMM_FREE`. The last six are included because `keyval` is an argument to the copy and delete functions for attributes.

7.7.6 Attributes Example

Advice to users. This example shows how to write a collective communication operation that uses caching to be more efficient after the first call. (*End of advice to users.*)

```

16 /* key for this module's stuff: */
17 static int gop_key = MPI_KEYVAL_INVALID;
18
19 typedef struct
20 {
21     int ref_count;          /* reference count */
22     /* other stuff, whatever else we want */
23 } gop_stuff_type;
24
25 void Efficient_Collective_Op(MPI_Comm comm, ...)
26 {
27     gop_stuff_type *gop_stuff;
28     MPI_Group      group;
29     int            foundflag;
30
31     MPI_Comm_group(comm, &group);
32
33     if (gop_key == MPI_KEYVAL_INVALID) /* get a key on first call ever */
34     {
35         if ( ! MPI_Comm_create_keyval(gop_stuff_copier,
36                                     gop_stuff_destructor,
37                                     &gop_key, NULL)) {
38             /* get the key while assigning its copy and delete callback
39              behavior. */
40         } else
41             MPI_Abort(comm, 99);
42     }
43
44     MPI_Comm_get_attr(comm, gop_key, &gop_stuff, &foundflag);
45     if (foundflag)
46     { /* This module has executed in this group before.
47        We will use the cached information */
48     }
49     else
50     { /* This is a group that we have not yet cached anything in.
51        We will now do so.
52     */

```



```

1
2  /* First, allocate storage for the stuff we want,
3     and initialize the reference count */
4
5  gop_stuff = (gop_stuff_type *) malloc(sizeof(gop_stuff_type));
6  if (gop_stuff == NULL) { /* abort on out-of-memory error */ }
7
8  gop_stuff->ref_count = 1;
9
10 /* Second, fill in *gop_stuff with whatever we want.
11     This part isn't shown here */
12
13 /* Third, store gop_stuff as the attribute value */
14 MPI_Comm_set_attr(comm, gop_key, gop_stuff);
15 }
16 /* Then, in any case, use contents of *gop_stuff
17     to do the global op ... */
18 }
19
20 /* The following routine is called by MPI when a group is freed */
21
22 int gop_stuff_destructor(MPI_Comm comm, int keyval, void *gop_stuffP,
23                          void *extra)
24 {
25     gop_stuff_type *gop_stuff = (gop_stuff_type *)gop_stuffP;
26     if (keyval != gop_key) { /* abort -- programming error */ }
27
28     /* The group's being freed removes one reference to gop_stuff */
29     gop_stuff->ref_count -= 1;
30
31     /* If no references remain, then free the storage */
32     if (gop_stuff->ref_count == 0) {
33         free((void *)gop_stuff);
34     }
35     return MPI_SUCCESS;
36 }
37
38 /* The following routine is called by MPI when a group is copied */
39
40 int gop_stuff_copier(MPI_Comm comm, int keyval, void *extra,
41                      void *gop_stuff_inP, void *gop_stuff_outP, int *flag)
42 {
43     gop_stuff_type *gop_stuff_in = (gop_stuff_type *)gop_stuff_inP;
44     gop_stuff_type **gop_stuff_out = (gop_stuff_type **)gop_stuff_outP;
45     if (keyval != gop_key) { /* abort -- programming error */ }
46
47     /* The new group adds one reference to this gop_stuff */
48     gop_stuff_in->ref_count += 1;
49     *gop_stuff_out = gop_stuff_in;
50     return MPI_SUCCESS;
51 }

```

7.8 Naming Objects

There are many occasions on which it would be useful to allow a user to associate a printable identifier with an MPI communicator, window, or datatype, for instance error reporting, debugging, and profiling. The names attached to opaque objects do not propagate when the object is duplicated or copied by MPI routines. For communicators this can be achieved using the following two functions.

```
MPI_COMM_SET_NAME(comm, comm_name)
```

INOUT	comm	communicator whose identifier is to be set (handle)
IN	comm_name	the character string that is remembered as the name (string)

C binding

```
int MPI_Comm_set_name(MPI_Comm comm, const char *comm_name)
```

Fortran 2008 binding

```
MPI_Comm_set_name(comm, comm_name, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  CHARACTER(LEN=*), INTENT(IN) :: comm_name
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_SET_NAME(COMM, COMM_NAME, IERROR)
  INTEGER COMM, IERROR
  CHARACTER*(*) COMM_NAME
```

`MPI_COMM_SET_NAME` allows a user to associate a name string with a communicator. The character string that is passed to `MPI_COMM_SET_NAME` will be saved inside the MPI library (so it can be freed by the caller immediately after the call, or allocated on the stack). Leading spaces in name are significant but trailing ones are not.

`MPI_COMM_SET_NAME` is a local (noncollective) operation, which only affects the name of the communicator as seen in the MPI process that made the `MPI_COMM_SET_NAME` call. There is no requirement that the same (or any) name be assigned to a communicator in every MPI process where it exists.

Advice to users. Since `MPI_COMM_SET_NAME` is provided to help debug code, it is sensible to give the same name to a communicator in all of the MPI processes where it exists, to avoid confusion. (*End of advice to users.*)

The length of the name that can be stored is limited to the value of `MPI_MAX_OBJECT_NAME` in Fortran and `MPI_MAX_OBJECT_NAME-1` in C to allow for the null terminator. Attempts to put names longer than this will result in truncation of the name. `MPI_MAX_OBJECT_NAME` must have a value of at least 64.

Advice to users. Under circumstances of store exhaustion an attempt to put a name of any length could fail, therefore the value of `MPI_MAX_OBJECT_NAME` should be viewed only as a strict upper bound on the name length, not a guarantee that setting names of less than this length will always succeed. (*End of advice to users.*)

Advice to implementors. Implementations that pre-allocate a fixed size space for a name should use the length of that allocation as the value of `MPI_MAX_OBJECT_NAME`. Implementations that allocate space for the name from the heap should still define `MPI_MAX_OBJECT_NAME` to be a relatively small value, since the user has to allocate space for a string of up to this size when calling `MPI_COMM_GET_NAME`. (*End of advice to implementors.*)

`MPI_COMM_GET_NAME(comm, comm_name, resultlen)`

IN	<code>comm</code>	communicator whose name is to be returned (handle)
OUT	<code>comm_name</code>	the name previously stored on the communicator, or an empty string if no such name exists (string)
OUT	<code>resultlen</code>	length of returned name (integer)

C binding

```
int MPI_Comm_get_name(MPI_Comm comm, char *comm_name, int *resultlen)
```

Fortran 2008 binding

```
MPI_Comm_get_name(comm, comm_name, resultlen, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  CHARACTER(LEN=MPI_MAX_OBJECT_NAME), INTENT(OUT) :: comm_name
  INTEGER, INTENT(OUT) :: resultlen
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_GET_NAME(COMM, COMM_NAME, RESULTLEN, IERROR)
  INTEGER COMM, RESULTLEN, IERROR
  CHARACTER*(*) COMM_NAME
```

`MPI_COMM_GET_NAME` returns the last name that has previously been associated with the given communicator. The name may be set and retrieved from any language. The same name will be returned independent of the language used. `comm_name` should be allocated so that it can hold a resulting string of length `MPI_MAX_OBJECT_NAME` characters. `MPI_COMM_GET_NAME` returns a copy of the set name in `comm_name`.

In C, a null character is additionally stored at `comm_name[resultlen]`. The value of `resultlen` cannot be larger than `MPI_MAX_OBJECT_NAME-1`. In Fortran, `comm_name` is padded on the right with blank characters. The value of `resultlen` cannot be larger than `MPI_MAX_OBJECT_NAME`.

If the user has not associated a name with a communicator, or an error occurs, `MPI_COMM_GET_NAME` will return an empty string (all spaces in Fortran, "" in C). The three predefined communicators will have predefined names associated with them. Thus, the names of `MPI_COMM_WORLD`, `MPI_COMM_SELF`, and the communicator returned by `MPI_COMM_GET_PARENT` (if not `MPI_COMM_NULL`) will have the default of "MPI_COMM_WORLD", "MPI_COMM_SELF", and "MPI_COMM_PARENT". Passing `MPI_COMM_NULL` as `comm` will return the string "MPI_COMM_NULL". The fact that the system may have chosen to give a default name to a communicator does not prevent the user from setting a name on the same communicator; doing this removes the old name and assigns the new one.

Rationale. We provide separate functions for setting and getting the name of a communicator, rather than simply providing a predefined attribute key for the following reasons:

- It is not, in general, possible to store a string as an attribute from Fortran.
- It is not easy to set up the delete function for a string attribute unless it is known to have been allocated from the heap.
- To make the attribute key useful additional code to call `strdup` is necessary. If this is not standardized then users have to write it. This is extra unneeded work that we can easily eliminate.
- The Fortran binding is not trivial to write (it will depend on details of the Fortran compilation system), and will not be portable. Therefore it should be in the library rather than in user code.

(End of rationale.)

Advice to users. The above definition means that it is safe simply to print the string returned by `MPI_COMM_GET_NAME`, as it is always a valid string even if there was no name.

Note that associating a name with a communicator has no effect on the semantics of an MPI program, and will (necessarily) increase the store requirement of the program, since the names must be saved. Therefore there is no requirement that users use these functions to associate names with communicators. However debugging and profiling MPI applications may be made easier if names are associated with communicators, since the debugger or profiler should then be able to present information in a less cryptic manner. *(End of advice to users.)*

The following functions are used for setting and getting names of datatypes. The constant `MPI_MAX_OBJECT_NAME` also applies to these names.

`MPI_TYPE_SET_NAME(datatype, type_name)`

<code>INOUT</code>	<code>datatype</code>	datatype whose identifier is to be set (handle)
<code>IN</code>	<code>type_name</code>	the character string that is remembered as the name (string)

C binding

```
int MPI_Type_set_name(MPI_Datatype datatype, const char *type_name)
```

Fortran 2008 binding

```
MPI_Type_set_name(datatype, type_name, ierror)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  CHARACTER(LEN=*) , INTENT(IN) :: type_name
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_SET_NAME(DATATYPE, TYPE_NAME, IERROR)
  INTEGER DATATYPE, IERROR
  CHARACTER*(*) TYPE_NAME
```

MPI_TYPE_GET_NAME(datatype, type_name, resultlen)			1
IN	datatype	datatype whose name is to be returned (handle)	2
OUT	type_name	the name previously stored on the datatype, or an empty string if no such name exists (string)	3
OUT	resultlen	length of returned name (integer)	4

C binding

```
int MPI_Type_get_name(MPI_Datatype datatype, char *type_name, int *resultlen)
```

Fortran 2008 binding

```
MPI_Type_get_name(datatype, type_name, resultlen, ierror)
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
CHARACTER(LEN=MPI_MAX_OBJECT_NAME), INTENT(OUT) :: type_name
```

```
INTEGER, INTENT(OUT) :: resultlen
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_GET_NAME(DATATYPE, TYPE_NAME, RESULTLEN, IERROR)
```

```
INTEGER DATATYPE, RESULTLEN, IERROR
```

```
CHARACTER*(*) TYPE_NAME
```

Named predefined datatypes have the default names of the datatype name. For example, MPI_WCHAR has the default name of "MPI_WCHAR". Passing MPI_DATATYPE_NULL as datatype will return the string "MPI_DATATYPE_NULL".

The following functions are used for setting and getting names of windows. The constant MPI_MAX_OBJECT_NAME also applies to these names.

```
MPI_WIN_SET_NAME(win, win_name)
```

INOUT	win	window whose identifier is to be set (handle)
-------	-----	---

IN	win_name	the character string that is remembered as the name (string)
----	----------	--

C binding

```
int MPI_Win_set_name(MPI_Win win, const char *win_name)
```

Fortran 2008 binding

```
MPI_Win_set_name(win, win_name, ierror)
```

```
TYPE(MPI_Win), INTENT(IN) :: win
```

```
CHARACTER(LEN=*), INTENT(IN) :: win_name
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_SET_NAME(WIN, WIN_NAME, IERROR)
```

```
INTEGER WIN, IERROR
```

```
CHARACTER*(*) WIN_NAME
```

```

1 MPI_WIN_GET_NAME(win, win_name, resultlen)
2     IN        win                window whose name is to be returned (handle)
3
4     OUT       win_name           the name previously stored on the window, or an
5                                     empty string if no such name exists (string)
6
7     OUT       resultlen         length of returned name (integer)

```

C binding

```

9 int MPI_Win_get_name(MPI_Win win, char *win_name, int *resultlen)

```

Fortran 2008 binding

```

12 MPI_Win_get_name(win, win_name, resultlen, ierror)
13     TYPE(MPI_Win), INTENT(IN) :: win
14     CHARACTER(LEN=MPI_MAX_OBJECT_NAME), INTENT(OUT) :: win_name
15     INTEGER, INTENT(OUT) :: resultlen
16     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

18 MPI_WIN_GET_NAME(WIN, WIN_NAME, RESULTLEN, IERROR)
19     INTEGER WIN, RESULTLEN, IERROR
20     CHARACTER*(*) WIN_NAME
21
22     Passing MPI_WIN_NULL as win will return the string "MPI_WIN_NULL".

```

7.9 Formalizing the Loosely Synchronous Model

In this section, we make further statements about the loosely synchronous model, with particular attention to intra-communication.

7.9.1 Basic Statements

When a caller passes a communicator (that contains a context and group) to a callee, that communicator must be free of side effects throughout execution of the subprogram: there should be no active operations on that communicator that might involve the MPI process. This provides one model in which libraries can be written, and work “safely.” For libraries so designated, the callee has permission to do whatever communication it likes with the communicator, and under the above guarantee knows that no other communications will interfere. Since we permit good implementations to create new communicators without synchronization (such as by preallocated contexts on communicators), this does not impose a significant overhead.

This form of safety is analogous to other common computer-science usages, such as passing a descriptor of an array to a library routine. The library routine has every right to expect such a descriptor to be valid and modifiable.

7.9.2 Models of Execution

In the loosely synchronous model, transfer of control to a **parallel procedure** is effected by having each executing MPI process invoke the procedure. The invocation is a collective operation: it is executed by all MPI processes in the execution group, and invocations are similarly ordered at all MPI processes. However, the invocation need not be synchronized.

We say that a parallel procedure is *active* in an MPI process if the MPI process belongs to a group that may collectively execute the procedure, and some member of that group is currently executing the procedure code. If a parallel procedure is active in an MPI process, then this MPI process may be receiving messages pertaining to this procedure, even if it does not currently execute the code of this procedure.

Static Communicator Allocation

This covers the case where, at any point in time, at most one invocation of a parallel procedure can be active at any MPI process, and the group of executing MPI processes is fixed. For example, all invocations of parallel procedures involve all MPI processes, MPI processes are single-threaded, and there are no recursive invocations.

In such a case, a communicator can be statically allocated to each procedure. The static allocation can be done in a preamble, as part of initialization code. If the parallel procedures can be organized into libraries, so that only one procedure of each library can be concurrently active in each processor, then it is sufficient to allocate one communicator per library.

Dynamic Communicator Allocation

Calls of parallel procedures are well-nested if a new parallel procedure is always invoked in a subset of a group executing the same parallel procedure. Thus, MPI processes that execute the same parallel procedure have the same execution stack.

In such a case, a new communicator needs to be dynamically allocated for each new invocation of a parallel procedure. The allocation is done by the caller. A new communicator can be generated by a call to `MPI_COMM_DUP`, if the callee execution group is identical to the caller execution group, or by a call to `MPI_COMM_SPLIT` if the caller execution group is split into several subgroups executing distinct parallel routines. The new communicator is passed as an argument to the invoked routine.

The need for generating a new communicator at each invocation can be alleviated or avoided altogether in some cases: If the execution group is not split, then one can allocate a stack of communicators in a preamble, and next manage the stack in a way that mimics the stack of recursive calls.

One can also take advantage of the well-ordering property of communication to avoid confusing caller and callee communication, even if both use the same communicator. To do so, one needs to abide by the following two rules:

- messages sent before a procedure call (or before a return from the procedure) are also received before the matching call (or return) at the receiving end;
- messages are always selected by source (no use is made of `MPI_ANY_SOURCE`).

The General Case

In the general case, there may be multiple concurrently active invocations of the same parallel procedure within the same group; invocations may not be well-nested. A new communicator needs to be created for each invocation. It is the user's responsibility to make sure that, should two distinct parallel procedures be invoked concurrently on overlapping sets of MPI processes, communicator creation is properly coordinated.

DRAFT

Chapter 8

Virtual Topologies for MPI Processes

8.1 Introduction

This chapter discusses the MPI *virtual topology* mechanism. A *virtual topology* is an extra, optional attribute that one can give to an intra-communicator; *virtual topologies* cannot be added to inter-communicators. A *virtual topology* can provide a convenient naming mechanism for the MPI processes of a group (within a communicator), and additionally, may assist the runtime system in mapping the processes onto hardware.

As stated in Chapter 7, a group in MPI is an ordered set of n process identifiers (henceforth MPI processes). Each MPI process in the group is assigned a rank between 0 and $n-1$. In many parallel applications, a linear assignment of integer ranks to the MPI processes does not adequately reflect the logical communication pattern of the MPI processes (which is usually determined by the underlying problem geometry and the numerical algorithm used). Often the MPI processes are arranged in topological patterns such as two- or three-dimensional grids. More generally, the logical MPI process arrangement is described by a graph. In this chapter we will refer to this logical MPI process arrangement as the *virtual topology*.

A clear distinction must be made between the *virtual topology* and the topology of the underlying, physical hardware. The *virtual topology* can be exploited by the system in the assignment of processes to physical processors, if this helps to improve the communication performance on a given machine. How this mapping is done, however, is outside the scope of MPI. The description of the *virtual topology*, on the other hand, depends only on the application, and is machine-independent. The functions that are described in this chapter deal with machine-independent mapping and communication on *virtual topologies*.

Rationale. Though physical mapping is not discussed, the existence of the *virtual topology* information may be used as advice by the runtime system. There are well-known techniques for mapping grid/torus structures to hardware topologies such as hypercubes or grids. For more complicated graph structures good heuristics often yield nearly optimal results [51]. On the other hand, if there is no way for the user to specify the logical process arrangement as a *virtual topology*, a random mapping is most likely to result. On some machines, this will lead to unnecessary contention in the interconnection network. Some details about predicted and measured performance improvements that result from good process-to-processor mapping on wormhole-routing architectures can be found in [13, 14].

Besides possible performance benefits, the *virtual topology* can function as a convenient, process-naming structure, with significant benefits for program readability and notational power in message-passing programming. (*End of rationale.*)

8.2 Virtual Topologies

The communication pattern of a set of MPI processes can be represented by a graph. The nodes represent MPI processes, and the edges connect MPI processes that communicate with each other. MPI provides message-passing between any pair of MPI processes in a group. There is no requirement for opening a channel explicitly. Therefore, a “missing link” in the user-defined graph of MPI processes does not prevent the corresponding MPI processes from exchanging messages. It means rather that this connection is neglected in the *virtual topology*. This strategy implies that the *virtual topology* gives no convenient way of naming this pathway of communication. Another possible consequence is that an automatic mapping tool (if one exists for the runtime environment) will not take account of this edge when mapping.

Specifying the *virtual topology* in terms of a graph is sufficient for all applications. However, in many applications the graph structure is regular, and the detailed set-up of the graph would be inconvenient for the user and might be less efficient at run time. A large fraction of all parallel applications use MPI process topologies like rings, two- or higher-dimensional grids, or tori. These structures are completely defined by the number of dimensions and the numbers of MPI processes in each coordinate direction. Also, the mapping of grids and tori is generally an easier problem than that of general graphs. Thus, it is desirable to address these cases explicitly.

The coordinates of MPI processes in a Cartesian structure begin their numbering at 0. Row-major numbering is always used for the MPI processes in a Cartesian structure. This means that, for example, for four MPI processes in a (2×2) grid, the relationship between their ranks in the group and their coordinates in the *virtual topology* is as follows:

```
coord (0,0):  rank 0
coord (0,1):  rank 1
coord (1,0):  rank 2
coord (1,1):  rank 3
```

8.3 Embedding in MPI

The support for *virtual topologies* as defined in this chapter is consistent with other parts of MPI, and, whenever possible, makes use of functions that are defined elsewhere. Topology information is associated with communicators. It is added to communicators using the caching mechanism described in Chapter 7.

Information representing a *virtual topology* may be added to a communicator at the time of its creation. If a communicator creation function adds information representing a *virtual topology* to the output communicator it creates, then it either propagates the topology representation from the input communicator to the output communicator, or adds a new topology representation generated from the input parameters that describe a *virtual topology*. The description of every MPI communicator creation function explicitly states how topology information is handled. Communicator creation functions that create new topology representations are described in Section 8.5.

8.4 Overview of the Functions

MPI supports three types of *virtual topology*: **Cartesian**, **graph**, and **distributed graph**. The function `MPI_CART_CREATE` can be used to create Cartesian topologies, the function `MPI_GRAPH_CREATE` can be used to create graph topologies, and the functions `MPI_DIST_GRAPH_CREATE_ADJACENT` and `MPI_DIST_GRAPH_CREATE` can be used to create distributed graph topologies. These topology creation functions are collective. As with other collective calls, the program must be written to work correctly, whether the call synchronizes or not.

The above topology creation functions take as input an existing communicator `comm_old`, which defines the set of MPI processes on which the topology is to be mapped. For `MPI_GRAPH_CREATE` and `MPI_CART_CREATE`, all input arguments must have identical values on all MPI processes of the group of `comm_old`. When calling `MPI_GRAPH_CREATE`, each MPI process specifies all nodes and edges in the graph. In contrast, the functions `MPI_DIST_GRAPH_CREATE_ADJACENT` or `MPI_DIST_GRAPH_CREATE` are used to specify the graph in a distributed fashion, whereby each MPI process only specifies a subset of the edges in the graph such that the entire graph structure is defined collectively across the set of MPI processes. Therefore the MPI processes provide different values for the arguments specifying the graph. However, all MPI processes must give the same value for `reorder` and the `info` argument. In all cases, a new communicator `comm_topol` is created that carries the topological structure as cached information (see Chapter 7). In analogy to function `MPI_COMM_CREATE`, no cached information propagates from `comm_old` to `comm_topol`.

`MPI_CART_CREATE` can be used to describe Cartesian structures of arbitrary dimension. For each coordinate direction one specifies whether the MPI process structure is periodic or not. Note that an n -dimensional hypercube is an n -dimensional torus with two processes per coordinate direction. Thus, special support for hypercube structures is not necessary. The local auxiliary function `MPI_DIMS_CREATE` can be used to compute a balanced distribution of MPI processes among a given number of dimensions.

MPI defines functions to query a communicator for topology information. The function `MPI_TOPO_TEST` is used to query for the type of topology associated with a communicator. Depending on the topology type, different information can be extracted. For a graph topology, the functions `MPI_GRAPHDIMS_GET` and `MPI_GRAPH_GET` retrieve the graph topology information that is associated with the communicator. Additionally, the functions `MPI_GRAPH_NEIGHBORS_COUNT` and `MPI_GRAPH_NEIGHBORS` can be used to obtain the neighbors of an arbitrary node in the graph. For a distributed graph topology, the functions `MPI_DIST_GRAPH_NEIGHBORS_COUNT` and `MPI_DIST_GRAPH_NEIGHBORS` can be used to obtain the neighbors of the calling MPI process. For a Cartesian topology, the function `MPI_CARTDIM_GET` returns the number of dimensions and `MPI_CART_GET` returns the numbers of MPI processes in each dimension and periodicity of the associated Cartesian topology. Additionally, the functions `MPI_CART_RANK` and `MPI_CART_COORDS` translate Cartesian coordinates into a group rank, and vice-versa. The function `MPI_CART_SHIFT` provides the information needed to communicate with neighbors along a Cartesian dimension. All of these query functions are local.

For Cartesian topologies, the function `MPI_CART_SUB` can be used to extract a Cartesian subspace (analogous to `MPI_COMM_SPLIT`). This function is collective over the input communicator's group.

The two additional functions, `MPI_GRAPH_MAP` and `MPI_CART_MAP`, are, in gen-

eral, not called by the user directly. However, together with the communicator manipulation functions presented in Chapter 7, they are sufficient to implement all other topology functions. Section 8.5.8 outlines such an implementation.

The neighborhood collective communication routines `MPI_NEIGHBOR_ALLGATHER`, `MPI_NEIGHBOR_ALLGATHERV`, `MPI_NEIGHBOR_ALLTOALL`, `MPI_NEIGHBOR_ALLTOALLV`, and `MPI_NEIGHBOR_ALLTOALLW` communicate with the nearest neighbors on the topology associated with the communicator. The nonblocking variants are `MPI_INEIGHBOR_ALLGATHER`, `MPI_INEIGHBOR_ALLGATHERV`, `MPI_INEIGHBOR_ALLTOALL`, `MPI_INEIGHBOR_ALLTOALLV`, and `MPI_INEIGHBOR_ALLTOALLW`.

8.5 Topology Constructors

8.5.1 Cartesian Constructor

<code>MPI_CART_CREATE(comm_old, ndims, dims, periods, reorder, comm_cart)</code>			
IN	<code>comm_old</code>		input communicator (handle)
IN	<code>ndims</code>		number of dimensions of Cartesian grid (integer)
IN	<code>dims</code>		integer array of size <code>ndims</code> specifying the number of processes in each dimension
IN	<code>periods</code>		logical array of size <code>ndims</code> specifying whether the grid is periodic (<code>true</code>) or not (<code>false</code>) in each dimension
IN	<code>reorder</code>		ranks may be reordered (<code>true</code>) or not (<code>false</code>) (logical)
OUT	<code>comm_cart</code>		new communicator with associated Cartesian topology (handle)

C binding

```
int MPI_Cart_create(MPI_Comm comm_old, int ndims, const int dims[],
                   const int periods[], int reorder, MPI_Comm *comm_cart)
```

Fortran 2008 binding

```
MPI_Cart_create(comm_old, ndims, dims, periods, reorder, comm_cart, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm_old
  INTEGER, INTENT(IN) :: ndims, dims(ndims)
  LOGICAL, INTENT(IN) :: periods(ndims), reorder
  TYPE(MPI_Comm), INTENT(OUT) :: comm_cart
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_CART_CREATE(COMM_OLD, NDIMS, DIMS, PERIODS, REORDER, COMM_CART, IERROR)
  INTEGER COMM_OLD, NDIMS, DIMS(*), COMM_CART, IERROR
  LOGICAL PERIODS(*), REORDER
```

`MPI_CART_CREATE` returns a handle to a new communicator to which the Cartesian topology information is attached. If `reorder = false` then the rank of each MPI process in the group of the new communicator is identical to its rank in the group of the old communicator.

If `reorder = true` then the procedure may reorder the ranks of the MPI processes (possibly so as to choose a good embedding of the *virtual topology* onto the physical machine). If the total size of the Cartesian grid is smaller than the size of the group of `comm_old`, then some MPI processes return `MPI_COMM_NULL`, in analogy to `MPI_COMM_SPLIT`. If `ndims` is zero then a zero-dimensional Cartesian topology is created. The call is erroneous if it specifies a grid that is larger than the group size or if `ndims` is negative. `MPI_CART_CREATE` will associate information representing a Cartesian topology with the specified number of dimensions, numbers of MPI processes in each coordinate direction, and periodicity with the new communicator.

8.5.2 Cartesian Convenience Function: `MPI_DIMS_CREATE`

For Cartesian topologies, the function `MPI_DIMS_CREATE` helps the user select a balanced distribution of MPI processes per coordinate direction, depending on the number of MPI processes in the group to be balanced and optional constraints that can be specified by the user.

`MPI_DIMS_CREATE`(`nnodes`, `ndims`, `dims`)

IN	<code>nnodes</code>	number of nodes in a grid (integer)
IN	<code>ndims</code>	number of Cartesian dimensions (integer)
INOUT	<code>dims</code>	integer array of size <code>ndims</code> specifying the number of nodes in each dimension

C binding

```
int MPI_Dims_create(int nnodes, int ndims, int dims[])
```

Fortran 2008 binding

```
MPI_Dims_create(nnodes, ndims, dims, ierror)
  INTEGER, INTENT(IN) :: nnodes, ndims
  INTEGER, INTENT(INOUT) :: dims(ndims)
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_DIMS_CREATE(NNODES, NDIMS, DIMS, IERROR)
  INTEGER NNODES, NDIMS, DIMS(*), IERROR
```

The entries in the array `dims` are set to describe a Cartesian grid with `ndims` dimensions and a total of `nnodes` nodes. The dimensions are set to be as close to each other as possible, using an appropriate divisibility algorithm. The caller may further constrain the operation of this routine by specifying elements of array `dims`. If `dims[i]` is set to a positive number, the routine will not modify the number of nodes in dimension `i`; only those entries where `dims[i] = 0` are modified by the call.

Negative input values of `dims[i]` are erroneous. An error will occur if `nnodes` is not a multiple of

$$\prod_{i, \text{dims}[i] \neq 0} \text{dims}[i].$$

For `dims[i]` set by the call, `dims[i]` will be ordered in nonincreasing order. Array `dims` is suitable for use as input to routine `MPI_CART_CREATE`. `MPI_DIMS_CREATE` is local. If `ndims` is zero and `nnodes` is one, `MPI_DIMS_CREATE` returns `MPI_SUCCESS`.

Example 8.1. The use of the array argument `dims` in `MPI_DIMS_CREATE`.

dims before call	function call	dims on return
(0,0)	<code>MPI_DIMS_CREATE(6, 2, dims)</code>	(3,2)
(0,0)	<code>MPI_DIMS_CREATE(7, 2, dims)</code>	(7,1)
(0,3,0)	<code>MPI_DIMS_CREATE(6, 3, dims)</code>	(2,3,1)
(0,3,0)	<code>MPI_DIMS_CREATE(7, 3, dims)</code>	erroneous call

8.5.3 Graph Constructor

`MPI_GRAPH_CREATE(comm_old, nnodes, index, edges, reorder, comm_graph)`

IN	<code>comm_old</code>	input communicator (handle)
IN	<code>nnodes</code>	number of nodes in graph (integer)
IN	<code>index</code>	array of integers describing node degrees (see below)
IN	<code>edges</code>	array of integers describing graph edges (see below)
IN	<code>reorder</code>	ranks may be reordered (<code>true</code>) or not (<code>false</code>) (logical)
OUT	<code>comm_graph</code>	new communicator with associated graph topology (handle)

C binding

```
int MPI_Graph_create(MPI_Comm comm_old, int nnodes, const int index[],
                    const int edges[], int reorder, MPI_Comm *comm_graph)
```

Fortran 2008 binding

```
MPI_Graph_create(comm_old, nnodes, index, edges, reorder, comm_graph, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm_old
  INTEGER, INTENT(IN) :: nnodes, index(nnodes), edges(*)
  LOGICAL, INTENT(IN) :: reorder
  TYPE(MPI_Comm), INTENT(OUT) :: comm_graph
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GRAPH_CREATE(COMM_OLD, NNODES, INDEX, EDGES, REORDER, COMM_GRAPH, IERROR)
  INTEGER COMM_OLD, NNODES, INDEX(*), EDGES(*), COMM_GRAPH, IERROR
  LOGICAL REORDER
```

`MPI_GRAPH_CREATE` returns a handle to a new communicator to which the graph topology information is attached. If `reorder = false` then the rank of each MPI process in the group of the new communicator is identical to its rank in the group of the old communicator. If `reorder = true` then the procedure may reorder the ranks of the MPI processes. If the number of nodes in the graph (`nnodes`) is smaller than the size of the

group of `comm_old`, then `MPI_COMM_NULL` is returned by some MPI processes, in analogy to `MPI_CART_CREATE` and `MPI_COMM_SPLIT`. If the graph is empty, i.e., `nnodes = 0`, then `MPI_COMM_NULL` is returned in all MPI processes. The call is erroneous if it specifies a graph that is larger than the group size of the input communicator.

The three parameters `nnodes`, `index` and `edges` define the graph structure. `nnodes` is the number of nodes of the graph. The nodes are numbered from 0 to `nnodes-1`. The i -th entry of array `index` stores the total number of neighbors of the first i graph nodes. The lists of neighbors of nodes 0, 1, ..., `nnodes-1` are stored in consecutive locations in array `edges`. The array `edges` is a flattened representation of the edge lists. The total number of entries in `index` is `nnodes` and the total number of entries in `edges` is equal to the number of graph edges.

The definitions of the arguments `nnodes`, `index`, and `edges` are illustrated with the following simple example.

Example 8.2. Specification of the adjacency matrix for `MPI_GRAPH_CREATE`. Assume there are four MPI processes with ranks 0, 1, 2, 3 in the input communicator with the following adjacency matrix:

MPI process	neighbors
0	1, 3
1	0
2	3
3	0, 2

Then, the input arguments are:

```
nnodes = 4
index = 2, 3, 4, 6
edges = 1, 3, 0, 3, 0, 2
```

Thus, in C, `index[0]` is the degree of node zero, and `index[i] - index[i-1]` is the degree of node i , $i=1, \dots, nnodes-1$; the list of neighbors of node zero is stored in `edges[j]`, for $0 \leq j \leq index[0] - 1$ and the list of neighbors of node i , $i > 0$, is stored in `edges[j]`, $index[i-1] \leq j \leq index[i] - 1$.

In Fortran, `index(1)` is the degree of node zero, and `index(i+1) - index(i)` is the degree of node i , $i=1, \dots, nnodes-1$; the list of neighbors of node zero is stored in `edges(j)`, for $1 \leq j \leq index(1)$ and the list of neighbors of node i , $i > 0$, is stored in `edges(j)`, $index(i)+1 \leq j \leq index(i+1)$.

A single MPI process is allowed to be defined multiple times in the list of neighbors of an MPI process (i.e., there may be multiple edges between two MPI processes). An MPI process is also allowed to be a neighbor to itself (i.e., a self loop in the graph). The adjacency matrix is allowed to be nonsymmetric.

Advice to users. Performance implications of using multiple edges or a nonsymmetric adjacency matrix are not defined. The definition of a node-neighbor edge does not imply a direction of the communication. (*End of advice to users.*)

Advice to implementors. The following topology information is likely to be stored with a communicator:

- 1 • Type of topology (Cartesian/graph)
- 2
- 3 • For a Cartesian topology:
 - 4 1. ndims (number of dimensions)
 - 5 2. dims (numbers of MPI processes per coordinate direction)
 - 6 3. periods (periodicity information)
 - 7
 - 8 4. own_position (own position in grid, could also be computed from rank and
 - 9 dims)
- 10 • For a graph topology:
 - 11 1. index
 - 12 2. edges
 - 13
 - 14 which are the arrays defining the graph structure.

15 For a graph structure the number of nodes is equal to the number of MPI processes
 16 in the group. Therefore, the number of nodes does not have to be stored explicitly.
 17 An additional zero entry at the start of array `index` simplifies access to the topology
 18 information. (*End of advice to implementors.*)

20 8.5.4 Distributed Graph Constructor

21
 22 MPI_GRAPH_CREATE requires that each MPI process passes the full (global) communica-
 23 tion graph to the call. This limits the scalability of this constructor. With the distributed
 24 graph interface, the communication graph is specified in a fully distributed fashion. Each
 25 MPI process specifies only the part of the communication graph of which it is aware. Typ-
 26 ically, this could be the set of MPI processes from which the MPI process will eventually
 27 receive or get data, or the set of MPI processes to which the MPI process will send or put
 28 data, or some combination of such edges. Two different interfaces can be used to create a
 29 distributed graph topology. MPI_DIST_GRAPH_CREATE_ADJACENT creates a distributed
 30 graph communicator with each MPI process specifying each of its incoming and outgoing
 31 (adjacent) edges in the logical communication graph and thus requires minimal communi-
 32 cation during creation. MPI_DIST_GRAPH_CREATE provides full flexibility such that any
 33 MPI process can indicate that communication will occur between any pair of MPI processes
 34 in the graph.

35 To provide better possibilities for optimization by the MPI library, the distributed
 36 graph constructors permit weighted communication edges and take an `info` argument that
 37 can further influence process reordering or other optimizations performed by the MPI library.
 38 For example, hints can be provided on how edge weights are to be interpreted, the quality
 39 of the reordering, and/or the time permitted for the MPI library to process the graph.

40
 41
 42 MPI_DIST_GRAPH_CREATE_ADJACENT(comm_old, indegree, sources, sourceweights,
 43 outdegree, destinations, destweights, info, reorder, comm_dist_graph)

44 IN comm_old input communicator (handle)

45 IN indegree size of sources and sourceweights arrays (non-negative

46 integer)

47

48

IN	sources	ranks of MPI processes for which the calling process is a destination (array of non-negative integers)	1 2
IN	sourceweights	weights of the edges into the calling MPI process (array of non-negative integers)	3 4 5
IN	outdegree	size of destinations and destweights arrays (non-negative integer)	6 7
IN	destinations	ranks of MPI processes for which the calling MPI process is a source (array of non-negative integers)	8 9
IN	destweights	weights of the edges out of the calling MPI process (array of non-negative integers)	10 11 12
IN	info	hints on optimization and interpretation of weights (handle)	13 14
IN	reorder	ranks may be reordered (<code>true</code>) or not (<code>false</code>) (logical)	15 16
OUT	comm_dist_graph	new communicator with associated distributed graph topology (handle)	17 18

C binding

```
int MPI_Dist_graph_create_adjacent(MPI_Comm comm_old, int indegree,
    const int sources[], const int sourceweights[], int outdegree,
    const int destinations[], const int destweights[], MPI_Info info,
    int reorder, MPI_Comm *comm_dist_graph)
```

Fortran 2008 binding

```
MPI_Dist_graph_create_adjacent(comm_old, indegree, sources, sourceweights,
    outdegree, destinations, destweights, info, reorder,
    comm_dist_graph, ierror)
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm_old
```

```
INTEGER, INTENT(IN) :: indegree, sources(indegree), sourceweights(*),
    outdegree, destinations(outdegree), destweights(*)
```

```
TYPE(MPI_Info), INTENT(IN) :: info
```

```
LOGICAL, INTENT(IN) :: reorder
```

```
TYPE(MPI_Comm), INTENT(OUT) :: comm_dist_graph
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_DIST_GRAPH_CREATE_ADJACENT(COMM_OLD, INDEGREE, SOURCES, SOURCEWEIGHTS,
    OUTDEGREE, DESTINATIONS, DESTWEIGHTS, INFO, REORDER,
    COMM_DIST_GRAPH, IERROR)
```

```
INTEGER COMM_OLD, INDEGREE, SOURCES(*), SOURCEWEIGHTS(*), OUTDEGREE,
```

```
DESTINATIONS(*), DESTWEIGHTS(*), INFO, COMM_DIST_GRAPH, IERROR
```

```
LOGICAL REORDER
```

`MPI_DIST_GRAPH_CREATE_ADJACENT` returns a handle to a new communicator to which the distributed graph topology information is attached. Each MPI process passes all information about its incoming and outgoing edges in the virtual distributed graph topology. The calling MPI processes must ensure that each edge of the graph is described in the source and in the destination process with the same weights. If there are multiple

edges for a given (source,dest) pair, then the sequence of the weights of these edges does not matter. The complete communication topology is the combination of all edges shown in the sources arrays of all MPI processes in comm_old, which must be identical to the combination of all edges shown in the destinations arrays. Source and destination MPI processes must be specified by their rank in the group of comm_old. This allows a fully distributed specification of the communication graph. Isolated MPI processes (i.e., MPI processes with no outgoing or incoming edges, that is, MPI processes that have specified indegree and outdegree as zero and thus do not occur as source or destination in the graph specification) are allowed.

The call creates a new communicator comm_dist_graph of distributed graph topology type to which topology information has been attached. The number of MPI processes in comm_dist_graph is identical to the number of MPI processes in comm_old. The call to MPI_DIST_GRAPH_CREATE_ADJACENT is collective.

Weights are specified as nonnegative integers and can be used to influence the process mapping strategy and other internal MPI optimizations. For instance, approximate count arguments of later communication calls along specific edges could be used as their edge weights. Multiplicity of edges can likewise indicate more intense communication between pairs of MPI processes. However, the exact meaning of edge weights is not specified by the MPI standard and is left to the implementation. In C or Fortran, an application can supply the special value MPI_UNWEIGHTED for the weight array to indicate that all edges have the same (effectively no) weight. It is erroneous to supply MPI_UNWEIGHTED for some but not all MPI processes of comm_old. If the graph is weighted but indegree or outdegree is zero, then MPI_WEIGHTS_EMPTY or any arbitrary array may be passed to sourceweights or destweights respectively. Note that MPI_UNWEIGHTED and MPI_WEIGHTS_EMPTY are not special weight values; rather they are special values for the total array argument. In Fortran, MPI_UNWEIGHTED and MPI_WEIGHTS_EMPTY are objects like MPI_BOTTOM (not usable for initialization or assignment). See Section 2.5.4.

Advice to users. In the case of an empty weights array argument passed while constructing a weighted graph, one should not pass NULL because the value of MPI_UNWEIGHTED may be equal to NULL. The value of this argument would then be indistinguishable from MPI_UNWEIGHTED to the implementation. In this case MPI_WEIGHTS_EMPTY should be used instead. (*End of advice to users.*)

Advice to implementors. It is recommended that MPI_UNWEIGHTED not be implemented as NULL. (*End of advice to implementors.*)

Rationale. To ensure backward compatibility, MPI_UNWEIGHTED may still be implemented as NULL. See Annex B.4. (*End of rationale.*)

The meaning of the info and reorder arguments is defined in the description of the following routine.

```
MPI_DIST_GRAPH_CREATE(comm_old, n, sources, degrees, destinations, weights, info,
                      reorder, comm_dist_graph)
```

IN	comm_old	input communicator (handle)
IN	n	number of source nodes for which this MPI process specifies edges (non-negative integer)

IN	sources	array containing the n source nodes for which this MPI process specifies edges (array of non-negative integers)	1 2 3
IN	degrees	array specifying the number of destinations for each source node in the source node array (array of non-negative integers)	4 5 6
IN	destinations	destination nodes for the source nodes in the source node array (array of non-negative integers)	7 8 9
IN	weights	weights for source to destination edges (array of non-negative integers)	10 11
IN	info	hints on optimization and interpretation of weights (handle)	12 13 14
IN	reorder	ranks may be reordered (<code>true</code>) or not (<code>false</code>) (logical)	15
OUT	comm_dist_graph	new communicator with associated distributed graph topology (handle)	16 17 18

C binding

```
int MPI_Dist_graph_create(MPI_Comm comm_old, int n, const int sources[],
                        const int degrees[], const int destinations[],
                        const int weights[], MPI_Info info, int reorder,
                        MPI_Comm *comm_dist_graph)
```

Fortran 2008 binding

```
MPI_Dist_graph_create(comm_old, n, sources, degrees, destinations, weights,
                    info, reorder, comm_dist_graph, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm_old
INTEGER, INTENT(IN) :: n, sources(n), degrees(n), destinations(*),
                    weights(*)
TYPE(MPI_Info), INTENT(IN) :: info
LOGICAL, INTENT(IN) :: reorder
TYPE(MPI_Comm), INTENT(OUT) :: comm_dist_graph
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_DIST_GRAPH_CREATE(COMM_OLD, N, SOURCES, DEGREES, DESTINATIONS, WEIGHTS,
                    INFO, REORDER, COMM_DIST_GRAPH, IERROR)
INTEGER COMM_OLD, N, SOURCES(*), DEGREES(*), DESTINATIONS(*), WEIGHTS(*),
                    INFO, COMM_DIST_GRAPH, IERROR
LOGICAL REORDER
```

MPI_DIST_GRAPH_CREATE returns a handle to a new communicator to which the distributed graph topology information is attached. Concretely, each MPI process calls the constructor with a set of directed (`source,destination`) communication edges as described below. Every MPI process passes an array of n source nodes in the `sources` array. For each source node, a nonnegative number of destination nodes is specified in the `degrees` array. The destination nodes are stored in the corresponding consecutive segment of the `destinations` array. More precisely, if the i-th node in `sources` is `s`, this specifies `degrees[i]` edges

1 (s,d) with d of the j-th such edge stored in destinations[degrees[0]+...+degrees[i-1]+j]. The
2 weight of this edge is stored in weights[degrees[0]+...+degrees[i-1]+j]. Both the sources and
3 the destinations arrays may contain the same node more than once, and the order in which
4 nodes are listed as destinations or sources is not significant. Similarly, different processes
5 may specify edges with the same source and destination nodes. Source and destination
6 nodes must be specified by their rank in the group of comm_old. Different MPI processes
7 may specify different numbers of source and destination nodes, as well as different source to
8 destination edges. This allows a fully distributed specification of the communication graph.
9 Isolated MPI processes (i.e., MPI processes with no outgoing or incoming edges, that is, MPI
10 processes that do not occur as source or destination node in the graph specification) are
11 allowed.

12 The call creates a new communicator comm_dist_graph of distributed graph topology
13 type to which topology information has been attached. The number of MPI processes in
14 comm_dist_graph is identical to the number of MPI processes in comm_old. The call to
15 MPI_DIST_GRAPH_CREATE is collective.

16 If reorder = false, all MPI processes will have the same rank in comm_dist_graph as in
17 comm_old. If reorder = true then the MPI library is free to remap to other MPI processes (of
18 comm_old) in order to improve communication on the edges of the communication graph.
19 The weight associated with each edge is a hint to the MPI library about the amount or
20 intensity of communication on that edge, and may be used to compute a “best” reordering.

21 Weights are specified as nonnegative integers and can be used to influence the MPI
22 process remapping strategy and other internal MPI optimizations. For instance, approxi-
23 mate count arguments of later communication calls along specific edges could be used as
24 their edge weights. Multiplicity of edges can likewise indicate more intense communication
25 between pairs of MPI processes. However, the exact meaning of edge weights and mul-
26 tiplicity of edges is not specified by the MPI standard and is left to the implementation.
27 In C or Fortran, an application can supply the special value MPI_UNWEIGHTED for the
28 weight array to indicate that all edges have the same (effectively no) weight. It is erro-
29 neous to supply MPI_UNWEIGHTED for some but not all MPI processes of comm_old. If the
30 graph is weighted but n = 0, then MPI_WEIGHTS_EMPTY or any arbitrary array may be
31 passed to weights. Note that MPI_UNWEIGHTED and MPI_WEIGHTS_EMPTY are not spe-
32 cial weight values; rather they are special values for the total array argument. In Fortran,
33 MPI_UNWEIGHTED and MPI_WEIGHTS_EMPTY are objects like MPI_BOTTOM (not usable
34 for initialization or assignment). See Section 2.5.4.

35
36 *Advice to users.* In the case of an empty weights array argument passed while
37 constructing a weighted graph, one should not pass NULL because the value of
38 MPI_UNWEIGHTED may be equal to NULL. The value of this argument would then
39 be indistinguishable from MPI_UNWEIGHTED to the implementation.
40 MPI_WEIGHTS_EMPTY should be used instead. (*End of advice to users.*)

41
42 *Advice to implementors.* It is recommended that MPI_UNWEIGHTED not be imple-
43 mented as NULL. (*End of advice to implementors.*)

44
45 *Rationale.* To ensure backward compatibility, MPI_UNWEIGHTED may still be imple-
46 mented as NULL. See Annex B.4. (*End of rationale.*)

47 The meaning of the weights argument can be influenced by the
48 info argument. The info argument can be used to guide the mapping of MPI processes to

the hardware; possible options include minimizing the maximum number of edges between processes on different SMP nodes, or minimizing the sum of all such edges. As described in Section 10, an MPI implementation is not obliged to follow specific hints, and it is valid for an MPI implementation not to do any reordering. An MPI implementation may specify more info (key,value) pairs. All MPI processes must specify the same set of (key,value) info pairs.

Advice to implementors. MPI implementations must document any additionally supported (key,value) info pairs. MPI_INFO_NULL is always valid, and may indicate the default creation of the distributed graph topology to the MPI library.

An implementation does not explicitly need to construct the topology from its distributed parts. However, all MPI processes can construct the full topology from the distributed specification and use this in a call to MPI_GRAPH_CREATE to create the topology. This may serve as a reference implementation of the functionality, and may be acceptable for small communicators. However, a scalable high-quality implementation would save the topology graph in a distributed way. (*End of advice to implementors.*)

Example 8.3. Several ways to specify the adjacency matrix for MPI_DIST_GRAPH_CREATE and MPI_DIST_GRAPH_CREATE_ADJACENT. As for Example 8.2, assume there are four MPI processes with ranks 0, 1, 2, 3 in the input communicator with the following adjacency matrix and unit edge weights:

MPI process	neighbors
0	1, 3
1	0
2	3
3	0, 2

With MPI_DIST_GRAPH_CREATE, this graph could be constructed in many different ways. One way would be that each MPI process specifies its outgoing edges. The arguments per MPI process would be:

MPI process	n	sources	degrees	destinations	weights
0	1	0	2	1,3	1,1
1	1	1	1	0	1
2	1	2	1	3	1
3	1	3	2	0,2	1,1

Another way would be to pass the whole graph on MPI process with rank 0 in the input communicator, which could be done with the following arguments per MPI process:

MPI process	n	sources	degrees	destinations	weights
0	4	0,1,2,3	2,1,1,2	1,3,0,3,0,2	1,1,1,1,1,1
1	0	-	-	-	-
2	0	-	-	-	-
3	0	-	-	-	-

In both cases above, the application could supply `MPI_UNWEIGHTED` instead of explicitly providing identical weights.

`MPI_DIST_GRAPH_CREATE_ADJACENT` could be used to specify this graph using the following arguments:

MPI process	indegree	sources	sourceweights	outdegree	destinations	destweights
0	2	1,3	1,1	2	1,3	1,1
1	1	0	1	1	0	1
2	1	3	1	1	3	1
3	2	0,2	1,1	2	0,2	1,1

Example 8.4. Cartesian grid plus diagonals specified with `MPI_DIST_GRAPH_CREATE`. A two-dimensional $P \times Q$ torus where all MPI processes communicate along the dimensions and along the diagonal edges cannot be modeled with Cartesian topologies, but can easily be captured with `MPI_DIST_GRAPH_CREATE` as shown in the following code. In this example, the communication along the dimensions is twice as heavy as the communication along the diagonals:

```

19 /*
20 Input:      dimensions P, Q
21 Condition: number of MPI processes equal to P*Q
22 */
23 int rank, x, y;
24 int sources[1], degrees[1];
25 int destinations[8], weights[8];
26 MPI_Comm comm_dist_graph;
27
28 MPI_Comm_rank(MPI_COMM_WORLD, &rank);
29
30 /* get x and y dimension */
31 y=rank/P; x=rank%P;
32
33 /* get my communication partners along x dimension */
34 destinations[0] = P*y+(x+1)%P; weights[0] = 2;
35 destinations[1] = P*y+(P+x-1)%P; weights[1] = 2;
36
37 /* get my communication partners along y dimension */
38 destinations[2] = P*((y+1)%Q)+x; weights[2] = 2;
39 destinations[3] = P*((Q+y-1)%Q)+x; weights[3] = 2;
40
41 /* get my communication partners along diagonals */
42 destinations[4] = P*((y+1)%Q)+(x+1)%P; weights[4] = 1;
43 destinations[5] = P*((Q+y-1)%Q)+(x+1)%P; weights[5] = 1;
44 destinations[6] = P*((y+1)%Q)+(P+x-1)%P; weights[6] = 1;
45 destinations[7] = P*((Q+y-1)%Q)+(P+x-1)%P; weights[7] = 1;
46
47 sources[0] = rank;
48 degrees[0] = 8;
49 MPI_Dist_graph_create(MPI_COMM_WORLD, 1, sources, degrees, destinations,
50                      weights, MPI_INFO_NULL, 1, &comm_dist_graph);

```

8.5.5 Topology Inquiry Functions

If a *virtual topology* has been defined with one of the above functions, then the topology information can be looked up using inquiry functions. They all are local calls.

`MPI_TOPO_TEST(comm, status)`

IN	comm	communicator (handle)
OUT	status	topology type of communicator comm (state)

C binding

```
int MPI_Topo_test(MPI_Comm comm, int *status)
```

Fortran 2008 binding

```
MPI_Topo_test(comm, status, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(OUT) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TOPO_TEST(COMM, STATUS, IERROR)
  INTEGER COMM, STATUS, IERROR
```

The function `MPI_TOPO_TEST` returns the type of topology that is associated with a communicator.

The output value `status` is one of the following:

<code>MPI_GRAPH</code>	graph topology
<code>MPI_CART</code>	Cartesian topology
<code>MPI_DIST_GRAPH</code>	distributed graph topology
<code>MPI_UNDEFINED</code>	no topology

`MPI_GRAPHDIMS_GET(comm, nnodes, nedges)`

IN	comm	communicator with associated graph topology (handle)
OUT	nnodes	number of nodes in graph (same as number of MPI processes in the group of comm) (integer)
OUT	nedges	number of edges in graph (integer)

C binding

```
int MPI_Graphdims_get(MPI_Comm comm, int *nnodes, int *nedges)
```

Fortran 2008 binding

```
MPI_Graphdims_get(comm, nnodes, nedges, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(OUT) :: nnodes, nedges
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GRAPHDIMS_GET(COMM, NNODES, NEDGES, IERROR)
```

1 INTEGER COMM, NNODES, NEDGES, IERROR

2 The functions `MPI_GRAPHDIMS_GET` and `MPI_GRAPH_GET` retrieve the graph topology information that is associated with the communicator. The information provided by `MPI_GRAPHDIMS_GET` can be used to dimension the vectors `index` and `edges` correctly for the following call to `MPI_GRAPH_GET`.

3
4
5
6
7
8 **MPI_GRAPH_GET**(comm, maxindex, maxedges, index, edges)
9
10 IN comm communicator with associated graph topology
11 (handle)
12 IN maxindex length of vector `index` in the calling program (integer)
13 IN maxedges length of vector `edges` in the calling program (integer)
14 OUT index array of integers containing the graph structure (for
15 details see the definition of `MPI_GRAPH_CREATE`)
16
17 OUT edges array of integers containing the graph structure

18
19 **C binding**

20 int MPI_Graph_get(MPI_Comm comm, int maxindex, int maxedges, int index[],
21 int edges[])

22
23 **Fortran 2008 binding**

24 MPI_Graph_get(comm, maxindex, maxedges, index, edges, ierror)
25 TYPE(MPI_Comm), INTENT(IN) :: comm
26 INTEGER, INTENT(IN) :: maxindex, maxedges
27 INTEGER, INTENT(OUT) :: index(maxindex), edges(maxedges)
28 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

29 **Fortran binding**

30 MPI_GRAPH_GET(COMM, MAXINDEX, MAXEDGES, INDEX, EDGES, IERROR)
31 INTEGER COMM, MAXINDEX, MAXEDGES, INDEX(*), EDGES(*), IERROR

32
33
34 **MPI_CARTDIM_GET**(comm, ndims)

35 IN comm communicator with associated Cartesian topology
36 (handle)
37
38 OUT ndims number of dimensions of the Cartesian structure
39 (integer)

40
41 **C binding**

42 int MPI_Cartdim_get(MPI_Comm comm, int *ndims)

43 **Fortran 2008 binding**

44 MPI_Cartdim_get(comm, ndims, ierror)
45 TYPE(MPI_Comm), INTENT(IN) :: comm
46 INTEGER, INTENT(OUT) :: ndims
47 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48

Fortran binding

```
MPI_CARTDIM_GET(COMM, NDIMS, IERROR)
    INTEGER COMM, NDIMS, IERROR
```

The functions `MPI_CARTDIM_GET` and `MPI_CART_GET` return the Cartesian topology information that is associated with the communicator. If `comm` is associated with a zero-dimensional Cartesian topology, `MPI_CARTDIM_GET` returns `ndims = 0` and `MPI_CART_GET` will keep all output arguments unchanged.

```
MPI_CART_GET(comm, maxdims, dims, periods, coords)
```

IN	<code>comm</code>	communicator with associated Cartesian topology (handle)
IN	<code>maxdims</code>	length of vectors <code>dims</code> , <code>periods</code> , and <code>coords</code> in the calling program (integer)
OUT	<code>dims</code>	number of MPI processes for each Cartesian dimension (array of integers)
OUT	<code>periods</code>	periodicity (<code>true/false</code>) for each Cartesian dimension (array of logicals)
OUT	<code>coords</code>	coordinates of calling MPI process in Cartesian structure (array of integers)

C binding

```
int MPI_Cart_get(MPI_Comm comm, int maxdims, int dims[], int periods[],
    int coords[])
```

Fortran 2008 binding

```
MPI_Cart_get(comm, maxdims, dims, periods, coords, ierror)
    TYPE(MPI_Comm), INTENT(IN) :: comm
    INTEGER, INTENT(IN) :: maxdims
    INTEGER, INTENT(OUT) :: dims(maxdims), coords(maxdims)
    LOGICAL, INTENT(OUT) :: periods(maxdims)
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_CART_GET(COMM, MAXDIMS, DIMS, PERIODS, COORDS, IERROR)
    INTEGER COMM, MAXDIMS, DIMS(*), COORDS(*), IERROR
    LOGICAL PERIODS(*)
```

If `maxdims` in a call to `MPI_CART_GET` is less than the number of dimensions of the Cartesian topology associated with the communicator `comm`, the outcome is unspecified.

```
MPI_CART_RANK(comm, coords, rank)
```

IN	<code>comm</code>	communicator with associated Cartesian topology (handle)
IN	<code>coords</code>	integer array (of size <code>ndims</code>) specifying the Cartesian coordinates of an MPI process

```

1      OUT      rank                rank of specified MPI process within group of comm
2                                     (integer)
3

```

C binding

```

5      int MPI_Cart_rank(MPI_Comm comm, const int coords[], int *rank)
6

```

Fortran 2008 binding

```

8      MPI_Cart_rank(comm, coords, rank, ierror)
9          TYPE(MPI_Comm), INTENT(IN) :: comm
10         INTEGER, INTENT(IN) :: coords(*)
11         INTEGER, INTENT(OUT) :: rank
12         INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

14     MPI_CART_RANK(COMM, COORDS, RANK, IERROR)
15         INTEGER COMM, COORDS(*), RANK, IERROR

```

For a communicator with an associated Cartesian topology, the function `MPI_CART_RANK` translates the logical coordinates of an MPI process to the corresponding rank in the group of the communicator. For dimension i with `periods(i) = true`, if the coordinate, `coords(i)`, is out of range, that is, `coords(i) < 0` or `coords(i) ≥ dims(i)`, it is shifted back to the interval $0 ≤ \text{coords}(i) < \text{dims}(i)$ automatically. Out-of-range coordinates are erroneous for nonperiodic dimensions.

If `comm` is associated with a zero-dimensional Cartesian topology, `coords` is not significant and 0 is returned in `rank`.

```

26     MPI_CART_COORDS(comm, rank, maxdims, coords)

```

28	IN	comm	communicator with associated Cartesian topology (handle)
30	IN	rank	rank of an MPI process within group of comm (integer)
32	IN	maxdims	length of vector <code>coords</code> in the calling program (integer)
34	OUT	coords	coordinates of the MPI process with the rank <code>rank</code> in Cartesian structure (array of integers)

C binding

```

39     int MPI_Cart_coords(MPI_Comm comm, int rank, int maxdims, int coords[])

```

Fortran 2008 binding

```

41     MPI_Cart_coords(comm, rank, maxdims, coords, ierror)
42         TYPE(MPI_Comm), INTENT(IN) :: comm
43         INTEGER, INTENT(IN) :: rank, maxdims
44         INTEGER, INTENT(OUT) :: coords(maxdims)
45         INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

47     MPI_CART_COORDS(COMM, RANK, MAXDIMS, COORDS, IERROR)

```

INTEGER COMM, RANK, MAXDIMS, COORDS(*), IERROR

The inverse mapping, rank-to-coordinates translation is provided by MPI_CART_COORDS. If comm is associated with a zero-dimensional Cartesian topology, coords will be unchanged. If maxdims is less than the number of dimensions of the Cartesian topology associated with the communicator comm, the outcome is unspecified.

MPI_GRAPH_NEIGHBORS_COUNT(comm, rank, nneighbors)

IN	comm	communicator with associated graph topology (handle)
IN	rank	rank of MPI process in group of comm (integer)
OUT	nneighbors	number of neighbors of specified MPI process (integer)

C binding

int MPI_Graph_neighbors_count(MPI_Comm comm, int rank, int *nneighbors)

Fortran 2008 binding

```
MPI_Graph_neighbors_count(comm, rank, nneighbors, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: rank
  INTEGER, INTENT(OUT) :: nneighbors
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GRAPH_NEIGHBORS_COUNT(COMM, RANK, NNEIGHBORS, IERROR)
  INTEGER COMM, RANK, NNEIGHBORS, IERROR
```

MPI_GRAPH_NEIGHBORS(comm, rank, maxneighbors, neighbors)

IN	comm	communicator with associated graph topology (handle)
IN	rank	rank of MPI process in group of comm (integer)
IN	maxneighbors	size of array neighbors (integer)
OUT	neighbors	ranks of MPI processes that are neighbors to specified MPI process (array of integers)

C binding

```
int MPI_Graph_neighbors(MPI_Comm comm, int rank, int maxneighbors,
  int neighbors[])
```

Fortran 2008 binding

```
MPI_Graph_neighbors(comm, rank, maxneighbors, neighbors, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: rank, maxneighbors
  INTEGER, INTENT(OUT) :: neighbors(maxneighbors)
```

1 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

2 **Fortran binding**

3 MPI_GRAPH_NEIGHBORS(COMM, RANK, MAXNEIGHBORS, NEIGHBORS, IERROR)

4 INTEGER COMM, RANK, MAXNEIGHBORS, NEIGHBORS(*), IERROR

5
6 MPI_GRAPH_NEIGHBORS_COUNT and MPI_GRAPH_NEIGHBORS provide adjacency
7 information for a graph topology. The returned count and array of neighbors for the queried
8 rank will both include *all* neighbors and reflect the same edge ordering as was specified by
9 the original call to MPI_GRAPH_CREATE. Specifically, MPI_GRAPH_NEIGHBORS_COUNT
10 and MPI_GRAPH_NEIGHBORS will return values based on the original index and edges
11 array passed to MPI_GRAPH_CREATE (for the purpose of this example, we assume that
12 `index[-1]` is zero):

- 13
- 14 • The number of neighbors (`nneighbors`) returned from
- 15 MPI_GRAPH_NEIGHBORS_COUNT will be $(\text{index}[\text{rank}] - \text{index}[\text{rank}-1])$.
- 16
- 17 • The neighbors array returned from MPI_GRAPH_NEIGHBORS will be
- 18 edges[`index[rank-1]`] through edges[`index[rank]-1`].
- 19

20 **Example 8.5.** Inquiry of graph topology information.

21 Assume there are four MPI processes with ranks 0, 1, 2, 3 in the input communicator with
22 the following adjacency matrix (note that some neighbors are listed multiple times):

MPI process	neighbors
0	1, 1, 3
1	0, 0
2	3
3	0, 2, 2

23
24
25
26
27
28
29 Thus, the input arguments to MPI_GRAPH_CREATE are:

30
31 nnodes = 4
32 index = 3, 5, 6, 9
33 edges = 1, 1, 3, 0, 0, 3, 0, 2, 2
34

35 Therefore, calling MPI_GRAPH_NEIGHBORS_COUNT and MPI_GRAPH_NEIGHBORS for
36 each of the four MPI processes will return:

Input rank	Count	Neighbors
0	3	1, 1, 3
1	2	0, 0
2	1	3
3	3	0, 2, 2

37
38
39
40
41
42
43 **Example 8.6.** Using a communicator with an associated graph topology that represents
44 a shuffle-exchange network.

45 Suppose that `comm` is a communicator with a shuffle-exchange topology. The group has 2^n
46 members. Each MPI process is labeled by a_1, \dots, a_n with $a_i \in \{0, 1\}$, and has three neigh-
47
48

bors: $\text{exchange}(a_1, \dots, a_n) = a_1, \dots, a_{n-1}, \bar{a}_n$ ($\bar{a} = 1 - a$), $\text{shuffle}(a_1, \dots, a_n) = a_2, \dots, a_n, a_1$, and $\text{unshuffle}(a_1, \dots, a_n) = a_n, a_1, \dots, a_{n-1}$. The graph adjacency list is illustrated below for $n = 3$.

node	exchange neighbors(1)	shuffle neighbors(2)	unshuffle neighbors(3)
0 (000)	1	0	0
1 (001)	0	2	4
2 (010)	3	4	1
3 (011)	2	6	5
4 (100)	5	1	2
5 (101)	4	3	6
6 (110)	7	5	3
7 (111)	6	7	7

Suppose that the communicator `comm` has this topology associated with it. The following code fragment cycles through the three types of neighbors and performs an appropriate permutation for each.

```

! assume: each MPI process has stored a real number A.
! extract neighborhood information
CALL MPI_COMM_RANK(comm, myrank, ierr)
CALL MPI_GRAPH_NEIGHBORS(comm, myrank, 3, neighbors, ierr)
! perform exchange permutation
CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, neighbors(1), 0, &
                        neighbors(1), 0, comm, status, ierr)
! perform shuffle permutation
CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, neighbors(2), 0, &
                        neighbors(3), 0, comm, status, ierr)
! perform unshuffle permutation
CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, neighbors(3), 0, &
                        neighbors(2), 0, comm, status, ierr)

```

`MPI_DIST_GRAPH_NEIGHBORS_COUNT` and `MPI_DIST_GRAPH_NEIGHBORS` provide adjacency information for a distributed graph topology.

`MPI_DIST_GRAPH_NEIGHBORS_COUNT(comm, indegree, outdegree, weighted)`

IN	comm	communicator with associated distributed graph topology (handle)
OUT	indegree	number of edges into this MPI process (non-negative integer)
OUT	outdegree	number of edges out of this MPI process (non-negative integer)
OUT	weighted	false if <code>MPI_UNWEIGHTED</code> was supplied during creation, true otherwise (logical)

C binding

```
int MPI_Dist_graph_neighbors_count(MPI_Comm comm, int *indegree,
```

```
1         int *outdegree, int *weighted)
```

```
2
3 Fortran 2008 binding
```

```
4 MPI_Dist_graph_neighbors_count(comm, indegree, outdegree, weighted, ierror)
5     TYPE(MPI_Comm), INTENT(IN) :: comm
6     INTEGER, INTENT(OUT) :: indegree, outdegree
7     LOGICAL, INTENT(OUT) :: weighted
8     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
9 Fortran binding
```

```
10 MPI_DIST_GRAPH_NEIGHBORS_COUNT(COMM, INDEGREE, OUTDEGREE, WEIGHTED, IERROR)
11     INTEGER COMM, INDEGREE, OUTDEGREE, IERROR
12     LOGICAL WEIGHTED
```

```
13
14
15 MPI_DIST_GRAPH_NEIGHBORS(comm, maxindegree, sources, sourceweights,
16     maxoutdegree, destinations, destweights)
```

```
17     IN     comm           communicator with associated distributed graph
18                       topology (handle)
19
20     IN     maxindegree    size of sources and sourceweights arrays
21                       (non-negative integer)
22
23     OUT    sources        ranks of MPI processes for which the calling MPI
24                       process is a destination (array of non-negative
25                       integers)
26
27     OUT    sourceweights  weights of the edges into the calling MPI process
28                       (array of non-negative integers)
29
30     IN     maxoutdegree   size of destinations and destweights arrays
31                       (non-negative integer)
32
33     OUT    destinations   ranks of MPI processes for which the calling MPI
34                       process is a source (array of non-negative integers)
35
36     OUT    destweights    weights of the edges out of the calling MPI process
37                       (array of non-negative integers)
```

```
35 C binding
```

```
36 int MPI_Dist_graph_neighbors(MPI_Comm comm, int maxindegree, int sources[],
37     int sourceweights[], int maxoutdegree, int destinations[],
38     int destweights[])
39
```

```
40 Fortran 2008 binding
```

```
41 MPI_Dist_graph_neighbors(comm, maxindegree, sources, sourceweights,
42     maxoutdegree, destinations, destweights, ierror)
43     TYPE(MPI_Comm), INTENT(IN) :: comm
44     INTEGER, INTENT(IN) :: maxindegree, maxoutdegree
45     INTEGER, INTENT(OUT) :: sources(maxindegree), destinations(maxoutdegree)
46     INTEGER :: sourceweights(*), destweights(*)
47     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48
```

Fortran binding

```

MPI_DIST_GRAPH_NEIGHBORS(COMM, MAXINDEGREE, SOURCES, SOURCEWEIGHTS,
                          MAXOUTDEGREE, DESTINATIONS, DESTWEIGHTS, IERROR)
INTEGER COMM, MAXINDEGREE, SOURCES(*), SOURCEWEIGHTS(*), MAXOUTDEGREE,
          DESTINATIONS(*), DESTWEIGHTS(*), IERROR

```

These calls are local. The number of edges into and out of the MPI process returned by `MPI_DIST_GRAPH_NEIGHBORS_COUNT` are the total number of such edges given in the call to `MPI_DIST_GRAPH_CREATE_ADJACENT` or `MPI_DIST_GRAPH_CREATE` (potentially by MPI processes other than the calling MPI process in the case of `MPI_DIST_GRAPH_CREATE`). Multiply-defined edges are all counted and returned by `MPI_DIST_GRAPH_NEIGHBORS` in some order. If `MPI_UNWEIGHTED` is supplied for `sourceweights` or `destweights` or both, or if `MPI_UNWEIGHTED` was supplied during the construction of the graph then no weight information is returned in that array or those arrays. If the communicator was created with `MPI_DIST_GRAPH_CREATE_ADJACENT` then for each MPI process in `comm`, the order of the values in `sources` and `destinations` is identical to the input that was used by the MPI process with the same rank in `comm_old` in the creation call. If the communicator was created with `MPI_DIST_GRAPH_CREATE` then the only requirement on the order of values in `sources` and `destinations` is that two calls to the routine with same input argument `comm` will return the same sequence of edges. If `maxindegree` or `maxoutdegree` is smaller than the numbers returned by `MPI_DIST_GRAPH_NEIGHBORS_COUNT`, then only the first part of the full list is returned.

Advice to implementors. Since the query calls are defined to be local, each MPI process needs to store the list of its neighbors with incoming and outgoing edges. Communication is required at the collective `MPI_DIST_GRAPH_CREATE` call in order to compute the neighbor lists for each MPI process from the distributed graph specification. (*End of advice to implementors.*)

8.5.6 Cartesian Shift Coordinates

If the MPI process topology is a Cartesian structure, an `MPI_SENDRECV` operation may be used along a coordinate direction to perform a shift of data. As input, `MPI_SENDRECV` takes the rank of a source MPI process for the receive, and the rank of a destination MPI process for the send. If the function `MPI_CART_SHIFT` is called for a communicator with an associated Cartesian topology, it provides the calling MPI process with the above identifiers, which then can be passed to `MPI_SENDRECV`. The user specifies the coordinate direction and the size of the step (positive or negative, but not zero). The function is local.

```

MPI_CART_SHIFT(comm, direction, disp, rank_source, rank_dest)

```

IN	<code>comm</code>	communicator with associated Cartesian topology (handle)
IN	<code>direction</code>	coordinate dimension of shift (integer)
IN	<code>disp</code>	displacement (> 0: upwards shift, < 0: downwards shift) (integer)
OUT	<code>rank_source</code>	rank of source MPI process (integer)
OUT	<code>rank_dest</code>	rank of destination MPI process (integer)

C binding

```

1 int MPI_Cart_shift(MPI_Comm comm, int direction, int disp, int *rank_source,
2                   int *rank_dest)
3
4

```

Fortran 2008 binding

```

5 MPI_Cart_shift(comm, direction, disp, rank_source, rank_dest, ierror)
6     TYPE(MPI_Comm), INTENT(IN) :: comm
7     INTEGER, INTENT(IN) :: direction, disp
8     INTEGER, INTENT(OUT) :: rank_source, rank_dest
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10

```

Fortran binding

```

11 MPI_CART_SHIFT(COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR)
12     INTEGER COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR
13

```

14 The direction argument indicates the coordinate dimension to be traversed by the shift.
15 The dimensions are numbered from 0 to `ndims-1`, where `ndims` is the number of dimensions.

16 Depending on the periodicity of the Cartesian topology in the specified coordinate
17 direction, `MPI_CART_SHIFT` provides the identifiers for a circular or an end-off shift. In the
18 case of an end-off shift, the value `MPI_PROC_NULL` is returned in `rank_source` or `rank_dest`,
19 indicating that the source or the destination for the shift is out of range.

20 It is erroneous to call `MPI_CART_SHIFT` with a direction that is either negative or
21 greater than or equal to the number of dimensions in the Cartesian communicator. This
22 implies that it is erroneous to call `MPI_CART_SHIFT` with a `comm` that is associated with
23 a zero-dimensional Cartesian topology.
24

Example 8.7. Using `MPI_CART_SHIFT` for a Cartesian topology.

25 The communicator, `comm`, has a two-dimensional, periodic, Cartesian topology associated
26 with it. A two-dimensional array of REALs is stored one element per MPI process, in variable
27 `A`. One wishes to skew this array, by shifting column `i` (vertically, i.e., along the column) by
28 `i` steps.
29

```

30 ...
31 ! find MPI process rank
32 CALL MPI_COMM_RANK(comm, rank, ierr)
33 ! find Cartesian coordinates
34 CALL MPI_CART_COORDS(comm, rank, maxdims, coords, ierr)
35 ! compute shift source and destination
36 CALL MPI_CART_SHIFT(comm, 0, coords(2), source, dest, ierr)
37 ! skew array
38 CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, dest, 0, source, 0, comm, &
39                          status, ierr)
40

```

41 *Advice to users.* In Fortran, the dimension indicated by `DIRECTION = i` has
42 `DIMS(i+1)` nodes, where `DIMS` is the array that was used to create the grid. In C, the
43 dimension indicated by `direction = i` is the dimension specified by `dims[i]`. (*End*
44 *of advice to users.*)
45
46
47
48

8.5.7 Partitioning of Cartesian Structures

`MPI_CART_SUB(comm, remain_dims, newcomm)`

IN	comm	communicator with associated Cartesian topology (handle)
IN	remain_dims	the <i>i</i> -th entry of <code>remain_dims</code> specifies whether the <i>i</i> -th dimension is kept in the subgrid (<code>true</code>) or is dropped (<code>false</code>) (array of logicals)
OUT	newcomm	new communicator with associated Cartesian topology containing the subgrid that includes the calling MPI process (handle)

C binding

```
int MPI_Cart_sub(MPI_Comm comm, const int remain_dims[], MPI_Comm *newcomm)
```

Fortran 2008 binding

```
MPI_Cart_sub(comm, remain_dims, newcomm, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  LOGICAL, INTENT(IN) :: remain_dims(*)
  TYPE(MPI_Comm), INTENT(OUT) :: newcomm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_CART_SUB(COMM, REMAIN_DIMS, NEWCOMM, IERROR)
  INTEGER COMM, NEWCOMM, IERROR
  LOGICAL REMAIN_DIMS(*)
```

`MPI_CART_SUB` can be used to partition the group associated with a communicator that has an associated Cartesian topology into subgroups that form lower-dimensional Cartesian subgrids, and to create for each subgroup a communicator with the associated subgrid Cartesian topology. The topologies of the new communicators describe the subgrids. The number of dimensions of the subgrids is the number of remaining dimensions, i.e., the number of `true` values in `remain_dims`. The numbers of MPI processes in each coordinate direction of the subgrids are the remaining numbers of MPI processes in each coordinate direction of the grid associated with the original communicator, i.e., the values of the original grid dimensions for which the corresponding entry in `remain_dims` is `true`. The periodicity for the remaining dimensions in the new communicator is preserved from the original communicator. If all entries in `remain_dims` are `false` or `comm` is already associated with a zero-dimensional Cartesian topology then `newcomm` is associated with a zero-dimensional Cartesian topology. (This function is closely related to `MPI_COMM_SPLIT`.)

Example 8.8. Creation of nonoverlapping Cartesian subcommunicators with `MPI_CART_SUB`.

Assume that `MPI_Cart_create(..., comm)` has defined a $(2 \times 3 \times 4)$ grid. Let `remain_dims = (true, false, true)`. Then a call to

```
MPI_Cart_sub(comm, remain_dims, &newcomm);
```

will create three communicators each with eight MPI processes in a 2×4 Cartesian topology. If `remain_dims = (false, false, true)` then the call to

```
MPI_Cart_sub(comm, remain_dims, &newcomm);
```

will create six nonoverlapping communicators, each with four MPI processes, in a one-dimensional Cartesian topology.

8.5.8 Low-Level Topology Functions

The two additional functions introduced in this section can be used to implement all other topology functions. In general they will not be called by the user directly, except when creating additional *virtual topology* capabilities other than those provided by MPI. The two calls are both local.

MPI_CART_MAP(comm, ndims, dims, periods, newrank)

IN	comm	input communicator (handle)
IN	ndims	number of dimensions of Cartesian structure (integer)
IN	dims	integer array of size ndims specifying the number of processes in each coordinate direction
IN	periods	logical array of size ndims specifying the periodicity specification in each coordinate direction
OUT	newrank	reordered rank of the calling MPI process; MPI_UNDEFINED if calling MPI process does not belong to grid (integer)

C binding

```
int MPI_Cart_map(MPI_Comm comm, int ndims, const int dims[],
                const int periods[], int *newrank)
```

Fortran 2008 binding

```
MPI_Cart_map(comm, ndims, dims, periods, newrank, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: ndims, dims(ndims)
  LOGICAL, INTENT(IN) :: periods(ndims)
  INTEGER, INTENT(OUT) :: newrank
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_CART_MAP(COMM, NDIMS, DIMS, PERIODS, NEWRANK, IERROR)
  INTEGER COMM, NDIMS, DIMS(*), NEWRANK, IERROR
  LOGICAL PERIODS(*)
```

MPI_CART_MAP computes an “optimal” placement for the calling MPI process on the physical machine. A possible implementation of this function is to always return the rank of the calling MPI process, that is, not to perform any reordering.

Advice to implementors. The function `MPI_CART_CREATE(comm, ndims, dims, periods, reorder, comm_cart)`, with `reorder = true` can be implemented by calling `MPI_CART_MAP(comm, ndims, dims, periods, newrank)`, then calling `MPI_COMM_SPLIT(comm, color, key, comm_cart)`, with `color = 0` if `newrank` \neq `MPI_UNDEFINED`, `color = MPI_UNDEFINED` otherwise, and `key = newrank`. If `ndims` is zero then a zero-dimensional Cartesian topology is created.

The function `MPI_CART_SUB(comm, remain_dims, comm_new)` can be implemented by a call to `MPI_COMM_SPLIT(comm, color, key, comm_new)`, using a single number encoding of the lost dimensions as `color` and a single number encoding of the preserved dimensions as `key`.

All other Cartesian topology functions can be implemented locally, using the topology information that is cached with the communicator. (*End of advice to implementors.*)

The corresponding function for graph structures is as follows.

`MPI_GRAPH_MAP(comm, nnodes, index, edges, newrank)`

IN	<code>comm</code>	input communicator (handle)
IN	<code>nnodes</code>	number of graph nodes (integer)
IN	<code>index</code>	integer array specifying the graph structure (for details see the definition of <code>MPI_GRAPH_CREATE</code>)
IN	<code>edges</code>	integer array specifying the graph structure
OUT	<code>newrank</code>	reordered rank of the calling MPI process; <code>MPI_UNDEFINED</code> if the calling MPI process does not belong to graph (integer)

C binding

```
int MPI_Graph_map(MPI_Comm comm, int nnodes, const int index[],
                 const int edges[], int *newrank)
```

Fortran 2008 binding

```
MPI_Graph_map(comm, nnodes, index, edges, newrank, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: nnodes, index(nnodes), edges(*)
  INTEGER, INTENT(OUT) :: newrank
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GRAPH_MAP(COMM, NNODES, INDEX, EDGES, NEWRANK, IERROR)
  INTEGER COMM, NNODES, INDEX(*), EDGES(*), NEWRANK, IERROR
```

Advice to implementors. The function `MPI_GRAPH_CREATE(comm, nnodes, index, edges, reorder, comm_graph)`, with `reorder = true` can be implemented by calling `MPI_GRAPH_MAP(comm, nnodes, index, edges, newrank)`, then calling `MPI_COMM_SPLIT(comm, color, key, comm_graph)`, with `color = 0` if `newrank` \neq `MPI_UNDEFINED`, `color = MPI_UNDEFINED` otherwise, and `key = newrank`.

All other graph topology functions can be implemented locally, using the topology information that is cached with the communicator. (*End of advice to implementors.*)

8.6 Neighborhood Collective Communication on Virtual Topologies

Virtual topologies specify a communication graph, but they implement no communication function themselves. Many applications require sparse nearest neighbor communications that can be expressed as graph topologies. We now describe several collective operations that perform communication along the edges of a graph representing a *virtual topology*. All of these functions are collective; i.e., they must be called by all MPI processes in the specified communicator. See Section 6 for an overview of other dense (global) collective communication operations and the semantics of collective operations.

If the graph was created with `MPI_DIST_GRAPH_CREATE_ADJACENT` with `sources` and `destinations` containing `0, . . . , n-1`, where `n` is the number of MPI processes in the group of `comm_old` (i.e., the graph is fully connected and also includes an edge from each node to itself), then the sparse neighborhood communication routine performs the same data exchange as the corresponding dense (fully-connected) collective operation. In the case of a Cartesian communicator, only nearest neighbor communication is provided, corresponding to `rank_source` and `rank_dest` in `MPI_CART_SHIFT` with input `disp = 1`.

Rationale. Neighborhood collective communications enable communication on a *virtual topology*. This high-level specification of data exchange among neighboring MPI processes enables optimizations in the MPI library because the communication pattern is known statically (the topology). Thus, the implementation can compute optimized message schedules during creation of the topology [40]. This functionality can significantly simplify the implementation of neighbor exchanges [36]. (*End of rationale.*)

For a distributed graph topology, created with `MPI_DIST_GRAPH_CREATE`, the sequence of neighbors in the send and receive buffers at each MPI process is defined as the sequence returned by `MPI_DIST_GRAPH_NEIGHBORS` for destinations and sources, respectively. For a general graph topology, created with `MPI_GRAPH_CREATE`, the use of neighborhood collective communication is restricted to adjacency matrices, where the number of edges between any two MPI processes is defined to be the same for both MPI processes (i.e., with a symmetric adjacency matrix). In this case, the order of neighbors in the send and receive buffers is defined as the sequence of neighbors as returned by `MPI_GRAPH_NEIGHBORS`. Note that graph topologies should generally be replaced by the distributed graph topologies.

For a Cartesian topology, created with `MPI_CART_CREATE`, the sequence of neighbors in the send and receive buffers at each MPI process is defined by the order of the dimensions, first the neighbor in the negative direction and then in the positive direction with displacement 1. The numbers of sources and destinations in the communication routines are `2*ndims` with `ndims` defined in `MPI_CART_CREATE`. If a neighbor does not exist, i.e., at the border of a Cartesian topology in the case of a nonperiodic virtual grid dimension (i.e., `periods[. . .]=false`), then this neighbor is defined to be `MPI_PROC_NULL`.

If a neighbor in any of the functions is `MPI_PROC_NULL`, then the neighborhood collective communication behaves like a point-to-point communication with `MPI_PROC_NULL` in this direction. That is, the buffer is still part of the sequence of neighbors but it is neither communicated nor updated.

8.6.1 Neighborhood Gather

In the neighborhood gather operation, each MPI process i gathers data items from each MPI process j if an edge (j, i) exists in the topology graph, and each MPI process i sends the same data items to all MPI processes j where an edge (i, j) exists. The send buffer is sent to each neighboring MPI process and the l -th block in the receive buffer is received from the l -th neighbor.

MPI_NEIGHBOR_ALLGATHER(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm)

IN	sendbuf	starting address of send buffer (choice)
IN	sendcount	number of elements sent to each neighbor (non-negative integer)
IN	sendtype	datatype of send buffer elements (handle)
OUT	recvbuf	starting address of receive buffer (choice)
IN	recvcount	number of elements received from each neighbor (non-negative integer)
IN	recvtype	datatype of receive buffer elements (handle)
IN	comm	communicator with associated virtual topology (handle)

C binding

```
int MPI_Neighbor_allgather(const void *sendbuf, int sendcount,
    MPI_Datatype sendtype, void *recvbuf, int recvcount,
    MPI_Datatype recvtype, MPI_Comm comm)
```

```
int MPI_Neighbor_allgather_c(const void *sendbuf, MPI_Count sendcount,
    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
    MPI_Datatype recvtype, MPI_Comm comm)
```

Fortran 2008 binding

```
MPI_Neighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Neighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

2 Fortran binding

```
3 MPI_NEIGHBOR_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT,
4     RECVTYPE, COMM, IERROR)
```

```
5     <type> SENDBUF(*), RECVBUF(*)
```

```
6     INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, IERROR
```

```
7
8     The MPI_NEIGHBOR_ALLGATHER procedure supports Cartesian communicators,
9     graph communicators, and distributed graph communicators as described in Section 8.6. If
10    comm is a distributed graph communicator, the outcome is as if each MPI process executed
11    sends to each of its outgoing neighbors and receives from each of its incoming neighbors:
```

```
12    MPI_Dist_graph_neighbors_count(comm, &indegree, &outdegree, &weighted);
```

```
13    int *srcs=(int*)malloc(indegree*sizeof(int));
```

```
14    int *dsts=(int*)malloc(outdegree*sizeof(int));
```

```
15    MPI_Dist_graph_neighbors(comm, indegree, srcs, MPI_UNWEIGHTED,
16        outdegree, dsts, MPI_UNWEIGHTED);
```

```
17    int k;
```

```
18    /* assume sendbuf and recvbuf are of type (char*) */
```

```
19    for(k=0; k<outdegree; ++k)
```

```
20        MPI_Isend(sendbuf, sendcount, sendtype, dsts[k],...);
```

```
21
22    for(k=0; k<indegree; ++k)
```

```
23        MPI_Irecv(recvbuf+k*recvcount*extent(recvtype), recvcount, recvtype,
24            srcs[k],...);
```

```
25    MPI_Waitall(...);
```

```
26
27    Figure 8.1 shows the neighborhood gather communication of one MPI process with
28    outgoing neighbors  $d_0 \dots d_3$  and incoming neighbors  $s_0 \dots s_5$ . The MPI process will send
29    its sendbuf to all four destinations (outgoing neighbors) and it will receive the contribution
30    from all six sources (incoming neighbors) into separate locations of its receive buffer.
```

```
31    All arguments are significant on all MPI processes and the argument comm must have
32    identical values on all MPI processes.
```

```
33    The type signature associated with sendcount, sendtype at an MPI process must be
34    equal to the type signature associated with recvcount, recvtype at all other MPI processes.
35    This implies that the amount of data sent must be equal to the amount of data received,
36    pairwise between every pair of communicating MPI processes. Distinct type maps between
37    sender and receiver are still allowed.
```

```
38    Rationale. For optimization reasons, the same type signature is required indepen-
39    dently of whether the topology graph is connected or not. (End of rationale.)
```

```
40
41    The “in place” option is not meaningful for this operation.
```

```
42
43    Example 8.9. Buffer usage of MPI_NEIGHBOR_ALLGATHER in the case of a Cartesian
44    virtual topology.
```

```
45    On a Cartesian virtual topology, the buffer usage in a given direction d with dims[d]=3 and
46    1, respectively during creation of the communicator is described in Figure 8.2.
```

```
47    The figure may apply to any (or multiple) directions in the Cartesian topology. The grey
48    buffers are required in all cases but are only accessed if during creation of the communicator,
```




Figure 8.1: Neighborhood gather communication example




Figure 8.2: Cartesian neighborhood allgather example for 3 and 1 processes in a dimension

`periods[d]` was defined as nonzero (in C) or `.TRUE.` (in Fortran).

The vector variant of `MPI_NEIGHBOR_ALLGATHER` allows one to gather different numbers of elements from each neighbor.

`MPI_NEIGHBOR_ALLGATHERV`(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs, recvtype, comm)

IN	sendbuf	starting address of send buffer (choice)	34
IN	sendcount	number of elements sent to each neighbor (non-negative integer)	35
IN	sendtype	datatype of send buffer elements (handle)	36
OUT	recvbuf	starting address of receive buffer (choice)	37
IN	recvcnts	nonnegative integer array (of length indegree) containing the number of elements that are received from each neighbor	38
IN	displs	integer array (of length indegree). Entry <i>i</i> specifies the displacement (relative to <code>recvbuf</code>) at which to place the incoming data from neighbor <i>i</i>	39
IN	recvtype	datatype of receive buffer elements (handle)	40
			41
			42
			43
			44
			45
			46
			47
			48

```

1      IN      comm      communicator with associated virtual topology
2                      (handle)
3

```

4 C binding

```

5 int MPI_Neighbor_allgatherv(const void *sendbuf, int sendcount,
6     MPI_Datatype sendtype, void *recvbuf, const int recvcnts[],
7     const int displs[], MPI_Datatype recvtype, MPI_Comm comm)
8
9 int MPI_Neighbor_allgatherv_c(const void *sendbuf, MPI_Count sendcount,
10    MPI_Datatype sendtype, void *recvbuf,
11    const MPI_Count recvcnts[], const MPI_Aint displs[],
12    MPI_Datatype recvtype, MPI_Comm comm)

```

13 Fortran 2008 binding

```

14 MPI_Neighbor_allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts,
15     displs, recvtype, comm, ierror)
16     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
17     INTEGER, INTENT(IN) :: sendcount, recvcnts(*), displs(*)
18     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
19     TYPE(*), DIMENSION(..) :: recvbuf
20     TYPE(MPI_Comm), INTENT(IN) :: comm
21     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22
23 MPI_Neighbor_allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts,
24     displs, recvtype, comm, ierror) !(_c)
25     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
26     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcnts(*)
27     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
28     TYPE(*), DIMENSION(..) :: recvbuf
29     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
30     TYPE(MPI_Comm), INTENT(IN) :: comm
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

32 Fortran binding

```

33 MPI_NEIGHBOR_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS,
34     DISPLS, RECVTYPE, COMM, IERROR)
35     <type> SENDBUF(*), RECVBUF(*)
36     INTEGER SENDCOUNT, SENDTYPE, REVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
37     IERROR

```

38 The MPI_NEIGHBOR_ALLGATHERV procedure supports Cartesian communicators,
39 graph communicators, and distributed graph communicators as described in Section 8.6. If
40 `comm` is a distributed graph communicator, the outcome is as if each MPI process executed
41 sends to each of its outgoing neighbors and receives from each of its incoming neighbors:
42

```

43 MPI_Dist_graph_neighbors_count(comm, &indegree, &outdegree, &weighted);
44 int *srcs=(int*)malloc(indegree*sizeof(int));
45 int *dsts=(int*)malloc(outdegree*sizeof(int));
46 MPI_Dist_graph_neighbors(comm, indegree, srcs, MPI_UNWEIGHTED,
47     outdegree, dsts, MPI_UNWEIGHTED);
48 int k;

```



```

/* assume sendbuf and recvbuf are of type (char*) */
for(k=0; k<outdegree; ++k)
    MPI_Isend(sendbuf, sendcount, sendtype, dsts[k],...);

for(k=0; k<indegree; ++k)
    MPI_Irecv(recvbuf+displs[k]*extent(recvtype), recvcnts[k], recvtype,
              srcs[k],...);

MPI_Waitall(...);

```

The type signature associated with `sendcount`, `sendtype` at MPI process j must be equal to the type signature associated with `recvcnts[l]`, `recvtype` at any other MPI process with `srcs[l]=j`. This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of communicating MPI processes. Distinct type maps between sender and receiver are still allowed. The data received from the l -th neighbor is placed into `recvbuf` beginning at offset `displs[l]` elements (in terms of the `recvtype`).

The “in place” option is not meaningful for this operation.

All arguments are significant on all MPI processes and the argument `comm` must have identical values on all MPI processes.

8.6.2 Neighborhood Alltoall

In the neighborhood alltoall operation, each MPI process i receives data items from each MPI process j if an edge (j, i) exists in the topology graph or Cartesian topology. Similarly, each MPI process i sends data items to all MPI processes j where an edge (i, j) exists. This call is more general than `MPI_NEIGHBOR_ALLGATHER` in that different data items can be sent to each neighbor. The k -th block in send buffer is sent to the k -th neighboring MPI process and the l -th block in the receive buffer is received from the l -th neighbor.

`MPI_NEIGHBOR_ALLTOALL`(`sendbuf`, `sendcount`, `sendtype`, `recvbuf`, `recvcount`, `recvtype`,
`comm`)

IN	<code>sendbuf</code>	starting address of send buffer (choice)
IN	<code>sendcount</code>	number of elements sent to each neighbor (non-negative integer)
IN	<code>sendtype</code>	datatype of send buffer elements (handle)
OUT	<code>recvbuf</code>	starting address of receive buffer (choice)
IN	<code>recvcount</code>	number of elements received from each neighbor (non-negative integer)
IN	<code>recvtype</code>	datatype of receive buffer elements (handle)
IN	<code>comm</code>	communicator with associated virtual topology (handle)

C binding

```

int MPI_Neighbor_alltoall(const void *sendbuf, int sendcount,
                          MPI_Datatype sendtype, void *recvbuf, int recvcount,
                          MPI_Datatype recvtype, MPI_Comm comm)

```

```

1 int MPI_Neighbor_alltoall_c(const void *sendbuf, MPI_Count sendcount,
2     MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
3     MPI_Datatype recvtype, MPI_Comm comm)
4

```

Fortran 2008 binding

```

5 MPI_Neighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount,
6     recvtype, comm, ierror)
7     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
8     INTEGER, INTENT(IN) :: sendcount, recvcount
9     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
10    TYPE(*), DIMENSION(..) :: recvbuf
11    TYPE(MPI_Comm), INTENT(IN) :: comm
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13
14 MPI_Neighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount,
15     recvtype, comm, ierror) !(_c)
16    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
17    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
18    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
19    TYPE(*), DIMENSION(..) :: recvbuf
20    TYPE(MPI_Comm), INTENT(IN) :: comm
21    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22

```

Fortran binding

```

23 MPI_NEIGHBOR_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT,
24     RECVTYPE, COMM, IERROR)
25     <type> SENDBUF(*), RECVBUF(*)
26     INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, IERROR
27

```

The `MPI_NEIGHBOR_ALLTOALL` procedure supports Cartesian communicators, graph communicators, and distributed graph communicators as described in Section 8.6. If `comm` is a distributed graph communicator, the outcome is as if each MPI process executed sends to each of its outgoing neighbors and receives from each of its incoming neighbors:

```

32 MPI_Dist_graph_neighbors_count(comm, &indegree, &outdegree, &weighted);
33 int *srcs=(int*)malloc(indegree*sizeof(int));
34 int *dsts=(int*)malloc(outdegree*sizeof(int));
35 MPI_Dist_graph_neighbors(comm, indegree, srcs, MPI_UNWEIGHTED,
36     outdegree, dsts, MPI_UNWEIGHTED);
37 int k;
38 /* assume sendbuf and recvbuf are of type (char*) */
39 for(k=0; k<outdegree; ++k)
40     MPI_Isend(sendbuf+k*sendcount*extent(sendtype), sendcount, sendtype,
41         dsts[k],...);
42
43 for(k=0; k<indegree; ++k)
44     MPI_Irecv(recvbuf+k*recvcount*extent(recvtype), recvcount, recvtype,
45         srcs[k],...);
46 MPI_Waitall(...);
47
48

```

The type signature associated with `sendcount`, `sendtype` at an MPI process must be equal to the type signature associated with `recvcount`, `recvtype` at any other MPI process. This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of communicating MPI processes. Distinct type maps between sender and receiver are still allowed.

The “in place” option is not meaningful for this operation.

All arguments are significant on all MPI processes and the argument `comm` must have identical values on all MPI processes.

Example 8.10. Buffer usage of `MPI_NEIGHBOR_ALLTOALL` in the case of a Cartesian virtual topology.

For a halo communication on a Cartesian grid, the buffer usage in a given direction `d` with `dims[d]=3` and 1, respectively during creation of the communicator is described in Figure 8.3. The figure may apply to any (or multiple) directions in the Cartesian topology. The grey buffers are required in all cases but are only accessed if during creation of the communicator, `periods[d]` was defined as nonzero (in C) or `.TRUE.` (in Fortran).

If `sendbuf` and `recvbuf` are declared as `(char *)` and contain a sequence of buffers each described by `sendcount,sendtype` and `recvbuf,recvtype`, then after `MPI_NEIGHBOR_ALLTOALL` on a Cartesian communicator returned, the content of the `recvbuf` is as if the following code is executed:

```
MPI_Cartdim_get(comm, &ndims);
MPI_Type_get_extent(sendtype, &send_lb, &send_extent);
MPI_Type_get_extent(recvtype, &recv_lb, &recv_extent);
for( /*direction*/ d=0; d < ndims; d++) {
    MPI_Cart_shift(comm, /*direction*/ d, /*disp*/ 1, &rank_source, &rank_dest);
    MPI_Sendrecv(sendbuf+(d*2+0)*sendcount*send_extent,
                sendcount, sendtype, rank_source, /*sendtag*/ d*2,
                recvbuf+(d*2+1)*recvcount*recv_extent,
                recvcount, recvtype, rank_dest, /*recvtag*/ d*2,
                comm, &status); /*communication in direction of displacment -1*/
    MPI_Sendrecv(sendbuf+(d*2+1)*sendcount*send_extent,
                sendcount, sendtype, rank_dest, /*sendtag*/ d*2+1,
                recvbuf+(d*2+0)*recvcount*recv_extent,
                recvcount, recvtype, rank_source, /*recvtag*/ d*2+1,
                comm, &status); /*communication in direction of displacment +1*/
}
```

The first call to `MPI_Sendrecv` implements the solid arrows’ communication pattern in each diagram of Figure 8.3, whereas the second call is for the dashed arrows’ pattern.

Advice to implementors. For a Cartesian topology, if the grid in a direction `d` is periodic and `dims[d]` is equal to 1 or 2, then `rank_source` and `rank_dest` are identical, but still all `ndims` send and `ndims` receive operations use different buffers. If in this case, the two send and receive operations per direction or of all directions are internally parallelized, then the several send and receive operations for the same sender-receiver MPI process pair shall be initiated in the same sequence on sender and receiver side or they shall be distinguished by different tags. The code above shows a valid sequence of operations and tags. (*End of advice to implementors.*)

The vector variant of `MPI_NEIGHBOR_ALLTOALL` allows sending/receiving different numbers of elements to and from each neighbor.

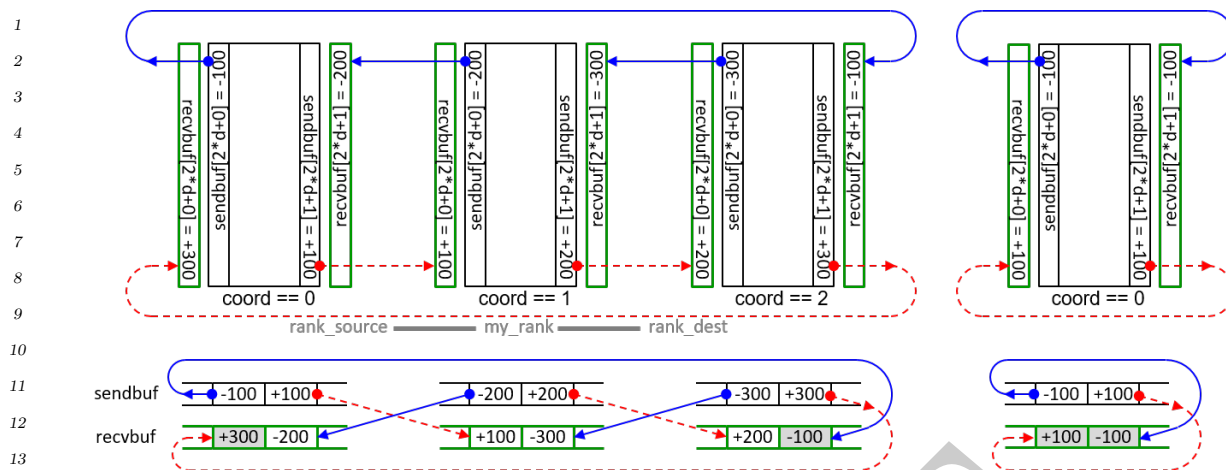


Figure 8.3: Cartesian neighborhood alltoall example for 3 and 1 MPI processes in a dimension

```

MPI_NEIGHBOR_ALLTOALLV(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcoun-
                        t, rdispls, recvtype, comm)
IN      sendbuf          starting address of send buffer (choice)
IN      sendcounts       nonnegative integer array (of length outdegree)
                        specifying the number of elements to send to each
                        neighbor
IN      sdispls          integer array (of length outdegree). Entry j specifies
                        the displacement (relative to sendbuf) from which to
                        send the outgoing data to neighbor j
IN      sendtype         datatype of send buffer elements (handle)
OUT     recvbuf          starting address of receive buffer (choice)
IN      recvcoun-       nonnegative integer array (of length indegree)
                        specifying the number of elements that are received
                        from each neighbor
IN      rdispls          integer array (of length indegree). Entry i specifies
                        the displacement (relative to recvbuf) at which to
                        place the incoming data from neighbor i
IN      recvtype         datatype of receive buffer elements (handle)
IN      comm             communicator with associated virtual topology
                        (handle)

```

C binding

```

int MPI_Neighbor_alltoallv(const void *sendbuf, const int sendcounts[],
                           const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
                           const int recvcoun-
                           t, const int rdispls[],
                           MPI_Datatype recvtype, MPI_Comm comm)

```

```

int MPI_Neighbor_alltoallv_c(const void *sendbuf, const MPI_Count sendcounts[],
                             const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf,
                             const MPI_Count recvcnts[], const MPI_Aint rdispls[],
                             MPI_Datatype recvtype, MPI_Comm comm)

```

Fortran 2008 binding

```

MPI_Neighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
                       recvcnts, rdispls, recvtype, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcounts(*), sdispls(*), recvcnts(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Neighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
                       recvcnts, rdispls, recvtype, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcnts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_NEIGHBOR_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF,
                       REVCOUNTS, RDISPLS, RECVTYPE, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, REVCOUNTS(*), RDISPLS(*),
RECVTYPE, COMM, IERROR

```

The `MPI_NEIGHBOR_ALLTOALLV` procedure supports Cartesian communicators, graph communicators, and distributed graph communicators as described in Section 8.6. If `comm` is a distributed graph communicator, the outcome is as if each MPI process executed sends to each of its outgoing neighbors and receives from each of its incoming neighbors:

```

MPI_Dist_graph_neighbors_count(comm, &indegree, &outdegree, &weighted);
int *srcs=(int*)malloc(indegree*sizeof(int));
int *dsts=(int*)malloc(outdegree*sizeof(int));
MPI_Dist_graph_neighbors(comm, indegree, srcs, MPI_UNWEIGHTED,
                        outdegree, dsts, MPI_UNWEIGHTED);

int k;

/* assume sendbuf and recvbuf are of type (char*) */
for(k=0; k<outdegree; ++k)
    MPI_Isend(sendbuf+sdispls[k]*extent(sendtype), sendcounts[k],
              sendtype, dsts[k],...);

for(k=0; k<indegree; ++k)
    MPI_Irecv(recvbuf+rdispls[k]*extent(recvtype), recvcnts[k],
              recvtype, srcs[k],...);

```

```
1
2 MPI_Waitall(...);
```

3 The type signature associated with `sendcounts[k]`, `sendtype` with `dsts[k]=j` at MPI process i must be equal to the type signature associated with `recvcounts[l]`, `recvtype` with `srcs[l]=i` at MPI process j . This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of communicating MPI processes. Distinct type maps between sender and receiver are still allowed. The data in the `sendbuf` beginning at offset `sdispls[k]` elements (in terms of the `sendtype`) is sent to the k -th outgoing neighbor. The data received from the l -th incoming neighbor is placed into `recvbuf` beginning at offset `rdispls[l]` elements (in terms of the `recvtype`).

4 The “in place” option is not meaningful for this operation.

5 All arguments are significant on all MPI processes and the argument `comm` must have identical values on all MPI processes.

6 `MPI_NEIGHBOR_ALLTOALLW` allows one to send and receive with different datatypes to and from each neighbor.

```
17
18 MPI_NEIGHBOR_ALLTOALLW(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcounts,
19 rdispls, recvtypes, comm)
```

20	IN	<code>sendbuf</code>	starting address of send buffer (choice)
21			
22	IN	<code>sendcounts</code>	nonnegative integer array (of length <code>outdegree</code>) specifying the number of elements to send to each neighbor
23			
24			
25	IN	<code>sdispls</code>	integer array (of length <code>outdegree</code>). Entry j specifies the displacement in bytes (relative to <code>sendbuf</code>) from which to take the outgoing data destined for neighbor j (array of integers)
26			
27			
28			
29	IN	<code>sendtypes</code>	array of datatypes (of length <code>outdegree</code>). Entry j specifies the type of data to send to neighbor j (array of handles)
30			
31			
32			
33	OUT	<code>recvbuf</code>	starting address of receive buffer (choice)
34	IN	<code>recvcounts</code>	nonnegative integer array (of length <code>indegree</code>) specifying the number of elements that are received from each neighbor
35			
36			
37	IN	<code>rdispls</code>	integer array (of length <code>indegree</code>). Entry i specifies the displacement in bytes (relative to <code>recvbuf</code>) at which to place the incoming data from neighbor i (array of integers)
38			
39			
40			
41			
42	IN	<code>recvtypes</code>	array of datatypes (of length <code>indegree</code>). Entry i specifies the type of data received from neighbor i (array of handles)
43			
44			
45	IN	<code>comm</code>	communicator with associated virtual topology (handle)
46			
47			
48			

C binding

```

int MPI_Neighbor_alltoallw(const void *sendbuf, const int sendcounts[],
                           const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
                           void *recvbuf, const int recvcounts[], const MPI_Aint rdispls[],
                           const MPI_Datatype recvtypes[], MPI_Comm comm)

```

```

int MPI_Neighbor_alltoallw_c(const void *sendbuf, const MPI_Count sendcounts[],
                              const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
                              void *recvbuf, const MPI_Count recvcounts[],
                              const MPI_Aint rdispls[], const MPI_Datatype recvtypes[],
                              MPI_Comm comm)

```

Fortran 2008 binding

```

MPI_Neighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
                       recvcounts, rdispls, recvtypes, comm, ierror)

```

```

TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcounts(*), recvcounts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Neighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
                       recvcounts, rdispls, recvtypes, comm, ierror) !(_c)

```

```

TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcounts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_NEIGHBOR_ALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF,
                       RECVCOUNTS, RDISPLS, RECVTYPES, COMM, IERROR)

```

```

<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SENDTYPES(*), RECVCOUNTS(*), RECVTYPES(*), COMM,
IERROR
INTEGER(KIND=MPI_ADDRESS_KIND) SDISPLS(*), RDISPLS(*)

```

The `MPI_NEIGHBOR_ALLTOALLW` procedure supports Cartesian communicators, graph communicators, and distributed graph communicators as described in Section 8.6. If `comm` is a distributed graph communicator, the outcome is as if each MPI process executed sends to each of its outgoing neighbors and receives from each of its incoming neighbors:

```

MPI_Dist_graph_neighbors_count(comm, &indegree, &outdegree, &weighted);
int *srcs=(int*)malloc(indegree*sizeof(int));
int *dsts=(int*)malloc(outdegree*sizeof(int));
MPI_Dist_graph_neighbors(comm, indegree, srcs, MPI_UNWEIGHTED,
                        outdegree, dsts, MPI_UNWEIGHTED);
int k;

```

```

1
2 /* assume sendbuf and recvbuf are of type (char*) */
3 for(k=0; k<outdegree; ++k)
4     MPI_Isend(sendbuf+sdispls[k], sendcounts[k], sendtypes[k],
5             dsts[k],...);
6
7 for(k=0; k<indegree; ++k)
8     MPI_Irecv(recvbuf+rdispls[k], recvcounts[k], recvtypes[k],
9             srcs[k],...);
10 MPI_Waitall(...);

```

The type signature associated with `sendcounts[k]`, `sendtypes[k]` with `dsts[k]=j` at MPI process i must be equal to the type signature associated with `recvcounts[l]`, `recvtypes[l]` with `srcs[l]=i` at MPI process j . This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of communicating MPI processes. Distinct type maps between sender and receiver are still allowed.

The “in place” option is not meaningful for this operation.

All arguments are significant on all MPI processes and the argument `comm` must have identical values on all MPI processes.

8.7 Nonblocking Neighborhood Communication on Process Topologies

Nonblocking variants of the neighborhood collective operations allow relaxed synchronization and overlapping of computation and communication. The semantics are similar to nonblocking collective operations as described in Section 6.12.

8.7.1 Nonblocking Neighborhood Gather

`MPI_INEIGHBOR_ALLGATHER(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm, request)`

IN	<code>sendbuf</code>	starting address of send buffer (choice)
IN	<code>sendcount</code>	number of elements sent to each neighbor (non-negative integer)
IN	<code>sendtype</code>	datatype of send buffer elements (handle)
OUT	<code>recvbuf</code>	starting address of receive buffer (choice)
IN	<code>recvcount</code>	number of elements received from each neighbor (non-negative integer)
IN	<code>recvtype</code>	datatype of receive buffer elements (handle)
IN	<code>comm</code>	communicator with associated virtual topology (handle)
OUT	<code>request</code>	communication request (handle)

C binding

```
int MPI_Ineighbor_allgather(const void *sendbuf, int sendcount,
```



```

        MPI_Datatype sendtype, void *recvbuf, int recvcount,
        MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
int MPI_Ineighbor_allgather_c(const void *sendbuf, MPI_Count sendcount,
        MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
        MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Ineighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
        recvtype, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ineighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
        recvtype, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_INEIGHBOR_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT,
        RECVTYPE, COMM, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, REQUEST, IERROR
MPI_INEIGHBOR_ALLGATHER starts a nonblocking variant of
MPI_NEIGHBOR_ALLGATHER.

```

```

MPI_INEIGHBOR_ALLGATHERV(sendbuf, sendcount, sendtype, recvbuf, recvcoun
ts, displs,
        recvtype, comm, request)
IN        sendbuf        starting address of send buffer (choice)
IN        sendcount      number of elements sent to each neighbor
                        (non-negative integer)
IN        sendtype       datatype of send buffer elements (handle)
OUT       recvbuf        starting address of receive buffer (choice)
IN        recvcoun
ts        nonnegative integer array (of length indegree)
                        containing the number of elements that are received
                        from each neighbor

```

1	IN	displs	integer array (of length indegree). Entry <i>i</i> specifies
2			the displacement (relative to <code>recvbuf</code>) at which to
3			place the incoming data from neighbor <i>i</i>
4	IN	recvtype	datatype of receive buffer elements (handle)
5			
6	IN	comm	communicator with associated virtual topology
7			(handle)
8	OUT	request	communication request (handle)
9			

C binding

```
11 int MPI_Ineighbor_allgatherv(const void *sendbuf, int sendcount,
12     MPI_Datatype sendtype, void *recvbuf, const int recvcounts[],
13     const int displs[], MPI_Datatype recvtype, MPI_Comm comm,
14     MPI_Request *request)
```

```
15
16 int MPI_Ineighbor_allgatherv_c(const void *sendbuf, MPI_Count sendcount,
17     MPI_Datatype sendtype, void *recvbuf,
18     const MPI_Count recvcounts[], const MPI_Aint displs[],
19     MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
```

Fortran 2008 binding

```
20
21 MPI_Ineighbor_allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts,
22     displs, recvtype, comm, request, ierror)
23     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
24     INTEGER, INTENT(IN) :: sendcount
25     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
26     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
27     INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*), displs(*)
28     TYPE(MPI_Comm), INTENT(IN) :: comm
29     TYPE(MPI_Request), INTENT(OUT) :: request
30     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
31
32 MPI_Ineighbor_allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts,
33     displs, recvtype, comm, request, ierror) !(_c)
34     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
35     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
36     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
37     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
38     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
39     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
40     TYPE(MPI_Comm), INTENT(IN) :: comm
41     TYPE(MPI_Request), INTENT(OUT) :: request
42     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
43
44 MPI_INEIGHBOR_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS,
45     DISPLS, RECVTYPE, COMM, REQUEST, IERROR)
46     <type> SENDBUF(*), RECVBUF(*)
47     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
48     REQUEST, IERROR
```

MPI_INEIGHBOR_ALLGATHERV starts a nonblocking variant of MPI_NEIGHBOR_ALLGATHERV.

8.7.2 Nonblocking Neighborhood Alltoall

MPI_INEIGHBOR_ALLTOALL(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm, request)

IN	sendbuf	starting address of send buffer (choice)	10
IN	sendcount	number of elements sent to each neighbor (non-negative integer)	11
IN	sendtype	datatype of send buffer elements (handle)	12
OUT	recvbuf	starting address of receive buffer (choice)	13
IN	recvcount	number of elements received from each neighbor (non-negative integer)	14
IN	recvtype	datatype of receive buffer elements (handle)	15
IN	comm	communicator with associated virtual topology (handle)	16
OUT	request	communication request (handle)	17

C binding

```
int MPI_Ineighbor_alltoall(const void *sendbuf, int sendcount,
    MPI_Datatype sendtype, void *recvbuf, int recvcount,
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
```

```
int MPI_Ineighbor_alltoall_c(const void *sendbuf, MPI_Count sendcount,
    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Ineighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, request, ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER, INTENT(IN) :: sendcount, recvcount
```

```
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```
TYPE(MPI_Comm), INTENT(IN) :: comm
```

```
TYPE(MPI_Request), INTENT(OUT) :: request
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Ineighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, request, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
```

```
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
```

```

1     TYPE(MPI_Comm), INTENT(IN) :: comm
2     TYPE(MPI_Request), INTENT(OUT) :: request
3     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

5 MPI_INEIGHBOR_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT,
6     RECVTYPE, COMM, REQUEST, IERROR)
7     <type> SENDBUF(*), RECVBUF(*)
8     INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, REQUEST, IERROR
9

```

10 MPI_INEIGHBOR_ALLTOALL starts a nonblocking variant of
11 MPI_NEIGHBOR_ALLTOALL.

```

13 MPI_INEIGHBOR_ALLTOALLV(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcoun
14     t, rdispls, recvtype, comm, request)
15
16 IN     sendbuf           starting address of send buffer (choice)
17 IN     sendcounts       nonnegative integer array (of length outdegree)
18                               specifying the number of elements to send to each
19                               neighbor
20 IN     sdispls          integer array (of length outdegree). Entry j specifies
21                               the displacement (relative to sendbuf) from which
22                               send the outgoing data to neighbor j
23
24 IN     sendtype         datatype of send buffer elements (handle)
25 OUT    recvbuf         starting address of receive buffer (choice)
26 IN     recvcoun
27 t                     nonnegative integer array (of length indegree)
28                               specifying the number of elements that are received
29                               from each neighbor
30 IN     rdispls          integer array (of length indegree). Entry i specifies
31                               the displacement (relative to recvbuf) at which to
32                               place the incoming data from neighbor i
33 IN     recvtype         datatype of receive buffer elements (handle)
34 IN     comm             communicator with associated virtual topology
35                               (handle)
36 OUT    request         communication request (handle)
37

```

C binding

```

39 int MPI_Ineighbor_alltoallv(const void *sendbuf, const int sendcounts[],
40     const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
41     const int recvcoun
42 t, const int rdispls[],
43     MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
44
45 int MPI_Ineighbor_alltoallv_c(const void *sendbuf,
46     const MPI_Count sendcounts[], const MPI_Aint sdispls[],
47     MPI_Datatype sendtype, void *recvbuf,
48     const MPI_Count recvcoun
49 t, const MPI_Aint rdispls[],
50     MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Ineighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
    recvcoun
    counts, rdispls, recvtype, comm, request, ierror)
    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
    INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
    recvcoun
    counts(*), rdispls(*)
    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
    TYPE(MPI_Comm), INTENT(IN) :: comm
    TYPE(MPI_Request), INTENT(OUT) :: request
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Ineighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
    recvcoun
    counts, rdispls, recvtype, comm, request, ierror) !(_c)
    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
    recvcoun
    counts(*)
    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
    rdispls(*)
    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
    TYPE(MPI_Comm), INTENT(IN) :: comm
    TYPE(MPI_Request), INTENT(OUT) :: request
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_INEIGHBOR_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF,
    RECVCOUNTS, RDISPLS, RECVTYPE, COMM, REQUEST, IERROR)
    <type> SENDBUF(*), RECVBUF(*)
    INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, REVCOUNTS(*), RDISPLS(*),
    RECVTYPE, COMM, REQUEST, IERROR

```

MPI_INEIGHBOR_ALLTOALLV starts a nonblocking variant of MPI_NEIGHBOR_ALLTOALLV.

```

MPI_INEIGHBOR_ALLTOALLW(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcoun
    ts, rdispls, recvtypes, comm, request)

```

IN	sendbuf	starting address of send buffer (choice)	37
IN	sendcounts	nonnegative integer array (of length outdegree) specifying the number of elements to send to each neighbor	38 39 40 41
IN	sdispls	integer array (of length outdegree). Entry j specifies the displacement in bytes (relative to sendbuf) from which to take the outgoing data destined for neighbor j (array of integers)	42 43 44 45
IN	sendtypes	array of datatypes (of length outdegree). Entry j specifies the type of data to send to neighbor j (array of handles)	46 47 48

1	OUT	recvbuf	starting address of receive buffer (choice)
2	IN	recvcounts	nonnegative integer array (of length indegree)
3			specifying the number of elements that are received
4			from each neighbor
5	IN	rdispls	integer array (of length indegree). Entry <i>i</i> specifies
6			the displacement in bytes (relative to <code>recvbuf</code>) at
7			which to place the incoming data from neighbor <i>i</i>
8			(array of integers)
9			
10	IN	recvtypes	array of datatypes (of length indegree). Entry <i>i</i>
11			specifies the type of data received from neighbor <i>i</i>
12			(array of handles)
13	IN	comm	communicator with associated virtual topology
14			(handle)
15			
16	OUT	request	communication request (handle)
17			

C binding

```

19 int MPI_Ineighbor_alltoallw(const void *sendbuf, const int sendcounts[],
20     const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
21     void *recvbuf, const int recvcounts[], const MPI_Aint rdispls[],
22     const MPI_Datatype recvtypes[], MPI_Comm comm,
23     MPI_Request *request)
24
25 int MPI_Ineighbor_alltoallw_c(const void *sendbuf,
26     const MPI_Count sendcounts[], const MPI_Aint sdispls[],
27     const MPI_Datatype sendtypes[], void *recvbuf,
28     const MPI_Count recvcounts[], const MPI_Aint rdispls[],
29     const MPI_Datatype recvtypes[], MPI_Comm comm,
30     MPI_Request *request)

```

Fortran 2008 binding

```

32 MPI_Ineighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
33     recvcounts, rdispls, recvtypes, comm, request, ierror)
34     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
35     INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), recvcounts(*)
36     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
37     rdispls(*)
38     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
39     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
40     TYPE(MPI_Comm), INTENT(IN) :: comm
41     TYPE(MPI_Request), INTENT(OUT) :: request
42     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
43
44 MPI_Ineighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
45     recvcounts, rdispls, recvtypes, comm, request, ierror) !(_c)
46     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
47     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
48     recvcounts(*)

```

```

INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
    rdispls(*)
TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_INEIGHBOR_ALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF,
    RECVCOUNTS, RDISPLS, RECVTYPES, COMM, REQUEST, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SENDTYPES(*), RECVCOUNTS(*), RECVTYPES(*), COMM,
    REQUEST, IERROR
INTEGER(KIND=MPI_ADDRESS_KIND) SDISPLS(*), RDISPLS(*)
MPI_INEIGHBOR_ALLTOALLW starts a nonblocking variant of
MPI_NEIGHBOR_ALLTOALLW.

```

8.8 Persistent Neighborhood Communication on Process Topologies

Persistent variants of the neighborhood collective operations can offer significant performance benefits for programs with repetitive communication patterns. The semantics are similar to persistent collective operations as described in Section 6.13.

8.8.1 Persistent Neighborhood Gather

```

MPI_NEIGHBOR_ALLGATHER_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcnt,
    recvtype, comm, info, request)

```

IN	sendbuf	starting address of send buffer (choice)
IN	sendcount	number of elements sent to each neighbor (non-negative integer)
IN	sendtype	datatype of send buffer elements (handle)
OUT	recvbuf	starting address of receive buffer (choice)
IN	recvcnt	number of elements received from each neighbor (non-negative integer)
IN	recvtype	datatype of receive buffer elements (handle)
IN	comm	communicator with associated virtual topology (handle)
IN	info	info argument (handle)
OUT	request	communication request (handle)

C binding

```

int MPI_Neighbor_allgather_init(const void *sendbuf, int sendcount,
    MPI_Datatype sendtype, void *recvbuf, int recvcnt,

```

```

1         MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
2         MPI_Request *request)
3
4 int MPI_Neighbor_allgather_init_c(const void *sendbuf, MPI_Count sendcount,
5         MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
6         MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
7         MPI_Request *request)

```

Fortran 2008 binding

```

9 MPI_Neighbor_allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount,
10        recvtype, comm, info, request, ierror)
11     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
12     INTEGER, INTENT(IN) :: sendcount, recvcount
13     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
14     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
15     TYPE(MPI_Comm), INTENT(IN) :: comm
16     TYPE(MPI_Info), INTENT(IN) :: info
17     TYPE(MPI_Request), INTENT(OUT) :: request
18     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19
20 MPI_Neighbor_allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount,
21        recvtype, comm, info, request, ierror) !(_c)
22     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
23     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
24     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
25     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
26     TYPE(MPI_Comm), INTENT(IN) :: comm
27     TYPE(MPI_Info), INTENT(IN) :: info
28     TYPE(MPI_Request), INTENT(OUT) :: request
29     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

30 MPI_NEIGHBOR_ALLGATHER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT,
31        RECVTYPE, COMM, INFO, REQUEST, IERROR)
32
33     <type> SENDBUF(*), RECVBUF(*)
34     INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, INFO, REQUEST,
35        IERROR

```

Creates a persistent collective communication request for the neighborhood allgather operation.

```

40 MPI_NEIGHBOR_ALLGATHERV_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcounts,
41        displ, recvtype, comm, info, request)

```

42	IN	sendbuf	starting address of send buffer (choice)
43	IN	sendcount	number of elements sent to each neighbor
44			(non-negative integer)
45			
46	IN	sendtype	datatype of send buffer elements (handle)
47	OUT	recvbuf	starting address of receive buffer (choice)
48			

IN	recvcounts	nonnegative integer array (of length indegree) containing the number of elements that are received from each neighbor	1 2 3
IN	displs	integer array (of length indegree). Entry <i>i</i> specifies the displacement (relative to <i>recvbuf</i>) at which to place the incoming data from neighbor <i>i</i>	4 5 6 7
IN	recvtype	datatype of receive buffer elements (handle)	8
IN	comm	communicator with associated virtual topology (handle)	9 10
IN	info	info argument (handle)	11
OUT	request	communication request (handle)	12 13 14
C binding			15
int MPI_Neighbor_allgatherv_init(const void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, const int recvcounts[], const int displs[], MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info, MPI_Request *request)			16 17 18 19
int MPI_Neighbor_allgatherv_init_c(const void *sendbuf, MPI_Count sendcount, MPI_Datatype sendtype, void *recvbuf, const MPI_Count recvcounts[], const MPI_Aint displs[], MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info, MPI_Request *request)			20 21 22 23 24 25
Fortran 2008 binding			26
MPI_Neighbor_allgatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs, recvtype, comm, info, request, ierror)			27 28
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf			29
INTEGER, INTENT(IN) :: sendcount, displs(*)			30
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype			31
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf			32
INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*)			33
TYPE(MPI_Comm), INTENT(IN) :: comm			34
TYPE(MPI_Info), INTENT(IN) :: info			35
TYPE(MPI_Request), INTENT(OUT) :: request			36
INTEGER, OPTIONAL, INTENT(OUT) :: ierror			37
MPI_Neighbor_allgatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs, recvtype, comm, info, request, ierror) !(_c)			38 39
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf			40
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount			41
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype			42
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf			43
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*)			44
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)			45
TYPE(MPI_Comm), INTENT(IN) :: comm			46
TYPE(MPI_Info), INTENT(IN) :: info			47
TYPE(MPI_Request), INTENT(OUT) :: request			48

1 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

2

3 Fortran binding

4

MPI_NEIGHBOR_ALLGATHERV_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS,
5 DISPLS, RECVTYPE, COMM, INFO, REQUEST, IERROR)

6

<type> SENDBUF(*), RECVBUF(*)

7

INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
8 INFO, REQUEST, IERROR

9

Creates a persistent collective communication request for the neighborhood allgather operation.

10

11

12 8.8.2 Persistent Neighborhood Alltoall

13

14

15

MPI_NEIGHBOR_ALLTOALL_INIT(sendbuf, sendcount, sendtype, recvbuf, recvcount,
16 recvtype, comm, info, request)

17

18

IN sendbuf starting address of send buffer (choice)

19

IN sendcount number of elements sent to each neighbor
20 (non-negative integer)

21

IN sendtype datatype of send buffer elements (handle)

22

OUT recvbuf starting address of receive buffer (choice)

23

IN recvcount number of elements received from each neighbor
24 (non-negative integer)

25

IN recvtype datatype of receive buffer elements (handle)

26

IN comm communicator with associated virtual topology
27 (handle)

28

IN info info argument (handle)

29

OUT request communication request (handle)

30

31

32

33

34 C binding

35

int MPI_Neighbor_alltoall_init(const void *sendbuf, int sendcount,
36 MPI_Datatype sendtype, void *recvbuf, int recvcount,
37 MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
38 MPI_Request *request)

39

int MPI_Neighbor_alltoall_init_c(const void *sendbuf, MPI_Count sendcount,
40 MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
41 MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
42 MPI_Request *request)

43

44 Fortran 2008 binding

45

MPI_Neighbor_alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount,
46 recvtype, comm, info, request, ierror)

47

TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf

48

INTEGER, INTENT(IN) :: sendcount, recvcount

	TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype	1	
	TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf	2	
	TYPE(MPI_Comm), INTENT(IN) :: comm	3	
	TYPE(MPI_Info), INTENT(IN) :: info	4	
	TYPE(MPI_Request), INTENT(OUT) :: request	5	
	INTEGER, OPTIONAL, INTENT(OUT) :: ierror	6	
		7	
	MPI_Neighbor_alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm, info, request, ierror) !(_c)	8	
	TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf	9	
	INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount	10	
	TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype	11	
	TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf	12	
	TYPE(MPI_Comm), INTENT(IN) :: comm	13	
	TYPE(MPI_Info), INTENT(IN) :: info	14	
	TYPE(MPI_Request), INTENT(OUT) :: request	15	
	INTEGER, OPTIONAL, INTENT(OUT) :: ierror	16	
		17	
	Fortran binding	18	
	MPI_NEIGHBOR_ALLTOALL_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, COMM, INFO, REQUEST, IERROR)	19	
	<type> SENDBUF(*), RECVBUF(*)	20	
	INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, INFO, REQUEST, IERROR	21	
		22	
	Creates a persistent collective communication request for the neighborhood alltoall operation.	23	
		24	
		25	
		26	
		27	
	MPI_NEIGHBOR_ALLTOALLV_INIT(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts, rdispls, recvtype, comm, info, request)	28	
		29	
IN	sendbuf	starting address of send buffer (choice)	30
			31
IN	sendcounts	nonnegative integer array (of length outdegree) specifying the number of elements to send to each neighbor	32
			33
			34
IN	sdispls	integer array (of length outdegree). Entry j specifies the displacement (relative to sendbuf) from which send the outgoing data to neighbor j	35
			36
			37
IN	sendtype	datatype of send buffer elements (handle)	38
			39
OUT	recvbuf	starting address of receive buffer (choice)	40
			41
IN	recvcounts	nonnegative integer array (of length indegree) specifying the number of elements that are received from each neighbor	42
			43
			44
IN	rdispls	integer array (of length indegree). Entry i specifies the displacement (relative to recvbuf) at which to place the incoming data from neighbor i	45
			46
			47
IN	recvtype	datatype of receive buffer elements (handle)	48

```

1      IN      comm      communicator with associated virtual topology
2                          (handle)
3
4      IN      info      info argument (handle)
5
6      OUT     request   communication request (handle)

```

C binding

```

8      int MPI_Neighbor_alltoallv_init(const void *sendbuf, const int sendcounts[],
9                          const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
10                         const int rcvcounts[], const int rdispls[],
11                         MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
12                         MPI_Request *request)

```

```

13     int MPI_Neighbor_alltoallv_init_c(const void *sendbuf,
14                         const MPI_Count sendcounts[], const MPI_Aint sdispls[],
15                         MPI_Datatype sendtype, void *recvbuf,
16                         const MPI_Count rcvcounts[], const MPI_Aint rdispls[],
17                         MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
18                         MPI_Request *request)

```

Fortran 2008 binding

```

20     MPI_Neighbor_alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
21                         rcvcounts, rdispls, recvtype, comm, info, request, ierror)
22
23     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
24     INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
25     rcvcounts(*), rdispls(*)
26     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
27     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
28     TYPE(MPI_Comm), INTENT(IN) :: comm
29     TYPE(MPI_Info), INTENT(IN) :: info
30     TYPE(MPI_Request), INTENT(OUT) :: request
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33     MPI_Neighbor_alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
34                         rcvcounts, rdispls, recvtype, comm, info, request, ierror) !(_c)
35     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
36     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
37     rcvcounts(*)
38     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
39     rdispls(*)
40     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
41     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
42     TYPE(MPI_Comm), INTENT(IN) :: comm
43     TYPE(MPI_Info), INTENT(IN) :: info
44     TYPE(MPI_Request), INTENT(OUT) :: request
45     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

46     MPI_NEIGHBOR_ALLTOALLV_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF,
47     RECVCOUNTS, RDISPLS, RECVTYPE, COMM, INFO, REQUEST, IERROR)

```

```

<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*),
    RECVTYPE, COMM, INFO, REQUEST, IERROR

```

Creates a persistent collective communication request for the neighborhood alltoallv operation.

```

MPI_NEIGHBOR_ALLTOALLW_INIT(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
    recvcounts, rdispls, recvtypes, comm, info, request)

IN    sendbuf          starting address of send buffer (choice)
IN    sendcounts       nonnegative integer array (of length outdegree)
                          specifying the number of elements to send to each
                          neighbor
IN    sdispls          integer array (of length outdegree). Entry j specifies
                          the displacement in bytes (relative to sendbuf) from
                          which to take the outgoing data destined for
                          neighbor j (array of integers)
IN    sendtypes        array of datatypes (of length outdegree). Entry j
                          specifies the type of data to send to neighbor j (array
                          of handles)
OUT   recvbuf          starting address of receive buffer (choice)
IN    recvcounts       nonnegative integer array (of length indegree)
                          specifying the number of elements that are received
                          from each neighbor
IN    rdispls          integer array (of length indegree). Entry i specifies
                          the displacement in bytes (relative to recvbuf) at
                          which to place the incoming data from neighbor i
                          (array of integers)
IN    recvtypes        array of datatypes (of length indegree). Entry i
                          specifies the type of data received from neighbor i
                          (array of handles)
IN    comm             communicator with associated virtual topology
                          (handle)
IN    info             info argument (handle)
OUT   request          communication request (handle)

```

C binding

```

int MPI_Neighbor_alltoallw_init(const void *sendbuf, const int sendcounts[],
    const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
    void *recvbuf, const int recvcounts[], const MPI_Aint rdispls[],
    const MPI_Datatype recvtypes[], MPI_Comm comm, MPI_Info info,
    MPI_Request *request)

int MPI_Neighbor_alltoallw_init_c(const void *sendbuf,
    const MPI_Count sendcounts[], const MPI_Aint sdispls[],

```

```

1         const MPI_Datatype sendtypes[], void *recvbuf,
2         const MPI_Count recvcnts[], const MPI_Aint rdispls[],
3         const MPI_Datatype recvtypes[], MPI_Comm comm, MPI_Info info,
4         MPI_Request *request)
5
6 Fortran 2008 binding
7 MPI_Neighbor_alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
8     recvcnts, rdispls, recvtypes, comm, info, request, ierror)
9     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
10    INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), recvcnts(*)
11    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
12    rdispls(*)
13    TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
14    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
15    TYPE(MPI_Comm), INTENT(IN) :: comm
16    TYPE(MPI_Info), INTENT(IN) :: info
17    TYPE(MPI_Request), INTENT(OUT) :: request
18    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19 MPI_Neighbor_alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
20     recvcnts, rdispls, recvtypes, comm, info, request, ierror)
21     !(_c)
22     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
23     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
24     recvcnts(*)
25     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
26     rdispls(*)
27     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
28     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
29     TYPE(MPI_Comm), INTENT(IN) :: comm
30     TYPE(MPI_Info), INTENT(IN) :: info
31     TYPE(MPI_Request), INTENT(OUT) :: request
32     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

34 MPI_NEIGHBOR_ALLTOALLW_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF,
35     RECVCOUNTS, RDISPLS, RECVTYPES, COMM, INFO, REQUEST, IERROR)
36     <type> SENDBUF(*), RECVBUF(*)
37     INTEGER SENDCOUNTS(*), SENDTYPES(*), RECVCOUNTS(*), RECVTYPES(*), COMM,
38     INFO, REQUEST, IERROR
39     INTEGER(KIND=MPI_ADDRESS_KIND) SDISPLS(*), RDISPLS(*)

```

Creates a persistent collective communication request for the neighborhood alltoallw operation.

8.9 An Application Example

```

1  INTEGER ndims, num_neigh
2  LOGICAL reorder
3  PARAMETER (ndims=2, num_neigh=4, reorder=.true.)
4  INTEGER comm, comm_size, comm_cart, dims(ndims), ierr
5  INTEGER neigh_rank(num_neigh), own_coords(ndims), i, j, it
6  LOGICAL periods(ndims)
7  REAL u(0:101,0:101), f(0:101,0:101)
8  DATA dims / ndims * 0 /
9  comm = MPI_COMM_WORLD
10 CALL MPI_COMM_SIZE(comm, comm_size, ierr)
11 ! Set MPI process grid size and periodicity
12 CALL MPI_DIMS_CREATE(comm_size, ndims, dims, ierr)
13 periods(1) = .TRUE.
14 periods(2) = .TRUE.
15 ! Create a grid structure in WORLD group and inquire about own position
16 CALL MPI_CART_CREATE(comm, ndims, dims, periods, reorder, &
17     comm_cart, ierr)
18 CALL MPI_CART_GET(comm_cart, ndims, dims, periods, own_coords, ierr)
19 i = own_coords(1)
20 j = own_coords(2)
21 ! Look up the ranks for the neighbors. Own MPI process coordinates are (i,j).
22 ! Neighbors are (i-1,j), (i+1,j), (i,j-1), (i,j+1) modulo (dims(1),dims(2))
23 CALL MPI_CART_SHIFT(comm_cart, 0,1, neigh_rank(1), neigh_rank(2), ierr)
24 CALL MPI_CART_SHIFT(comm_cart, 1,1, neigh_rank(3), neigh_rank(4), ierr)
25 ! Initialize the grid functions and start the iteration
26 CALL init(u, f)
27 DO it=1,100
28     CALL relax(u, f)
29     ! Exchange data with neighbor processes
30     CALL exchange(u, comm_cart, neigh_rank, num_neigh)
31 END DO
32 CALL output(u)

```

Figure 8.4: Set-up of MPI process structure for two-dimensional parallel Poisson solver

Example 8.11. Neighborhood collective communication in a Cartesian virtual topology. The example in Figures 8.4–8.7 shows how the grid definition and inquiry functions can be used in an application program. A partial differential equation, for instance the Poisson equation, is to be solved on a rectangular domain. First, the MPI processes organize themselves in a two-dimensional structure. Each MPI process then inquires about the ranks of its neighbors in the four directions (up, down, right, left). The numerical problem is solved by an iterative method, the details of which are hidden in the subroutine `relax`. In each relaxation step each MPI process computes new values for the solution grid function at the points `u(1:100,1:100)` owned by the MPI process. Then the values at inter-process boundaries have to be exchanged with neighboring MPI processes. For example, the newly calculated values in `u(1,1:100)` must be sent into the halo cells `u(101,1:100)` of the left-hand neighbor with coordinates `(own_coord(1)-1,own_coord(2))`.

```

1  SUBROUTINE exchange(u, comm_cart, neigh_rank, num_neigh)
2  USE MPI
3  REAL u(0:101,0:101)
4  INTEGER comm_cart, num_neigh, neigh_rank(num_neigh)
5  REAL sndbuf(100,num_neigh), rcvbuf(100,num_neigh)
6  INTEGER ierr
7  sndbuf(1:100,1) = u( 1,1:100)
8  sndbuf(1:100,2) = u(100,1:100)
9  sndbuf(1:100,3) = u(1:100, 1)
10 sndbuf(1:100,4) = u(1:100,100)
11 CALL MPI_NEIGHBOR_ALLTOALL(sndbuf, 100, MPI_REAL, rcvbuf, 100, MPI_REAL, &
12                               comm_cart, ierr)
13 ! instead of
14 ! CALL MPI_IRECV(rcvbuf(1,1),100,MPI_REAL, neigh_rank(1),..., rq(1), ierr)
15 ! CALL MPI_ISEND(sndbuf(1,2),100,MPI_REAL, neigh_rank(2),..., rq(2), ierr)
16 ! Always pairing a receive from rank_source with a send to rank_dest
17 ! of the same direction in MPI_CART_SHIFT!
18 ! CALL MPI_IRECV(rcvbuf(1,2),100,MPI_REAL, neigh_rank(2),..., rq(3), ierr)
19 ! CALL MPI_ISEND(sndbuf(1,1),100,MPI_REAL, neigh_rank(1),..., rq(4), ierr)
20 ! CALL MPI_IRECV(rcvbuf(1,3),100,MPI_REAL, neigh_rank(3),..., rq(5), ierr)
21 ! CALL MPI_ISEND(sndbuf(1,4),100,MPI_REAL, neigh_rank(4),..., rq(6), ierr)
22 ! CALL MPI_IRECV(rcvbuf(1,4),100,MPI_REAL, neigh_rank(4),..., rq(7), ierr)
23 ! CALL MPI_ISEND(sndbuf(1,3),100,MPI_REAL, neigh_rank(3),..., rq(8), ierr)
24 ! Of course, one can first start all four IRECV and then all four ISEND,
25 ! Or vice versa, but both in the sequence shown above. Otherwise, the
26 ! matching would be wrong for 2 or only 1 MPI processes in a direction.
27 ! CALL MPI_WAITALL(2*num_neigh, rq, statuses, ierr)
28 u( 0,1:100) = rcvbuf(1:100,1)
29 u(101,1:100) = rcvbuf(1:100,2)
30 u(1:100, 0) = rcvbuf(1:100,3)
31 u(1:100,101) = rcvbuf(1:100,4)
32 END

```

Figure 8.5: Communication routine with local data copying and sparse neighborhood alltoall


```

SUBROUTINE exchange(u, comm_cart, neigh_rank, num_neigh) 1
USE MPI 2
IMPLICIT NONE 3
REAL u(0:101,0:101) 4
INTEGER comm_cart, num_neigh, neigh_rank(num_neigh) 5
INTEGER sndcounts(num_neigh), sndtypes(num_neigh) 6
INTEGER rcvcounts(num_neigh), rcvtypes(num_neigh) 7
INTEGER(KIND=MPI_ADDRESS_KIND) lb, sizeofreal 8
INTEGER(KIND=MPI_ADDRESS_KIND) sdispls(num_neigh), rdispls(num_neigh) 9
INTEGER type_vec, ierr 10
! The following initialization need to be done only once
! before the first call of exchange. 11
CALL MPI_TYPE_GET_EXTENT(MPI_REAL, lb, sizeofreal, ierr) 12
CALL MPI_TYPE_VECTOR(100, 1, 102, MPI_REAL, type_vec, ierr) 13
CALL MPI_TYPE_COMMIT(type_vec, ierr) 14
sndtypes(1:2) = type_vec 15
sndcounts(1:2) = 1 16
sndtypes(3:4) = MPI_REAL 17
sndcounts(3:4) = 100 18
rcvtypes = sndtypes 19
rcvcounts = sndcounts 20
sdispls(1) = ( 1 + 1*102) * sizeofreal ! first element of u( 1, 1:100) 21
sdispls(2) = (100 + 1*102) * sizeofreal ! first element of u(100, 1:100) 22
sdispls(3) = ( 1 + 1*102) * sizeofreal ! first element of u( 1:100, 1 ) 23
sdispls(4) = ( 1 + 100*102) * sizeofreal ! first element of u( 1:100,100 ) 24
rdispls(1) = ( 0 + 1*102) * sizeofreal ! first element of u( 0, 1:100) 25
rdispls(2) = (101 + 1*102) * sizeofreal ! first element of u(101, 1:100) 26
rdispls(3) = ( 1 + 0*102) * sizeofreal ! first element of u( 1:100, 0 ) 27
rdispls(4) = ( 1 + 101*102) * sizeofreal ! first element of u( 1:100,101 ) 28
! the following communication has to be done in each call of exchange 29
CALL MPI_NEIGHBOR_ALLTOALLW(u, sndcounts, sdispls, sndtypes, & 30
u, rcvcounts, rdispls, rcvtypes, & 31
comm_cart, ierr) 32
! The following finalizing need to be done only once 33
! after the last call of exchange. 34
CALL MPI_TYPE_FREE(type_vec, ierr) 35
END 36

```

Figure 8.6: Communication routine with sparse neighborhood alltoallw and without local data copying

```

1  INTEGER ndims, num_neigh
2  LOGICAL reorder
3  PARAMETER (ndims=2, num_neigh=4, reorder=.true.)
4  INTEGER comm, comm_size, comm_cart, dims(ndims), it, ierr
5  LOGICAL periods(ndims)
6  REAL u(0:101,0:101), f(0:101,0:101)
7  DATA dims / ndims * 0 /
8  INTEGER sndcounts(num_neigh), sndtypes(num_neigh)
9  INTEGER rcvcounts(num_neigh), rcvtypes(num_neigh)
10 INTEGER(KIND=MPI_ADDRESS_KIND) lb, sizeofreal
11 INTEGER(KIND=MPI_ADDRESS_KIND) sdispls(num_neigh), rdispls(num_neigh)
12 INTEGER type_vec, request, info, status(MPI_STATUS_SIZE)
13 comm = MPI_COMM_WORLD
14 CALL MPI_COMM_SIZE(comm, comm_size, ierr)
15 ! Set MPI process grid size and periodicity
16 CALL MPI_DIMS_CREATE(comm_size, ndims, dims, ierr)
17 periods(1) = .TRUE.
18 periods(2) = .TRUE.
19 ! Create a grid structure in WORLD group
20 CALL MPI_CART_CREATE(comm, ndims, dims, periods, reorder, &
21                      comm_cart, ierr)
22 ! Create datatypes for the neighborhood communication
23 !
24 ! Insert code from example in Figure 7.4 to create and initialize
25 ! sndcounts, sdispls, sndtypes, rcvcounts, rdispls, and rcvtypes
26 !
27 ! Initialize the neighborhood alltoallw operation
28 info = MPI_INFO_NULL
29 CALL MPI_NEIGHBOR_ALLTOALLW_INIT(u, sndcounts, sdispls, sndtypes, &
30                                 u, rcvcounts, rdispls, rcvtypes, &
31                                 comm_cart, info, request, ierr)
32 ! Initialize the grid functions and start the iteration
33 CALL init(u, f)
34 DO it=1,100
35 ! Start data exchange with neighbor processes
36 CALL MPI_START(request, ierr)
37 ! Compute inner cells
38 CALL relax_inner (u, f)
39 ! Check on completion of neighbor exchange
40 CALL MPI_WAIT(request, status, ierr)
41 ! Compute edge cells
42 CALL relax_edges(u, f)
43 END DO
44 CALL output(u)
45 CALL MPI_REQUEST_FREE(request, ierr)
46 CALL MPI_TYPE_FREE(type_vec, ierr)

```

Figure 8.7: Two-dimensional parallel Poisson solver with persistent sparse neighborhood alltoallw and without local data copying

Chapter 9

MPI Environmental Management

This chapter discusses routines for getting and, where appropriate, setting various parameters that relate to the MPI implementation and the execution environment (such as error handling). The procedures for entering and leaving the MPI execution environment are also described here.

9.1 Implementation Information

9.1.1 Version Inquiries

In order to cope with changes to the MPI standard, there are both compile-time and run-time ways to determine which version of the standard is in use in the environment one is using.

The “version” will be represented by two separate integers, for the version and subversion: In C,

```
#define MPI_VERSION 4
#define MPI_SUBVERSION 1
```

in Fortran,

```
INTEGER :: MPI_VERSION, MPI_SUBVERSION
PARAMETER (MPI_VERSION = 4)
PARAMETER (MPI_SUBVERSION = 1)
```

For runtime determination,

MPI_GET_VERSION(version, subversion)

OUT	version	version number (integer)
OUT	subversion	subversion number (integer)

C binding

```
int MPI_Get_version(int *version, int *subversion)
```

Fortran 2008 binding

```
MPI_Get_version(version, subversion, ierror)
  INTEGER, INTENT(OUT) :: version, subversion
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GET_VERSION(VERSION, SUBVERSION, IERROR)
  INTEGER VERSION, SUBVERSION, IERROR
```

MPI_GET_VERSION can be called at any time in an MPI program. This function must always be thread-safe, as defined in Section 11.6. Valid (MPI_VERSION, MPI_SUBVERSION) pairs in this and previous versions of the MPI standard are (4,1), (4,0), (3,1), (3,0), (2,2), (2,1), (2,0), and (1,2).

MPI_GET_LIBRARY_VERSION(version, resultlen)

OUT	version	version number (string)
OUT	resultlen	Length (in printable characters) of the result returned in version (integer)

C binding

```
int MPI_Get_library_version(char *version, int *resultlen)
```

Fortran 2008 binding

```
MPI_Get_library_version(version, resultlen, ierror)
  CHARACTER(LEN=MPI_MAX_LIBRARY_VERSION_STRING), INTENT(OUT) :: version
  INTEGER, INTENT(OUT) :: resultlen
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GET_LIBRARY_VERSION(VERSION, RESULTLEN, IERROR)
  CHARACTER*(*) VERSION
  INTEGER RESULTLEN, IERROR
```

This routine returns a string representing the version of the MPI library. The version argument is a character string for maximum flexibility.

Advice to implementors. An implementation of MPI should return a different string for every change to its source code or build that could be visible to the user. (*End of advice to implementors.*)

The argument `version` must represent storage that is `MPI_MAX_LIBRARY_VERSION_STRING` characters long. `MPI_GET_LIBRARY_VERSION` may write up to this many characters into `version`.

The number of characters actually written is returned in the output argument, `resultlen`. In C, a null character is additionally stored at `version[resultlen]`. The value of `resultlen` cannot be larger than `MPI_MAX_LIBRARY_VERSION_STRING - 1`. In Fortran, `version` is padded on the right with blank characters. The value of `resultlen` cannot be larger than `MPI_MAX_LIBRARY_VERSION_STRING`.

`MPI_GET_LIBRARY_VERSION` can be called at any time in an MPI program. This function must always be thread-safe, as defined in Section 11.6.

9.1.2 Environmental Inquiries

When using the World Model (Section 11.2), a set of attributes that describe the execution environment is attached to the communicator `MPI_COMM_WORLD` when MPI is initialized. The values of these attributes can be inquired by using the function `MPI_COMM_GET_ATTR` described in Section 7.7 and in Section 19.3.7. It is erroneous to delete these attributes, free their keys, or change their values.

The list of predefined attribute keys include

MPI_TAG_UB: Upper bound for tag value.

MPI_IO: Rank of a node that has regular I/O facilities (possibly myrank). Nodes in the same communicator may return different values for this parameter.

MPI_WTIME_IS_GLOBAL: Boolean variable that indicates whether clocks are synchronized.

When using the Sessions Model (Section 11.3), only the MPI_TAG_UB attribute is available. Vendors may add implementation-specific parameters (such as node number, real memory size, virtual memory size, etc.)

These predefined attributes do not change value between MPI initialization (MPI_INIT) and MPI completion (MPI_FINALIZE), and cannot be updated or deleted by users.

Advice to users. Note that in the C binding, the value returned by these attributes is a *pointer* to an int containing the requested value. (*End of advice to users.*)

The required parameter values are discussed in more detail below:

Tag Values

Tag values range from 0 to the value returned for MPI_TAG_UB, inclusive. These values are guaranteed to be unchanging during the execution of an MPI program. In addition, the tag upper bound value must be *at least* 32767. An MPI implementation is free to make the value of MPI_TAG_UB larger than this; for example, the value $2^{30} - 1$ is also a valid value for MPI_TAG_UB.

In the Sessions Model, the attribute MPI_TAG_UB is attached to all communicators created by MPI_COMM_CREATE_FROM_GROUP and MPI_INTERCOMM_CREATE_FROM_GROUPS, with the same value on all MPI processes in the communicator. In the World Model, the attribute MPI_TAG_UB has the same value on all processes of MPI_COMM_WORLD.

IO Rank

The value returned for MPI_IO is the rank of a processor that can provide language-standard I/O facilities. For Fortran, this means that all of the Fortran I/O operations are supported (e.g., OPEN, REWIND, WRITE). For C, this means that all of the ISO C I/O operations are supported (e.g., fopen, fprintf, lseek).

If every process can provide language-standard I/O, then the value MPI_ANY_SOURCE will be returned. Otherwise, if the calling process can provide language-standard I/O, then its rank will be returned. Otherwise, if some process can provide language-standard I/O then the rank of one such process will be returned. The same value need not be returned by all processes. If no process can provide language-standard I/O, then the value MPI_PROC_NULL will be returned.

Advice to users. Note that input is not collective, and this attribute does *not* indicate which process can or does provide input. (*End of advice to users.*)

Clock Synchronization

The value returned for `MPI_WTIME_IS_GLOBAL` is 1 if clocks at all processes in `MPI_COMM_WORLD` are synchronized, 0 otherwise. A collection of clocks is considered synchronized if explicit effort has been taken to synchronize them. The expectation is that the variation in time, as measured by calls to `MPI_WTIME`, will be less than one half the round-trip time for an MPI message of length zero. If time is measured at a process just before a send and at another process just after a matching receive, the second time should be always higher than the first one.

The attribute `MPI_WTIME_IS_GLOBAL` need not be present when the clocks are not synchronized (however, the attribute key `MPI_WTIME_IS_GLOBAL` is always valid). This attribute may be associated with communicators other than `MPI_COMM_WORLD`.

The attribute `MPI_WTIME_IS_GLOBAL` has the same value on all processes of `MPI_COMM_WORLD`.

Inquire Processor Name

`MPI_GET_PROCESSOR_NAME(name, resultlen)`

OUT	name	A unique specifier for the actual (as opposed to virtual) node.
OUT	resultlen	Length (in printable characters) of the result returned in name

C binding

```
int MPI_Get_processor_name(char *name, int *resultlen)
```

Fortran 2008 binding

```
MPI_Get_processor_name(name, resultlen, ierror)
  CHARACTER(LEN=MPI_MAX_PROCESSOR_NAME), INTENT(OUT) :: name
  INTEGER, INTENT(OUT) :: resultlen
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GET_PROCESSOR_NAME(NAME, RESULTLEN, IERROR)
  CHARACTER*(*) NAME
  INTEGER RESULTLEN, IERROR
```

This routine returns the name of the processor on which it was called at the moment of the call. The name is a character string for maximum flexibility. From this value it must be possible to identify a specific piece of hardware; possible values include “processor 9 in rack 4 of mpp.cs.org” and “231” (where 231 is the actual processor number in the running homogeneous system). The argument `name` must represent storage that is at least `MPI_MAX_PROCESSOR_NAME` characters long. `MPI_GET_PROCESSOR_NAME` may write up to this many characters into `name`.

The number of characters actually written is returned in the output argument, `resultlen`. In C, a null character is additionally stored at `name[resultlen]`. The value of `resultlen` cannot be larger than `MPI_MAX_PROCESSOR_NAME-1`. In Fortran, `name` is padded on the right with blank characters. The value of `resultlen` cannot be larger than `MPI_MAX_PROCESSOR_NAME`.

Rationale. This function allows MPI implementations that do process migration to return the current processor. Note that nothing in MPI *requires* or defines process migration; this definition of MPI_GET_PROCESSOR_NAME simply allows such an implementation. (*End of rationale.*)

Advice to users. The user must provide at least MPI_MAX_PROCESSOR_NAME space to write the processor name—processor names can be this long. The user should examine the output argument, `resultlen`, to determine the actual length of the name. (*End of advice to users.*)

Inquire Hardware Resource Information

MPI_GET_HW_RESOURCE_INFO(hw_info)

OUT hw_info info object created (handle)

C binding

```
int MPI_Get_hw_resource_info(MPI_Info *hw_info)
```

Fortran 2008 binding

```
MPI_Get_hw_resource_info(hw_info, ierror)
  TYPE(MPI_Info), INTENT(OUT) :: hw_info
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_GET_HW_RESOURCE_INFO(HW_INFO, IERROR)
  INTEGER HW_INFO, IERROR
```

MPI_GET_HW_RESOURCE_INFO is a local procedure that returns an info object containing information pertaining to the hardware platform on which the calling MPI process is executing at the moment of the call. This information is stored as (key,value) pairs where each key is the name of a hardware resource type and its value is set to "true" if the calling MPI process is restricted to a single instance of a hardware resource of that type and "false" otherwise. The order in which the keys are stored in `hw_info` is unspecified. This procedure will return different information for MPI processes that are restricted to different hardware resources. Otherwise, info objects with identical (key, value) pairs are returned. The user is responsible for freeing `hw_info` via MPI_INFO_FREE.

The keys stored in the `hw_info` object have a *Uniform Resource Identifier* (URI) format. The first part of the URI indicates the key provider and the second part conforms to the format used by this key provider. The key provider "mpi://" is reserved for exclusive use by the MPI standard.

Advice to implementors. Key provider names could be derived from MPI implementation names (e.g., "mpich://", "openmpi://"), from names of external libraries or pieces of software (e.g., "hwloc://", "pmix://"), from names of programming or execution models (e.g., "openmp://"), from resource manager names (e.g., "slurm://") or from hardware vendor names. (*End of advice to implementors.*)

Advice to users. Users should be cautious when using such keys as comparisons between different providers may not be always meaningful nor relevant. (*End of advice to users.*)

Advice to users. The keys stored in the info object returned by this procedure can be used in MPI_COMM_SPLIT_TYPE with the split_type value MPI_COMM_TYPE_HW_GUIDED or MPI_COMM_TYPE_RESOURCE_GUIDED as key values for the info key "mpi_hw_resource_type". (*End of advice to users.*)

Subsequent calls to MPI_GET_HW_RESOURCE_INFO may return different information throughout the execution of the program because an MPI process can be relocated (e.g., migrated or have its hardware restrictions changed).

9.2 Memory Allocation

In some systems, message-passing and remote-memory-access (RMA) operations run faster when accessing specially allocated memory (e.g., memory that is shared by the other processes in the communicating group on an SMP). MPI provides a mechanism for allocating and freeing such special memory. The use of such memory for message-passing or RMA is not mandatory, and this memory can be used without restrictions as any other dynamically allocated memory. However, implementations may restrict the use of some RMA functionality as defined in Section 12.5.3.

MPI_ALLOC_MEM(size, info, baseptr)

IN	size	size of memory segment in bytes (non-negative integer)
IN	info	info argument (handle)
OUT	baseptr	pointer to beginning of memory segment allocated

C binding

```
int MPI_Alloc_mem(MPI_Aint size, MPI_Info info, void *baseptr)
```

Fortran 2008 binding

```
MPI_Alloc_mem(size, info, baseptr, ierror)
  USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(C_PTR), INTENT(OUT) :: baseptr
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_ALLOC_MEM(SIZE, INFO, BASEPTR, IERROR)
  INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR
  INTEGER INFO, IERROR
```

If the Fortran compiler provides TYPE(C_PTR), then the following generic interface must be provided in the mpi module and should be provided in the (deprecated) mpif.h include

file through overloading, i.e., with the same routine name as the routine with
 INTEGER(KIND=MPI_ADDRESS_KIND) BASEPTR, but with a different specific procedure name:

```

INTERFACE MPI_ALLOC_MEM
  SUBROUTINE MPI_ALLOC_MEM(SIZE, INFO, BASEPTR, IERROR)
    IMPORT :: MPI_ADDRESS_KIND
    INTEGER :: INFO, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR
  END SUBROUTINE
  SUBROUTINE MPI_ALLOC_MEM_CPTR(SIZE, INFO, BASEPTR, IERROR)
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
    IMPORT :: MPI_ADDRESS_KIND
    INTEGER :: INFO, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE
    TYPE(C_PTR) :: BASEPTR
  END SUBROUTINE
END INTERFACE

```

The base procedure name of this overloaded function is `MPI_ALLOC_MEM_CPTR`. The implied specific procedure names are described in Section [19.1.5](#).

By default, the allocated memory shall be aligned to at least the alignment required for load/store accesses of any datatype corresponding to a predefined MPI datatype. The `info` argument may be used to specify a desired alternative minimum alignment in bytes for the allocated memory by setting the value of the key `"mpi_minimum_memory_alignment"` to an integral number equal to a power of two. An implementation may ignore values smaller than the default required alignment. The `info` argument can also be used to provide directives that control the desired location of the allocated memory. Such a directive does not affect the semantics of the call. The corresponding `info` values are implementation-dependent. A null directive value of `info = MPI_INFO_NULL` is always valid.

The function `MPI_ALLOC_MEM` may raise an error of class `MPI_ERR_NO_MEM` to indicate it failed because memory is exhausted.

`MPI_FREE_MEM(base)`

IN	base	initial address of memory segment allocated by
		<code>MPI_ALLOC_MEM</code> (choice)

C binding

```
int MPI_Free_mem(void *base)
```

Fortran 2008 binding

```

MPI_Free_mem(base, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: base
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FREE_MEM(BASE, IERROR)
  <type> BASE(*)
  INTEGER IERROR

```

The function `MPI_FREE_MEM` may raise an error of class `MPI_ERR_BASE` to indicate an invalid base argument.

Rationale. The C bindings of `MPI_ALLOC_MEM` and `MPI_FREE_MEM` are similar to the bindings for the `malloc` and `free` C library calls: a call to `MPI_Alloc_mem(..., &base)` should be paired with a call to `MPI_Free_mem(base)` (one less level of indirection). Both arguments are declared to be of same type `void*` so as to facilitate type casting. The Fortran binding is consistent with the C bindings: the Fortran `MPI_ALLOC_MEM` call returns in `baseptr` the `TYPE(C_PTR)` pointer or the (integer valued) address of the allocated memory. The `base` argument of `MPI_FREE_MEM` is a choice argument, which passes (a reference to) the variable stored at that location. (*End of rationale.*)

Advice to implementors. If `MPI_ALLOC_MEM` allocates special memory, then a design similar to the design of C `malloc` and `free` functions has to be used, in order to find out the size of a memory segment, when the segment is freed. If no special memory is used, `MPI_ALLOC_MEM` simply invokes `malloc`, and `MPI_FREE_MEM` invokes `free`.

A call to `MPI_ALLOC_MEM` can be used in shared memory systems to allocate memory in a shared memory segment. (*End of advice to implementors.*)

Example 9.1. Example of use of `MPI_ALLOC_MEM`, in Fortran with `TYPE(C_PTR)` pointers. We assume 4-byte REALs.

```

USE mpi_f08 ! or USE mpi (not guaranteed with INCLUDE 'mpif.h')
USE, INTRINSIC :: ISO_C_BINDING
TYPE(C_PTR) :: p
REAL, DIMENSION(:, :), POINTER :: a ! no memory is allocated
INTEGER, DIMENSION(2) :: shape
INTEGER(KIND=MPI_ADDRESS_KIND) :: size
shape = (/100,100/)
size = 4 * shape(1) * shape(2) ! assuming 4 bytes per REAL
CALL MPI_ALLOC_MEM(size, MPI_INFO_NULL, p, ierr) ! memory is allocated and
CALL C_F_POINTER(p, a, shape) ! intrinsic ! now accessible via a(i,j)
... ! in ISO_C_BINDING
a(3,5) = 2.71
...
CALL MPI_FREE_MEM(a, ierr) ! memory is freed

```

Example 9.2. Example of use of `MPI_ALLOC_MEM`, in Fortran with nonstandard **Cray-pointers**. We assume 4-byte REALs, and assume that these pointers are address-sized.

```

REAL A
POINTER (P, A(100,100)) ! no memory is allocated
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE
SIZE = 4*100*100
CALL MPI_ALLOC_MEM(SIZE, MPI_INFO_NULL, P, IERR)
! memory is allocated
...
A(3,5) = 2.71
...

```

```
CALL MPI_FREE_MEM(A, IERR) ! memory is freed
```

This code is not Fortran 77 or Fortran 90 code. Some compilers may not support this code or need a special option, e.g., the GNU gFortran compiler needs `-fcray-pointer`.

Advice to implementors. Some compilers map Cray-pointers to address-sized integers, some to `TYPE(C_PTR)` pointers (e.g., Cray Fortran, version 7.3.3). From the user's viewpoint, this mapping is irrelevant because Examples 9.2 should work correctly with an MPI-3.0 (or later) library if Cray-pointers are available. (*End of advice to implementors.*)

Example 9.3. Same example, in C.

```
float (* f)[100][100];
/* no memory is allocated */
MPI_Alloc_mem(sizeof(float)*100*100, MPI_INFO_NULL, &f);
/* memory allocated */
...
(*f)[5][3] = 2.71;
...
MPI_Free_mem(f);
```

9.3 Error Handling

An MPI implementation may be unable or choose not to handle some failures that occur during MPI calls. These can include failures that generate exceptions or traps, such as floating point errors or access violations. The set of failures that are handled by MPI is implementation-dependent. Each such failure causes an error to be raised.

The above text takes precedence over any text on error handling within this document. Specifically, text that states that errors *will* be handled should be read as *may* be handled. More background information about how MPI treats errors can be found in Section 2.8.

A user can associate error handlers to four types of objects: communicators, windows, files, and sessions. The specified error handling routine will be used for any error that occurs during an MPI procedure or an operation that refers to the respective object. Figure 9.1 presents a diagram of the error handler that is invoked in different situations. When the MPI procedure or operation refers to a communicator, window, or file, the error handler for that object will be invoked; otherwise, if the procedure or operation refers to a session, the error handler for the session will be invoked. Some MPI procedures have indirect references to these objects. For example, in a procedure that takes a request handle as a parameter, an error during the corresponding operation is raised on the communicator, window, or file on which the request has been initialized. Similarly, a group contains a reference to the session from which it was derived, and procedures on groups invoke the error handler from that session. The referenced object may have been destroyed before an error is raised (e.g., a procedure on a group derived from a session that has been finalized), in this case, the associated error handler for the object cannot be obtained.

MPI procedures that do not refer to an MPI object from which the associated error handler can be obtained, directly or indirectly, are considered to be attached to the communicator `MPI_COMM_SELF` when using the World Model (see Section 11.2). When

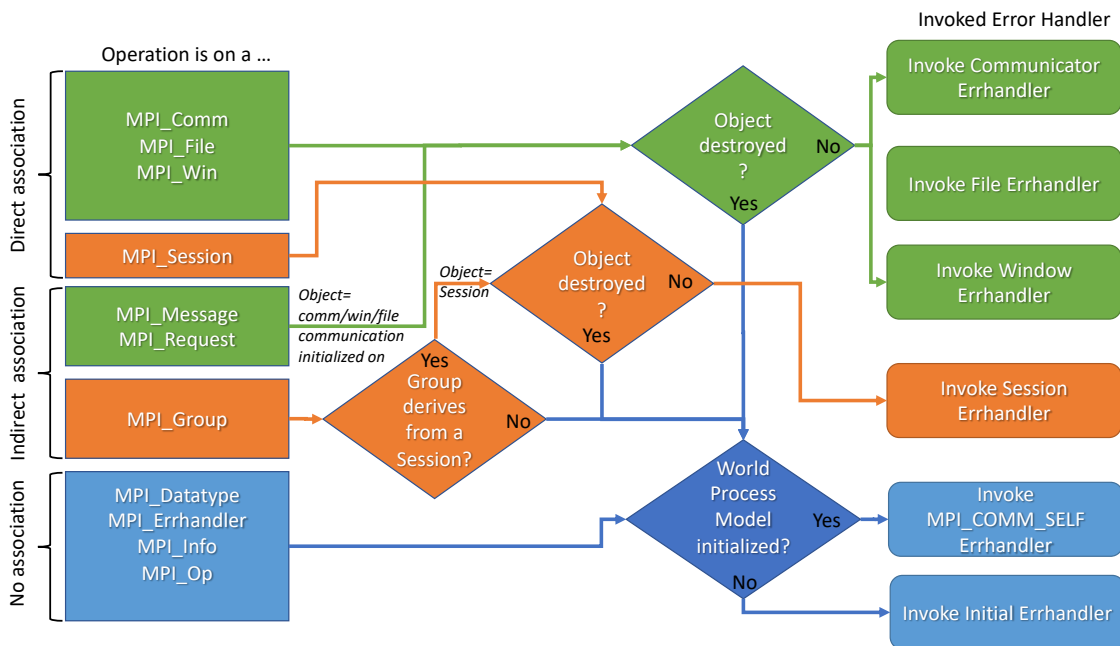


Figure 9.1: Diagram for deciding which error handler is invoked.

MPI_COMM_SELF is not initialized (i.e., before MPI_INIT / MPI_INIT_THREAD, after MPI_FINALIZE, or when using the Sessions Model exclusively) the error raises the initial error handler (set during the launch operation, see 11.8.4). The attachment of error handlers to objects is purely local: different processes may attach different error handlers to corresponding objects.

Several predefined error handlers are available in MPI:

MPI_ERRORS_ARE_FATAL: The handler, when called, causes the program to abort all connected MPI processes. This is similar to calling MPI_ABORT using a communicator containing all connected processes with an implementation-specific value as the errorcode argument.

MPI_ERRORS_ABORT: The handler, when called, is invoked on a communicator in a manner similar to calling MPI_ABORT on that communicator. If the error handler is invoked on an window or file, it is similar to calling MPI_ABORT using a communicator containing the group of MPI processes associated with the window or file, respectively. If the error handler is invoked on a session, the operation aborts only the local MPI process. In all cases, the value that would be provided as the errorcode argument to MPI_ABORT is implementation-specific.

MPI_ERRORS_RETURN: The handler has no effect other than returning the error code to the user.

Advice to implementors. The implementation-specific error information resulting from MPI_ERRORS_ARE_FATAL and MPI_ERRORS_ABORT provided to the invoking environment should be meaningful to the end-user, for example a predefined error class. (End of advice to implementors.)

Implementations may provide additional predefined error handlers and programmers can code their own error handlers.

Unless otherwise requested, the error handler `MPI_ERRORS_ARE_FATAL` is set as the default initial error handler and associated with predefined communicators. Thus, if the user chooses not to control error handling, every error that MPI handles is treated as fatal. Since (almost) all MPI calls return an error code, a user may choose to handle errors in its main code, by testing the return code of MPI calls and executing a suitable recovery code when the call was not successful. In this case, the error handler `MPI_ERRORS_RETURN` will be used. Usually it is more convenient and more efficient not to test for errors after each MPI call, and have such error handled by a nontrivial MPI error handler. Note that unlike predefined communicators, windows and files do not inherit from the initial error handler, as defined in Sections 12.6 and 14.7 respectively.

When an error is raised, MPI will provide the user information about that error using an error code. Some errors might prevent MPI from completing further API calls successfully and those functions will continue to report errors until the cause of the error is corrected or the user terminates the application. The user can make the determination of whether or not to attempt to continue when handling such an error.

Advice to users. For example, users may be unable to correct errors corresponding to some error classes, such as `MPI_ERR_INTERN`. Such errors may cause subsequent MPI calls to complete in error. (*End of advice to users.*)

Advice to implementors. A high-quality implementation will, to the greatest possible extent, circumscribe the impact of an error, so that normal processing can continue after an error handler was invoked. The implementation documentation will provide information on the possible effect of each class of errors and available recovery actions. (*End of advice to implementors.*)

An MPI error handler is an opaque object, which is accessed by a handle. MPI calls are provided to create new error handlers, to associate error handlers with objects, and to test which error handler is associated with an object. C has distinct typedefs for user defined error handling callback functions that accept communicator, file, window, and session arguments. In Fortran there are four user routines.

An error handler object is created by a call to `MPI_XXX_CREATE_ERRHANDLER`, where XXX is, respectively, `COMM`, `WIN`, `FILE`, or `SESSION`.

An error handler is attached to a communicator, window, file, or session by a call to `MPI_XXX_SET_ERRHANDLER`. The error handler must be either a predefined error handler, or an error handler that was created by a call to `MPI_XXX_CREATE_ERRHANDLER`, with matching XXX. An error handler can also be attached to a session using the `errorhandler` argument to `MPI_SESSION_INIT`. The predefined error handlers `MPI_ERRORS_RETURN` and `MPI_ERRORS_ARE_FATAL` can be attached to communicators, windows, files, or sessions.

The error handler currently associated with a communicator, window, file, or session can be retrieved by a call to `MPI_XXX_GET_ERRHANDLER`.

The MPI function `MPI_ERRHANDLER_FREE` can be used to free an error handler that was created by a call to `MPI_XXX_CREATE_ERRHANDLER`.

`MPI_XXX_GET_ERRHANDLER` behave as if a new error handler object is created. That is, once the error handler is no longer needed, `MPI_ERRHANDLER_FREE` should be called with the error handler returned from `MPI_XXX_GET_ERRHANDLER` to mark the error

handler for deallocation. This provides behavior similar to that of MPI_COMM_GROUP and MPI_GROUP_FREE.

Advice to implementors. High-quality implementations should raise an error when an error handler that was created by a call to MPI_XXX_CREATE_ERRHANDLER is attached to an object of the wrong type with a call to MPI_YYY_SET_ERRHANDLER. To do so, it is necessary to maintain, with each error handler, information on the typedef of the associated user function. (*End of advice to implementors.*)

The syntax for these calls is given below.

9.3.1 Error Handlers for Communicators

MPI_COMM_CREATE_ERRHANDLER(comm_errhandler_fn, errhandler)

IN	comm_errhandler_fn	user defined error handling procedure (function)
OUT	errhandler	MPI error handler (handle)

C binding

```
int MPI_Comm_create_errhandler(
    MPI_Comm_errhandler_function *comm_errhandler_fn,
    MPI_Errhandler *errhandler)
```

Fortran 2008 binding

```
MPI_Comm_create_errhandler(comm_errhandler_fn, errhandler, ierror)
    PROCEDURE(MPI_Comm_errhandler_function) :: comm_errhandler_fn
    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_CREATE_ERRHANDLER(COMM_ERRHANDLER_FN, ERRHANDLER, IERROR)
    EXTERNAL COMM_ERRHANDLER_FN
    INTEGER ERRHANDLER, IERROR
```

Creates an error handler that can be attached to communicators.

The user routine should be, in C, a function of type MPI_Comm_errhandler_function, which is defined as

```
typedef void MPI_Comm_errhandler_function(MPI_Comm *comm, int *error_code,
    . . . );
```

The first argument is the communicator in use. The second is the error code to be returned by the MPI routine that raised the error. If the routine would have returned MPI_ERR_IN_STATUS, it is the error code returned in the status for the request that caused the error handler to be invoked. The remaining arguments are “varargs” arguments whose number and meaning is implementation-dependent. An implementation should clearly document these arguments. Addresses are used so that the handler may be written in Fortran. With the Fortran mpi_f08 module, the user routine comm_errhandler_fn should be of the form:

```
ABSTRACT INTERFACE
```

```

SUBROUTINE MPI_Comm_errhandler_function(comm, error_code)
  TYPE(MPI_Comm) :: comm
  INTEGER :: error_code

```

With the Fortran `mpi` module and (deprecated) `mpif.h` include file, the user routine `COMM_ERRHANDLER_FN` should be of the form:

```

SUBROUTINE COMM_ERRHANDLER_FUNCTION(COMM, ERROR_CODE)
  INTEGER COMM, ERROR_CODE

```

Rationale. The variable argument list is provided because it provides an ISO-standard hook for providing additional information to the error handler; without this hook, ISO C prohibits additional arguments. (*End of rationale.*)

Advice to users. A newly created communicator inherits the error handler that is associated with the “parent” communicator. In particular, the user can specify a “global” error handler for all communicators by associating this handler with the communicator `MPI_COMM_WORLD` immediately after initialization. (*End of advice to users.*)

```

MPI_COMM_SET_ERRHANDLER(comm, errhandler)

```

INOUT	comm	communicator (handle)
IN	errhandler	new error handler for communicator (handle)

C binding

```

int MPI_Comm_set_errhandler(MPI_Comm comm, MPI_Errhandler errhandler)

```

Fortran 2008 binding

```

MPI_Comm_set_errhandler(comm, errhandler, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_COMM_SET_ERRHANDLER(COMM, ERRHANDLER, IERROR)
  INTEGER COMM, ERRHANDLER, IERROR

```

Attaches a new error handler to a communicator. The error handler must be either a predefined error handler, or an error handler created by a call to `MPI_COMM_CREATE_ERRHANDLER`.

```

MPI_COMM_GET_ERRHANDLER(comm, errhandler)

```

IN	comm	communicator (handle)
OUT	errhandler	error handler currently associated with communicator (handle)

C binding

```

int MPI_Comm_get_errhandler(MPI_Comm comm, MPI_Errhandler *errhandler)

```

Fortran 2008 binding

```

1 MPI_Comm_get_errhandler(comm, errhandler, ierror)
2     TYPE(MPI_Comm), INTENT(IN) :: comm
3     TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

6 MPI_COMM_GET_ERRHANDLER(COMM, ERRHANDLER, IERROR)
7     INTEGER COMM, ERRHANDLER, IERROR

```

Retrieves the error handler currently associated with a communicator. For example, a library function may register at its entry point the current error handler for a communicator, set its own private error handler for this communicator, and restore before exiting the previous error handler.

9.3.2 Error Handlers for Windows

```

15 MPI_WIN_CREATE_ERRHANDLER(win_errhandler_fn, errhandler)

```

IN	win_errhandler_fn	user defined error handling procedure (function)
OUT	errhandler	MPI error handler (handle)

C binding

```

18 int MPI_Win_create_errhandler(MPI_Win_errhandler_function *win_errhandler_fn,
19                             MPI_Errhandler *errhandler)

```

Fortran 2008 binding

```

20 MPI_Win_create_errhandler(win_errhandler_fn, errhandler, ierror)
21     PROCEDURE(MPI_Win_errhandler_function) :: win_errhandler_fn
22     TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
23     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

24 MPI_WIN_CREATE_ERRHANDLER(WIN_ERRHANDLER_FN, ERRHANDLER, IERROR)
25     EXTERNAL WIN_ERRHANDLER_FN
26     INTEGER ERRHANDLER, IERROR

```

Creates an error handler that can be attached to a window object. The user routine should be, in C, a function of type `MPI_Win_errhandler_function`, which is defined as `typedef void MPI_Win_errhandler_function(MPI_Win *win, int *error_code, ...);`

The first argument is the window in use, the second is the error code to be returned. The remaining arguments are “varargs” arguments whose number and meaning is implementation-dependent. An implementation should clearly document these arguments. With the Fortran `mpi_f08` module, the user routine `win_errhandler_fn` should be of the form:

```

27 ABSTRACT INTERFACE
28     SUBROUTINE MPI_Win_errhandler_function(win, error_code)
29         TYPE(MPI_Win) :: win
30         INTEGER :: error_code

```


With the Fortran `mpi` module and (deprecated) `mpif.h` include file, the user routine `WIN_ERRHANDLER_FN` should be of the form:

```
SUBROUTINE WIN_ERRHANDLER_FUNCTION(WIN, ERROR_CODE)
  INTEGER WIN, ERROR_CODE
```

`MPI_WIN_SET_ERRHANDLER(win, errhandler)`

INOUT	win	window object (handle)
IN	errhandler	new error handler for window (handle)

C binding

```
int MPI_Win_set_errhandler(MPI_Win win, MPI_Errhandler errhandler)
```

Fortran 2008 binding

```
MPI_Win_set_errhandler(win, errhandler, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_SET_ERRHANDLER(WIN, ERRHANDLER, IERROR)
  INTEGER WIN, ERRHANDLER, IERROR
```

Attaches a new error handler to a window. The error handler must be either a pre-defined error handler, or an error handler created by a call to `MPI_WIN_CREATE_ERRHANDLER`.

`MPI_WIN_GET_ERRHANDLER(win, errhandler)`

IN	win	window object (handle)
OUT	errhandler	error handler currently associated with window (handle)

C binding

```
int MPI_Win_get_errhandler(MPI_Win win, MPI_Errhandler *errhandler)
```

Fortran 2008 binding

```
MPI_Win_get_errhandler(win, errhandler, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_GET_ERRHANDLER(WIN, ERRHANDLER, IERROR)
  INTEGER WIN, ERRHANDLER, IERROR
```

Retrieves the error handler currently associated with a window.

9.3.3 Error Handlers for Files

`MPI_FILE_CREATE_ERRHANDLER(file_errhandler_fn, errhandler)`

IN	<code>file_errhandler_fn</code>	user defined error handling procedure (function)
OUT	<code>errhandler</code>	MPI error handler (handle)

C binding

```
int MPI_File_create_errhandler(
    MPI_File_errhandler_function *file_errhandler_fn,
    MPI_Errhandler *errhandler)
```

Fortran 2008 binding

```
MPI_File_create_errhandler(file_errhandler_fn, errhandler, ierror)
    PROCEDURE(MPI_File_errhandler_function) :: file_errhandler_fn
    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_CREATE_ERRHANDLER(FILE_ERRHANDLER_FN, ERRHANDLER, IERROR)
    EXTERNAL FILE_ERRHANDLER_FN
    INTEGER ERRHANDLER, IERROR
```

Creates an error handler that can be attached to a file object. The user routine should be, in C, a function of type `MPI_File_errhandler_function`, which is defined as

```
typedef void MPI_File_errhandler_function(MPI_File *file, int *error_code,
    . . . );
```

The first argument is the file in use, the second is the error code to be returned. The remaining arguments are “varargs” arguments whose number and meaning is implementation-dependent. An implementation should clearly document these arguments.

With the Fortran `mpi_f08` module, the user routine `file_errhandler_fn` should be of the form:

```
ABSTRACT INTERFACE
    SUBROUTINE MPI_File_errhandler_function(file, error_code)
        TYPE(MPI_File) :: file
        INTEGER :: error_code
```

With the Fortran `mpi` module and (deprecated) `mpif.h` include file, the user routine `FILE_ERRHANDLER_FN` should be of the form:

```
SUBROUTINE FILE_ERRHANDLER_FUNCTION(FILE, ERROR_CODE)
    INTEGER FILE, ERROR_CODE
```

`MPI_FILE_SET_ERRHANDLER(file, errhandler)`

INOUT	<code>file</code>	file (handle)
IN	<code>errhandler</code>	new error handler for file (handle)

C binding

```
int MPI_File_set_errhandler(MPI_File file, MPI_Errhandler errhandler)
```

Fortran 2008 binding

```

MPI_File_set_errhandler(file, errhandler, ierror)
    TYPE(MPI_File), INTENT(IN) :: file
    TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_SET_ERRHANDLER(FILE, ERRHANDLER, IERROR)
    INTEGER FILE, ERRHANDLER, IERROR

```

Attaches a new error handler to a file. The error handler must be either a predefined error handler, or an error handler created by a call to `MPI_FILE_CREATE_ERRHANDLER`.

```

MPI_FILE_GET_ERRHANDLER(file, errhandler)

```

IN	file	file (handle)
OUT	errhandler	error handler currently associated with file (handle)

C binding

```

int MPI_File_get_errhandler(MPI_File file, MPI_Errhandler *errhandler)

```

Fortran 2008 binding

```

MPI_File_get_errhandler(file, errhandler, ierror)
    TYPE(MPI_File), INTENT(IN) :: file
    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_GET_ERRHANDLER(FILE, ERRHANDLER, IERROR)
    INTEGER FILE, ERRHANDLER, IERROR

```

Retrieves the error handler currently associated with a file.

9.3.4 Error Handlers for Sessions

```

MPI_SESSION_CREATE_ERRHANDLER(session_errhandler_fn, errhandler)

```

IN	session_errhandler_fn	user defined error handling procedure (function)
OUT	errhandler	MPI error handler (handle)

C binding

```

int MPI_Session_create_errhandler(
    MPI_Session_errhandler_function *session_errhandler_fn,
    MPI_Errhandler *errhandler)

```

Fortran 2008 binding

```

MPI_Session_create_errhandler(session_errhandler_fn, errhandler, ierror)
    PROCEDURE(MPI_Session_errhandler_function) :: session_errhandler_fn
    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler

```

1 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

2

3 **Fortran binding**

4

MPI_SESSION_CREATE_ERRHANDLER(SESSION_ERRHANDLER_FN, ERRHANDLER, IERROR)

5

EXTERNAL SESSION_ERRHANDLER_FN

6

INTEGER ERRHANDLER, IERROR

7

Creates an error handler that can be attached to a session object. In C, the `session_errhandler_fn` argument should be a function of type `MPI_Session_errhandler_function`, which is defined as

10

```
10 typedef void MPI_Session_errhandler_function(MPI_Session *session,
11                                             int *error_code, . . . );
```

12

The first argument is the session in use, the second is the error code to be returned. The remaining arguments are “varargs” arguments whose number and meaning is implementation-dependent. An implementation should clearly document these arguments. With the Fortran `mpi_f08` module, the `session_errhandler_fn` argument should be of the form:

16

ABSTRACT INTERFACE

17

```
17 SUBROUTINE MPI_Session_errhandler_function(session, error_code)
```

18

```
18 TYPE(MPI_Session) :: session
```

19

```
19 INTEGER :: error_code
```

20

With the Fortran `mpi` module and (deprecated) `mpif.h` include file, the `SESSION_ERRHANDLER_FN` argument should be of the form:

22

```
22 SUBROUTINE SESSION_ERRHANDLER_FUNCTION(SESSION, ERROR_CODE)
```

23

```
23 INTEGER SESSION, ERROR_CODE
```

24

25

26

27

MPI_SESSION_SET_ERRHANDLER(session, errhandler)

28

```
28 INOUT session session (handle)
```

29

```
29 IN errhandler new error handler for session (handle)
```

30

31

32

32 **C binding**

33

```
33 int MPI_Session_set_errhandler(MPI_Session session, MPI_Errhandler errhandler)
```

34

34 **Fortran 2008 binding**

35

```
35 MPI_Session_set_errhandler(session, errhandler, ierror)
```

36

```
36 TYPE(MPI_Session), INTENT(IN) :: session
```

37

```
37 TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
```

38

```
38 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

39

40

40 **Fortran binding**

41

```
41 MPI_SESSION_SET_ERRHANDLER(SESSION, ERRHANDLER, IERROR)
```

42

```
42 INTEGER SESSION, ERRHANDLER, IERROR
```

43

Attaches a new error handler to a session. The error handler must be either a pre-defined error handler, or an error handler created by a call to `MPI_SESSION_CREATE_ERRHANDLER`.

46

47

48

MPI_SESSION_GET_ERRHANDLER(session, errhandler)		1
IN	session	session (handle)
OUT	errhandler	error handler currently associated with session (handle)

C binding

```
int MPI_Session_get_errhandler(MPI_Session session, MPI_Errhandler *errhandler)
```

Fortran 2008 binding

```
MPI_Session_get_errhandler(session, errhandler, ierror)
```

```
TYPE(MPI_Session), INTENT(IN) :: session
```

```
TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SESSION_GET_ERRHANDLER(SESSION, ERRHANDLER, IERROR)
```

```
INTEGER SESSION, ERRHANDLER, IERROR
```

Retrieves the error handler currently associated with a session.

9.3.5 Freeing Errorhandlers and Retrieving Error Strings

```
MPI_ERRHANDLER_FREE(errhandler)
```

```
INOUT errhandler MPI error handler (handle)
```

C binding

```
int MPI_Errhandler_free(MPI_Errhandler *errhandler)
```

Fortran 2008 binding

```
MPI_Errhandler_free(errhandler, ierror)
```

```
TYPE(MPI_Errhandler), INTENT(INOUT) :: errhandler
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_ERRHANDLER_FREE(ERRHANDLER, IERROR)
```

```
INTEGER ERRHANDLER, IERROR
```

Marks the error handler associated with `errhandler` for deallocation and sets `errhandler` to `MPI_ERRHANDLER_NULL`. The error handler will be deallocated after all the objects associated with it (communicator, window, or file) have been deallocated.

1 MPI_ERROR_STRING(errorcode, string, resultlen)

2	IN	errorcode	Error code returned by an MPI routine
3			
4	OUT	string	Text that corresponds to the errorcode
5	OUT	resultlen	Length (in printable characters) of the result returned in <code>string</code>
6			
7			

8 C binding

9 int MPI_Error_string(int errorcode, char *string, int *resultlen)

11 Fortran 2008 binding

12 MPI_Error_string(errorcode, string, resultlen, ierror)

13 INTEGER, INTENT(IN) :: errorcode

14 CHARACTER(LEN=MPI_MAX_ERROR_STRING), INTENT(OUT) :: string

15 INTEGER, INTENT(OUT) :: resultlen

16 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

17 Fortran binding

18 MPI_ERROR_STRING(ERRORCODE, STRING, RESULTLEN, IERROR)

19 INTEGER ERRORCODE, RESULTLEN, IERROR

20 CHARACTER*(*) STRING

21 Returns the error string associated with an error code or class. The argument
22 `string` must represent storage that is at least `MPI_MAX_ERROR_STRING` characters long. The
23 number of characters actually written is returned in the output argument, `resultlen`. This
24 function must always be thread-safe, as defined in Section 11.6. It is one of the few routines
25 that may be called before MPI is initialized or after MPI is finalized.

27 *Rationale.* The form of this function was chosen to make the Fortran and C bindings
28 similar. A version that returns a pointer to a string has two difficulties. First, the
29 return string must be statically allocated and different for each error message (allowing
30 the pointers returned by successive calls to `MPI_ERROR_STRING` to point to the
31 correct message). Second, in Fortran, a function declared as returning `CHARACTER*(*)`
32 can not be referenced in, for example, a `PRINT` statement. (*End of rationale.*)

35 9.4 Error Codes and Classes

36 The error codes returned by MPI are left entirely to the implementation (with the exception
37 of `MPI_SUCCESS`). This is done to allow an implementation to provide as much information
38 as possible in the error code (for use with `MPI_ERROR_STRING`).

39 All MPI function calls shall return `MPI_SUCCESS` if and only if the specification of
40 that function has been fulfilled at the point of return. For multiple completion functions,
41 if the function returns `MPI_ERR_IN_STATUS`, the error code in each status object shall be
42 set to `MPI_SUCCESS` if and only if the specification of the operation represented by the
43 corresponding `MPI_Request` has been fulfilled at the point of return.

44 When an operation raises an error, it may not satisfy its specification (for example, a
45 synchronizing operation may not have synchronized) and the content of the output buffers,
46 targeted memory, or output parameters is undefined. However, a valid error code shall
47

always be set when an operation raises an error, whether in the return value, error field in the status object, or element in an array of error codes.

To make it possible for an application to interpret an error code, the routine `MPI_ERROR_CLASS` converts any error code into one of a small set of standard error codes, called **error classes**. Valid error classes are shown in Table 9.1 and Table 9.2.

The error classes are a subset of the error codes: an MPI function may return an error class number; and the function `MPI_ERROR_STRING` can be used to compute the error string associated with an error class. The values defined for MPI error classes are valid MPI error codes.

The error codes satisfy,

$$0 = \text{MPI_SUCCESS} < \text{MPI_ERR_...} \leq \text{MPI_ERR_LASTCODE}.$$

Rationale. The difference between `MPI_ERR_UNKNOWN` and `MPI_ERR_OTHER` is that `MPI_ERROR_STRING` can return useful information about `MPI_ERR_OTHER`.

Note that `MPI_SUCCESS = 0` is necessary to be consistent with C practice; the separation of error classes and error codes allows us to define the error classes this way. Having a known `LASTCODE` is often a nice sanity check as well. (*End of rationale.*)

Advice to implementors. Note that all `MPI_T_` return codes, which must have the prefix `MPI_T_ERR_`, are also required to satisfy

$$0 = \text{MPI_SUCCESS} < \text{MPI_T_ERR_XXX} \leq \text{MPI_ERR_LASTCODE}.$$

as described in Section 15.3.10. (*End of advice to implementors.*)

`MPI_ERROR_CLASS(errorcode, errorclass)`

IN	errorcode	Error code returned by an MPI routine
OUT	errorclass	Error class associated with errorcode

C binding

```
int MPI_Error_class(int errorcode, int *errorclass)
```

Fortran 2008 binding

```
MPI_Error_class(errorcode, errorclass, ierror)
  INTEGER, INTENT(IN) :: errorcode
  INTEGER, INTENT(OUT) :: errorclass
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_ERROR_CLASS(ERRORCODE, ERRORCLASS, IERROR)
  INTEGER ERRORCODE, ERRORCLASS, IERROR
```

The function `MPI_ERROR_CLASS` maps each standard error code (error class) onto itself.

This function must always be thread-safe, as defined in Section 11.6. It is one of the few routines that may be called before MPI is initialized or after MPI is finalized.

Table 9.1: Error classes (Part 1)

MPI_SUCCESS	No error
MPI_ERR_ACCESS	Permission denied
MPI_ERR_AMODE	Error related to the amode passed to MPI_FILE_OPEN
MPI_ERR_ARG	Invalid argument of some other kind
MPI_ERR_ASSERT	Invalid assertion argument
MPI_ERR_BAD_FILE	Invalid file name (e.g., path name too long)
MPI_ERR_BASE	Invalid base passed to MPI_FREE_MEM
MPI_ERR_BUFFER	Invalid buffer pointer argument
MPI_ERR_COMM	Invalid communicator argument
MPI_ERR_CONVERSION	An error occurred in a user supplied data conversion function
MPI_ERR_COUNT	Invalid count argument
MPI_ERR_DIMS	Invalid dimension argument
MPI_ERR_DISP	Invalid displacement argument
MPI_ERR_DUP_DATAREP	Conversion functions could not be registered because a data representation identifier that was already defined was passed to MPI_REGISTER_DATAREP
MPI_ERR_ERRHANDLER	Invalid error handler argument
MPI_ERR_FILE	Invalid file handle argument
MPI_ERR_FILE_EXISTS	File exists
MPI_ERR_FILE_IN_USE	File operation could not be completed, as the file is currently open by some process
MPI_ERR_GROUP	Invalid group argument
MPI_ERR_INFO	Invalid info argument
MPI_ERR_INFO_KEY	Key longer than MPI_MAX_INFO_KEY
MPI_ERR_INFO_NOKEY	Invalid key passed to MPI_INFO_DELETE
MPI_ERR_INFO_VALUE	Value longer than MPI_MAX_INFO_VAL
MPI_ERR_IN_STATUS	Error code is in status
MPI_ERR_INTERN	Internal MPI (implementation) error
MPI_ERR_IO	Other I/O error
MPI_ERR_KEYVAL	Invalid keyval argument
MPI_ERR_LOCKTYPE	Invalid locktype argument
MPI_ERR_NAME	Invalid service name passed to MPI_LOOKUP_NAME
MPI_ERR_NO_MEM	MPI_ALLOC_MEM failed because memory is exhausted
MPI_ERR_NO_SPACE	Not enough space
MPI_ERR_NO_SUCH_FILE	File does not exist
MPI_ERR_NOT_SAME	Collective argument not identical on all processes, or collective routines called in a different order by different processes

Table 9.2: Error classes (Part 2)		1
MPI_ERR_OP	Invalid operation argument	3
MPI_ERR_OTHER	Known error not in this list	4
MPI_ERR_PENDING	Pending request	5
MPI_ERR_PORT	Invalid port name passed to MPI_COMM_CONNECT	7
MPI_ERR_PROC_ABORTED	Operation failed because a peer process has aborted	9
MPI_ERR_QUOTA	Quota exceeded	10
MPI_ERR_RANK	Invalid rank argument	11
MPI_ERR_READ_ONLY	Read-only file or file system	12
MPI_ERR_REQUEST	Invalid request argument	13
MPI_ERR_RMA_ATTACH	Memory cannot be attached (e.g., because of resource exhaustion)	15
MPI_ERR_RMA_CONFLICT	Conflicting accesses to window	16
MPI_ERR_RMA_FLAVOR	Passed window has the wrong flavor for the called function	18
MPI_ERR_RMA_RANGE	Target memory is not part of the win- dow (in the case of a window created with MPI_WIN_CREATE_DYNAMIC, tar- get memory is not attached)	22
MPI_ERR_RMA_SHARED	Memory cannot be shared (e.g., some pro- cess in the group of the specified commu- nicator cannot expose shared memory)	25
MPI_ERR_RMA_SYNC	Wrong synchronization of RMA calls	26
MPI_ERR_ROOT	Invalid root argument	27
MPI_ERR_SERVICE	Invalid service name passed to MPI_UNPUBLISH_NAME	29
MPI_ERR_SESSION	Invalid session argument	30
MPI_ERR_SIZE	Invalid size argument	31
MPI_ERR_SPAWN	Error in spawning processes	32
MPI_ERR_TAG	Invalid tag argument	33
MPI_ERR_TOPOLOGY	Invalid topology argument	34
MPI_ERR_TRUNCATE	Message truncated on receive	35
MPI_ERR_TYPE	Invalid datatype argument	36
MPI_ERR_UNKNOWN	Unknown error	37
MPI_ERR_UNSUPPORTED_DATAREP	Unsupported datarep passed to MPI_FILE_SET_VIEW	39
MPI_ERR_UNSUPPORTED_OPERATION	Unsupported operation, such as seeking on a file that supports sequential access only	41
MPI_ERR_VALUE_TOO_LARGE	Value is too large to store	42
MPI_ERR_WIN	Invalid window argument	43
MPI_ERR_LASTCODE	Last error code	44

9.5 Error Classes, Error Codes, and Error Handlers

Users may want to write a layered library on top of an existing MPI implementation, and this library may have its own set of error codes and classes. An example of such a library is an I/O library based on MPI, see Chapter 14. For this purpose, functions are needed to:

1. add a new error class and remove previously added user-defined error classes;
2. associate error codes with this error class, so that `MPI_ERROR_CLASS` works;
3. associate strings with these error codes, so that `MPI_ERROR_STRING` works;
4. remove such associations;
5. invoke the error handler associated with a communicator, window, file, or session object.

Several procedures are provided to do this. They are all local.

9.5.1 User-Defined Error Classes and Codes

The procedures that add and remove error classes, codes, or strings are thread-safe, as defined in Section 11.6. They are some of the few MPI procedures that may be called before MPI is initialized or after MPI is finalized, as defined in Section 11.4.1.

Advice to users. Note that despite the procedures being thread-safe, some concurrent calls can result in undefined behavior. Notably, the rules mandating that a call adding an error class/code/string must precede a call that removes that error class/code/string apply even when the procedures are called from different threads. Calling the procedures with different input values for the class/code parameters is always thread-safe. (*End of advice to users.*)

`MPI_ADD_ERROR_CLASS(errorclass)`

OUT errorclass value for the new error class (integer)

C binding

`int MPI_Add_error_class(int *errorclass)`

Fortran 2008 binding

`MPI_Add_error_class(errorclass, ierror)`
 INTEGER, INTENT(OUT) :: errorclass
 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

`MPI_ADD_ERROR_CLASS(ERRORCLASS, IERROR)`
 INTEGER ERRORCLASS, IERROR

Creates a new error class and returns the value for it.

Rationale. To avoid conflicts with existing error codes and classes, the value is set by the implementation and not by the user. (*End of rationale.*)

Advice to users. Since a call to `MPI_ADD_ERROR_CLASS` is local, the same `errorclass` may not be returned on all processes that make this call. Thus, it is not safe to assume that registering a new error on a set of processes at the same time will yield the same `errorclass` on all of the processes. Getting the “same” error on multiple processes may not cause the same value of error code to be generated. (*End of advice to users.*)

The value of `MPI_ERR_LASTCODE` is a constant value and is not affected by new user-defined error codes and classes. Instead, when using the World Model (Section 11.2), a predefined attribute key `MPI_LASTUSEDCLASS` is associated with `MPI_COMM_WORLD`. The attribute value corresponding to this key is the current maximum error class including the user-defined ones. This is a local value and may be different on different processes. The value returned by this key is always greater than or equal to `MPI_ERR_LASTCODE`.

Advice to users. The value returned by the key `MPI_LASTUSEDCLASS` will not change unless the user calls a procedure to explicitly add or remove an error class/code. In a multithreaded environment, the user must take extra care in assuming this value has not changed. Note that error codes and error classes are not necessarily dense. A user may not assume that each error class below `MPI_LASTUSEDCLASS` is valid. (*End of advice to users.*)

`MPI_REMOVE_ERROR_CLASS(errorclass)`

IN `errorclass` value for the error class to remove (integer)

C binding

```
int MPI_Remove_error_class(int errorclass)
```

Fortran 2008 binding

```
MPI_Remove_error_class(errorclass, ierror)
  INTEGER, INTENT(IN) :: errorclass
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_REMOVE_ERROR_CLASS(ERRORCLASS, IERROR)
  INTEGER ERRORCLASS, IERROR
```

Removes a user-created error class.

The value of the predefined attribute key `MPI_LASTUSEDCLASS` associated with `MPI_COMM_WORLD` is updated to reflect the maximum error class value. Note that there may be unused error classes that have a smaller value than `MPI_LASTUSEDCLASS`.

It is erroneous to call `MPI_REMOVE_ERROR_CLASS` with a value for `errorclass` that was not added by a call to `MPI_ADD_ERROR_CLASS`. Once an `errorclass` is removed by calling `MPI_REMOVE_ERROR_CLASS`, it is erroneous to remove it again without first obtaining the value from another call to `MPI_ADD_ERROR_CLASS`. It is erroneous to remove an error class when its associated error codes have not been removed before.

```

1 MPI_ADD_ERROR_CODE(errorclass, errorcode)
2     IN      errorclass          error class (integer)
3
4     OUT    errorcode           new error code to be associated with errorclass
5                                 (integer)
6

```

C binding

```
8 int MPI_Add_error_code(int errorclass, int *errorcode)
```

Fortran 2008 binding

```

10 MPI_Add_error_code(errorclass, errorcode, ierror)
11     INTEGER, INTENT(IN) :: errorclass
12     INTEGER, INTENT(OUT) :: errorcode
13     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14

```

Fortran binding

```

16 MPI_ADD_ERROR_CODE(ERRORCLASS, ERRORCODE, IERROR)
17     INTEGER ERRORCLASS, ERRORCODE, IERROR

```

Creates new error code associated with `errorclass` and returns its value in `errorcode`.

Rationale. To avoid conflicts with existing error codes and classes, the value of the new error code is set by the implementation and not by the user. (*End of rationale.*)

```

25 MPI_REMOVE_ERROR_CODE(errorcode)
26     IN      errorcode          error code to be removed (integer)
27

```

C binding

```
29 int MPI_Remove_error_code(int errorcode)
```

Fortran 2008 binding

```

32 MPI_Remove_error_code(errorcode, ierror)
33     INTEGER, INTENT(IN) :: errorcode
34     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

36 MPI_REMOVE_ERROR_CODE(ERRORCODE, IERROR)
37     INTEGER ERRORCODE, IERROR

```

Removes a user-created error code and all its associations with any error class.

It is erroneous to call `MPI_REMOVE_ERROR_CODE` with a value for `errorcode` that was not added by a call to `MPI_ADD_ERROR_CODE`. Once an `errorcode` is removed by calling `MPI_REMOVE_ERROR_CODE`, it is erroneous to remove it again without first obtaining the value from another call to `MPI_ADD_ERROR_CODE`. It is erroneous to remove an error code when its associated error string has not been removed before.

MPI_ADD_ERROR_STRING(errorcode, string) 1

IN errorcode error code or class (integer) 2

IN string text corresponding to errorcode (string) 3

C binding 5

int MPI_Add_error_string(int errorcode, const char *string) 6

Fortran 2008 binding 8

MPI_Add_error_string(errorcode, string, ierror) 9

INTEGER, INTENT(IN) :: errorcode 10

CHARACTER(LEN=*), INTENT(IN) :: string 11

INTEGER, OPTIONAL, INTENT(OUT) :: ierror 12

Fortran binding 13

MPI_ADD_ERROR_STRING(ERRORCODE, STRING, IERROR) 14

INTEGER ERRORCODE, IERROR 15

CHARACTER*(*) STRING 16

Associates a user-defined error string with an error code or class. The string must be no more than MPI_MAX_ERROR_STRING characters long. The length of the string is as defined in the calling language. The length of the string does not include the null terminator in C. Trailing blanks will be stripped in Fortran. Calling MPI_ADD_ERROR_STRING for an errorcode that already has a string will replace the old string with the new string. It is erroneous to call MPI_ADD_ERROR_STRING for an error code or class with a value \leq MPI_ERR_LASTCODE. 17-24

If MPI_ERROR_STRING is called when no string has been set, it will return a empty string (all spaces in Fortran, "" in C). 25-28

MPI_REMOVE_ERROR_STRING(errorcode) 29

IN errorcode error code or class (integer) 30

C binding 31

int MPI_Remove_error_string(int errorcode) 32

Fortran 2008 binding 33

MPI_Remove_error_string(errorcode, ierror) 34

INTEGER, INTENT(IN) :: errorcode 35

INTEGER, OPTIONAL, INTENT(OUT) :: ierror 36

Fortran binding 37

MPI_REMOVE_ERROR_STRING(ERRORCODE, IERROR) 38

INTEGER ERRORCODE, IERROR 39

Removes a user-defined association of an error string with an error code or class. 40

It is erroneous to call MPI_REMOVE_ERROR_STRING with a value for errorcode that does not have an error string added by a call to MPI_ADD_ERROR_STRING. 41-44

9.5.2 Calling Error Handlers

Section 9.3 describes the methods for creating and associating error handlers with communicators, files, windows, and sessions. Error handlers can be invoked implicitly when errors are raised during MPI operations, but can also be called by the user.

`MPI_COMM_CALL_ERRHANDLER(comm, errorcode)`

IN	comm	communicator with error handler (handle)
IN	errorcode	error code (integer)

C binding

```
int MPI_Comm_call_errhandler(MPI_Comm comm, int errorcode)
```

Fortran 2008 binding

```
MPI_Comm_call_errhandler(comm, errorcode, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: errorcode
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_CALL_ERRHANDLER(COMM, ERRORCODE, IERROR)
  INTEGER COMM, ERRORCODE, IERROR
```

This function invokes the error handler assigned to the communicator with the error code supplied. This function returns `MPI_SUCCESS` in C and the same value in `IERROR` if the error handler was successfully called (assuming the process is not aborted and the error handler returns).

`MPI_WIN_CALL_ERRHANDLER(win, errorcode)`

IN	win	window with error handler (handle)
IN	errorcode	error code (integer)

C binding

```
int MPI_Win_call_errhandler(MPI_Win win, int errorcode)
```

Fortran 2008 binding

```
MPI_Win_call_errhandler(win, errorcode, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  INTEGER, INTENT(IN) :: errorcode
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_CALL_ERRHANDLER(WIN, ERRORCODE, IERROR)
  INTEGER WIN, ERRORCODE, IERROR
```

This function invokes the error handler assigned to the window with the error code supplied. This function returns `MPI_SUCCESS` in C and the same value in `IERROR` if the error handler was successfully called (assuming the process is not aborted and the error handler returns).

Advice to users. In contrast to communicators, the error handler MPI_ERRORS_ARE_FATAL is associated with a window when it is created. (*End of advice to users.*)

MPI_FILE_CALL_ERRHANDLER(fh, errorcode)

IN	fh	file with error handler (handle)
IN	errorcode	error code (integer)

C binding

```
int MPI_File_call_errhandler(MPI_File fh, int errorcode)
```

Fortran 2008 binding

```
MPI_File_call_errhandler(fh, errorcode, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER, INTENT(IN) :: errorcode
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_CALL_ERRHANDLER(FH, ERRORCODE, IERROR)
  INTEGER FH, ERRORCODE, IERROR
```

This function invokes the error handler assigned to the file with the error code supplied. This function returns MPI_SUCCESS in C and the same value in IERROR if the error handler was successfully called (assuming the process is not aborted and the error handler returns).

Advice to users. The default error handler for files is MPI_ERRORS_RETURN. (*End of advice to users.*)

MPI_SESSION_CALL_ERRHANDLER(session, errorcode)

IN	session	session with error handler (handle)
IN	errorcode	error code (integer)

C binding

```
int MPI_Session_call_errhandler(MPI_Session session, int errorcode)
```

Fortran 2008 binding

```
MPI_Session_call_errhandler(session, errorcode, ierror)
  TYPE(MPI_Session), INTENT(IN) :: session
  INTEGER, INTENT(IN) :: errorcode
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SESSION_CALL_ERRHANDLER(SESSION, ERRORCODE, IERROR)
  INTEGER SESSION, ERRORCODE, IERROR
```

This function invokes the error handler assigned to the session with the error code supplied. This function returns MPI_SUCCESS in C and the same value in IERROR if the

error handler was successfully called (assuming the process is not aborted and the error handler returns).

Advice to users. Users are warned that handlers should not be called recursively with `MPI_COMM_CALL_ERRHANDLER`, `MPI_FILE_CALL_ERRHANDLER`, `MPI_WIN_CALL_ERRHANDLER`, or `MPI_SESSION_CALL_ERRHANDLER`. Doing this can create a situation where an infinite recursion is created. This can occur if `MPI_COMM_CALL_ERRHANDLER`, `MPI_FILE_CALL_ERRHANDLER`, `MPI_WIN_CALL_ERRHANDLER`, or `MPI_SESSION_CALL_ERRHANDLER` is called inside an error handler.

Error codes and classes are associated with a process. As a result, they may be used in any error handler. Error handlers should be prepared to deal with any error code they are given. Furthermore, it is good practice to only call an error handler with the appropriate error codes. For example, file errors would normally be sent to the file error handler. (*End of advice to users.*)

9.6 Timers and Synchronization

MPI defines a timer. A timer is specified even though it is not “message-passing,” because timing parallel programs is important in “performance debugging” and because existing timers (both in POSIX 1003.1-1988 and 1003.4D 14.1 and in Fortran 90) are either inconvenient or do not provide adequate access to high resolution timers.

`MPI_WTIME()`

C binding

`double MPI_Wtime(void)`

Fortran 2008 binding

`DOUBLE PRECISION MPI_Wtime()`

Fortran binding

`DOUBLE PRECISION MPI_WTIME()`

`MPI_WTIME` returns a floating-point number of seconds, representing elapsed wall-clock time since some time in the past. The “time in the past” is guaranteed not to change during the life of the process. The user is responsible for converting large numbers of seconds to other units if they are preferred. This function is portable (it returns seconds, not “ticks”), and it allows high-resolution. One would use it like this:

Example 9.4.

```
{
    double starttime, endtime;
    starttime = MPI_Wtime();
    ... stuff to be timed ...
    endtime = MPI_Wtime();
    printf("That took %f seconds\n", endtime-starttime);
}
```


The times returned are local to the node that called them. There is no requirement that different nodes return “the same time.” (But see also the discussion of `MPI_WTIME_IS_GLOBAL` in Section 9.1.2).

`MPI_WTICK()`

C binding

`double MPI_Wtick(void)`

Fortran 2008 binding

`DOUBLE PRECISION MPI_Wtick()`

Fortran binding

`DOUBLE PRECISION MPI_WTICK()`

`MPI_WTICK` returns the resolution of `MPI_WTIME` in seconds. That is, it returns, as a double precision value, the number of seconds between successive clock ticks. For example, if the clock is implemented by the hardware as a counter that is incremented every millisecond, the value returned by `MPI_WTICK` should be (10^{-3}) .

DRAFT

Chapter 10

The Info Object

Many of the procedures in MPI take an argument `info`. `info` is an opaque object with a handle of type `MPI_Info` in C and Fortran with the `mpi_f08` module, and `INTEGER` in Fortran with the `mpi` module or the (deprecated) `mpif.h` include file. It stores an unordered set of (`key,value`) pairs (both `key` and `value` are strings). A `key` can have only one `value`. MPI reserves several `keys` and requires that if an implementation uses a reserved `key`, it must provide the specified functionality. An implementation is not required to support these `keys` and may support any others not reserved by MPI.

Some `info` hints allow the MPI library to restrict its support for certain operations in order to improve performance or resource utilization. If an application provides such an `info` hint, it must be compatible with any changes in the behavior of the MPI library that are allowed by the `info` hint.

An implementation must support `info` objects as caches for arbitrary (`key,value`) pairs, regardless of whether it recognizes the `key`. Each procedure that takes hints in the form of an `MPI_Info` must be prepared to ignore any `key` it does not recognize. This description of `info` objects does not attempt to define how a particular procedure should react if it recognizes a `key` but not the associated `value`. `MPI_INFO_GET_NKEYS`, `MPI_INFO_GET_NTHKEY`, and `MPI_INFO_GET_STRING` must retain all (`key,value`) pairs so that layered functionality can also use the `Info` object.

`Keys` have an implementation-defined maximum length of `MPI_MAX_INFO_KEY`, which is at least 32 and at most 255. `Values` have an implementation-defined maximum length of `MPI_MAX_INFO_VAL`. In Fortran, leading and trailing spaces are stripped from both. Returned `values` will never be larger than these maximum lengths. Both `key` and `value` are case sensitive.

Rationale. `Keys` have a maximum length because the set of known `keys` will always be finite and known to the implementation and because there is no reason for `keys` to be complex. The small maximum size allows applications to declare `keys` of size `MPI_MAX_INFO_KEY`. The limitation on `value` sizes is so that an implementation is not forced to deal with arbitrarily long strings. (*End of rationale.*)

Advice to users. `MPI_MAX_INFO_VAL` might be very large, so it might not be wise to declare a string of that size. (*End of advice to users.*)

When `info` is used as an argument to any MPI procedure, it is interpreted before that procedure returns, so that it may be read, modified, or freed immediately after return. Changes to an `info` object after return from a procedure do not affect that interpretation.

Rationale. Prior to MPI-4.0, the above statement was restricted to nonblocking MPI procedures. For simplicity this restriction was removed, as it currently applies to

1 all MPI procedures that use info arguments. Note, this has to be revisited for new
 2 procedures added in the future, e.g., for future procedures that could return an info
 3 argument to be filled in after the return from the procedure. (*End of rationale.*)
 4

5 When the descriptions refer to a key or value as being a boolean, an integer, or a list,
 6 they mean the string representation of these types. An implementation may define its own
 7 rules for how info value strings are converted to other types, but to ensure portability, every
 8 implementation must support the following representations. Valid values for a boolean must
 9 include the strings "true" and "false" (all lowercase). For integers, valid values must include
 10 string representations of decimal values of integers that are within the range of a standard
 11 integer type in the program. (However it is possible that not every integer is a valid value
 12 for a given key.) On positive numbers, + signs are optional. No space may appear between
 13 a + or - sign and the leading digit of a number. For comma separated lists, the string
 14 must contain valid elements separated by commas. Leading and trailing spaces are stripped
 15 automatically from the types of info values described above and for each element of a comma
 16 separated list. These rules apply to all info values of these types. Implementations are free
 17 to specify a different interpretation for values of other info keys.
 18

19 MPI_INFO_CREATE(info)

20 OUT info info object created (handle)
 21

22
 23 **C binding**

24 int MPI_Info_create(MPI_Info *info)

25
 26 **Fortran 2008 binding**

27 MPI_Info_create(info, ierror)
 28 TYPE(MPI_Info), INTENT(OUT) :: info
 29 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

30 **Fortran binding**

31 MPI_INFO_CREATE(INFO, IERROR)
 32 INTEGER INFO, IERROR

33 MPI_INFO_CREATE creates a new info object. The newly created object contains no
 34 key/value pairs.
 35

36
 37 MPI_INFO_SET(info, key, value)

38 INOUT info info object (handle)
 39 IN key key (string)
 40 IN value value (string)
 41
 42

43 **C binding**

44 int MPI_Info_set(MPI_Info info, const char *key, const char *value)

45
 46 **Fortran 2008 binding**

47 MPI_Info_set(info, key, value, ierror)
 48 TYPE(MPI_Info), INTENT(IN) :: info

```

CHARACTER(LEN=*), INTENT(IN) :: key, value
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_INFO_SET(INFO, KEY, VALUE, IERROR)
  INTEGER INFO, IERROR
  CHARACTER*(*) KEY, VALUE

```

MPI_INFO_SET adds the (key,value) pair to info, and overrides the value if a value for the same key was previously set. key and value are null-terminated strings in C. In Fortran, leading and trailing spaces in key and value are stripped. If either key or value are longer than the respective maximum length, the call raises an error of class MPI_ERR_INFO_KEY or MPI_ERR_INFO_VALUE, respectively.

```

MPI_INFO_DELETE(info, key)

```

INOUT	info	info object (handle)
IN	key	key (string)

C binding

```

int MPI_Info_delete(MPI_Info info, const char *key)

```

Fortran 2008 binding

```

MPI_Info_delete(info, key, ierror)
  TYPE(MPI_Info), INTENT(IN) :: info
  CHARACTER(LEN=*), INTENT(IN) :: key
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_INFO_DELETE(INFO, KEY, IERROR)
  INTEGER INFO, IERROR
  CHARACTER*(*) KEY

```

MPI_INFO_DELETE deletes a (key,value) pair from info. If key is not defined in info, the call raises an error of class MPI_ERR_INFO_NOKEY.

```

MPI_INFO_GET_STRING(info, key, buflen, value, flag)

```

IN	info	info object (handle)
IN	key	key (string)
INOUT	buflen	length of buffer (integer)
OUT	value	value (string)
OUT	flag	true if key defined, false if not (logical)

C binding

```

int MPI_Info_get_string(MPI_Info info, const char *key, int *buflen,
  char *value, int *flag)

```

Fortran 2008 binding

```

1 MPI_Info_get_string(info, key, buflen, value, flag, ierror)
2   TYPE(MPI_Info), INTENT(IN) :: info
3   CHARACTER(LEN=*), INTENT(IN) :: key
4   INTEGER, INTENT(INOUT) :: buflen
5   CHARACTER(LEN=*), INTENT(OUT) :: value
6   LOGICAL, INTENT(OUT) :: flag
7   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

9 MPI_INFO_GET_STRING(INFO, KEY, BUFLen, VALUE, FLAG, IERROR)
10  INTEGER INFO, BUFLen, IERROR
11  CHARACTER*(*) KEY, VALUE
12  LOGICAL FLAG

```

This procedure retrieves the value associated with `key` from `info`, if any. If such a key exists in `info`, it sets `flag` to `true` and returns the value in `value`, otherwise it sets `flag` to `false` and leaves `value` unchanged. `buflen` on input is the size of the provided buffer, `value`, for the output of `buflen` it is the size of the buffer needed to store the value string. If the `buflen` passed into the procedure is less than the actual size needed to store the value string (including null terminator in C), the value is truncated. On return, the value of `buflen` will be set to the required buffer size to hold the value string. If `buflen` is set to 0, `value` is not changed. In C, `buflen` includes the required space for the null terminator. In C, this procedure returns a null terminated string in all cases where the `buflen` input value is greater than 0.

If `key` is larger than `MPI_MAX_INFO_KEY`, the call is erroneous.

Advice to users. The `MPI_INFO_GET_STRING` procedure can be used to obtain the size of the required buffer for a value string by setting the `buflen` to 0. The returned `buflen` can then be used to allocate memory before calling `MPI_INFO_GET_STRING` again to obtain the value string. (*End of advice to users.*)

```

33 MPI_INFO_GET_NKEYS(info, nkeys)

```

IN	info	info object (handle)
OUT	nkeys	number of defined keys (integer)

C binding

```

39 int MPI_Info_get_nkeys(MPI_Info info, int *nkeys)

```

Fortran 2008 binding

```

41 MPI_Info_get_nkeys(info, nkeys, ierror)
42   TYPE(MPI_Info), INTENT(IN) :: info
43   INTEGER, INTENT(OUT) :: nkeys
44   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

46 MPI_INFO_GET_NKEYS(INFO, NKEYS, IERROR)
47  INTEGER INFO, NKEYS, IERROR

```

MPI_INFO_GET_NKEYS returns the number of currently defined keys in info. 1

MPI_INFO_GET_NTHKEY(info, n, key) 2
3
4

IN	info	info object (handle) 5
IN	n	key number (integer) 6 7
OUT	key	key (string) 8 9

C binding 10

int MPI_Info_get_nthkey(MPI_Info info, int n, char *key) 11

Fortran 2008 binding 12

```
MPI_Info_get_nthkey(info, n, key, ierror) 13
  TYPE(MPI_Info), INTENT(IN) :: info 14
  INTEGER, INTENT(IN) :: n 15
  CHARACTER(LEN=*), INTENT(OUT) :: key 16
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror 17
18
```

Fortran binding 19

```
MPI_INFO_GET_NTHKEY(INFO, N, KEY, IERROR) 20
  INTEGER INFO, N, IERROR 21
  CHARACTER*(*) KEY 22
```

This procedure returns the *n*th defined key in *info*. Keys are numbered $0 \dots N - 1$ where *N* is the value returned by MPI_INFO_GET_NKEYS. All keys between 0 and $N - 1$ are guaranteed to be defined. The number of a given key does not change as long as *info* is not modified with MPI_INFO_SET or MPI_INFO_DELETE. 23
24
25
26
27

MPI_INFO_DUP(info, newinfo) 28
29

IN	info	info object (handle) 30
OUT	newinfo	info object created (handle) 31 32

C binding 33

int MPI_Info_dup(MPI_Info info, MPI_Info *newinfo) 34
35

Fortran 2008 binding 36

```
MPI_Info_dup(info, newinfo, ierror) 37
  TYPE(MPI_Info), INTENT(IN) :: info 38
  TYPE(MPI_Info), INTENT(OUT) :: newinfo 39
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror 40
41
```

Fortran binding 42

```
MPI_INFO_DUP(INFO, NEWINFO, IERROR) 43
  INTEGER INFO, NEWINFO, IERROR 44
```

MPI_INFO_DUP duplicates an existing info object, creating a new object, with the same (key,value) pairs and the same ordering of keys. 45
46

1 MPI_INFO_FREE(info)

2 INOUT info info object (handle)

4 **C binding**

5 int MPI_Info_free(MPI_Info *info)

7 **Fortran 2008 binding**

8 MPI_Info_free(info, ierror)

9 TYPE(MPI_Info), INTENT(INOUT) :: info

10 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

11 **Fortran binding**

12 MPI_INFO_FREE(INFO, IERROR)

13 INTEGER INFO, IERROR

14 This procedure frees `info` and sets it to `MPI_INFO_NULL`.

17 MPI_INFO_CREATE_ENV(info)

19 OUT info info object (handle)

21 **C binding**

22 int MPI_Info_create_env(int argc, char *argv[], MPI_Info *info)

24 **Fortran 2008 binding**

25 MPI_Info_create_env(info, ierror)

26 TYPE(MPI_Info), INTENT(OUT) :: info

27 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

28 **Fortran binding**

29 MPI_INFO_CREATE_ENV(INFO, IERROR)

30 INTEGER INFO, IERROR

31 This procedure creates an output object `info` with the same construction as
 32 `MPI_INFO_ENV` as created during `MPI_INIT` or `MPI_INIT_THREAD` when the same argu-
 33 ments are used. This construction is described in Section 11.2.1; however, this procedure
 34 can be called when not using the World Model, e.g., when using the Sessions Model. This
 35 object is not a direct copy or alias of the `MPI_INFO_ENV` object and could contain different
 36 values based on the input arguments and other sources. Multiple calls to this procedure
 37 that are given the same input arguments will produce `info` objects consistent with the defi-
 38 nition of `MPI_INFO_ENV`. The version for ISO C accepts the `argc` and `argv` that are provided
 39 by the arguments to `main` or `0` for `argc` and `NULL` for `argv`. The user is responsible for freeing
 40 the `info` object via `MPI_INFO_FREE`. This procedure is local.

41 This procedure must always be thread-safe, as defined in Section 11.6. It is one of the
 42 few procedures that may be called before MPI is initialized or after MPI is finalized.

44 *Advice to users.*

45 In some circumstances (e.g., when passing `0` to `argc` and `NULL` to `argv` in C or in Fortran
 46 where such arguments do not exist), the `info` object may not be populated or may be
 47 populated incompletely because this procedure is local and the implementation may
 48

not be able to determine the correct values. Note that this could result in different values in the resulting info object at different MPI processes.

(End of advice to users.)

DRAFT

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

DRAFT

Chapter 11

Process Initialization, Creation, and Management

11.1 Introduction

MPI is primarily concerned with communication rather than process or resource management. However, it is necessary to address these issues to some degree in order to define a useful framework for communication. This chapter presents a set of MPI interfaces that allows for several approaches to MPI initialization and process management while placing minimal restrictions on the execution environment.

One goal of MPI is to achieve *source code portability*. By this we mean that a program written using MPI and complying with the relevant language standards is portable as written, and must not require any source code changes when moved from one system to another. This explicitly does *not* say anything about how an MPI program is started or launched from the command line, nor what the user must do to set up the environment in which an MPI program will run. However, an implementation may require some setup or initialization procedure to be performed before the complete set of MPI routines may be called.

To this end, MPI presents two models for **MPI process initialization**. In the World Model, an initial set of processes is created that are related by their membership in a common MPI_COMM_WORLD (see Section 11.2) communicator. In the Sessions Model (Section 11.3), an initial set of processes is also created, but the application must explicitly manage the creation of MPI groups, and hence MPI communicators. MPI_COMM_WORLD is only valid for use as a communicator in the World Model, i.e., after a successful call to MPI_INIT or MPI_INIT_THREAD and before a call to MPI_FINALIZE. An application can employ both of these Process Models concurrently. In multi-component MPI applications, for example, a component such as a library can make use of the Sessions Model to instantiate MPI resources without impacting the rest of the application.

The Dynamic Process Model (see Section 11.7), provides for the creation and management of additional processes after an MPI application has been started. A major impetus for the Dynamic Process Model comes from the PVM [26] research effort. This work has provided a wealth of experience with process management and resource control that illustrates their benefits and potential pitfalls.

In developing the Dynamic Process Model, the MPI Forum decided not to address resource control because it was not able to design a portable interface that would be appropriate for the broad spectrum of existing and potential resource and process controllers. MPI assumes that resource control is provided externally.

Process management functionality is included in MPI to enable its use in classes of message-passing applications requiring process control. These include task farms, serial

1 applications with parallel modules, and problems that require a run-time assessment of the
 2 number and type of processes that should be started.

3 The following goals are central to the design of MPI process management:

- 4 ● The MPI process model must apply to the vast majority of current parallel environ-
 5 ments.
- 6 ● MPI must not take over operating system responsibilities. It should instead provide a
 7 clean interface between an application and system software.
- 8 ● MPI must guarantee communication determinism in the presence of dynamic pro-
 9 cesses, i.e., dynamic process management must not introduce unavoidable race condi-
 10 tions.
- 11 ● MPI must not contain features that compromise performance.
- 12
- 13
- 14

15 The Dynamic Process Model addresses these issues in two ways. First, MPI remains
 16 primarily a communication library. It does not manage the parallel environment in which
 17 a parallel program executes, though it provides a minimal interface between an application
 18 and external resource and process managers.

19 Second, MPI maintains a consistent concept of a communicator, regardless of how its
 20 members came into existence. A communicator is never changed once created, and it is
 21 always created using deterministic collective operations.

22 11.2 The World Model

23 11.2.1 Starting MPI Processes

24 When using the World Model, MPI is initialized by calling either `MPI_INIT` or
 25 `MPI_INIT_THREAD`.

26 `MPI_INIT()`

27 **C binding**

28 `int MPI_Init(int *argc, char ***argv)`

29 **Fortran 2008 binding**

30 `MPI_Init(ierror)`
 31 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

32 **Fortran binding**

33 `MPI_INIT(IERROR)`
 34 `INTEGER IERROR`

35 In the World Model, an MPI program must contain exactly one call to an MPI ini-
 36 tialization routine: `MPI_INIT` or `MPI_INIT_THREAD`. `MPI_COMM_WORLD` and
 37 `MPI_COMM_SELF` are not valid for use as communicators prior to invocation of `MPI_INIT` or
 38 `MPI_INIT_THREAD`. Subsequent calls to either of these initialization routines are erroneous.
 39 A subset of MPI functions may be invoked before MPI initialization routines are called, see
 40 Section 11.4.1. The procedures `MPI_INIT` and `MPI_INIT_THREAD` accept either the `argc`
 41 and `argv` that are provided by the arguments to `main` or `NULL`.
 42
 43
 44
 45
 46
 47
 48

Example 11.1. Initializing MPI using MPI_INIT

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
int main(int argc, char *argv[])
{
    MPI_Init(&argc, &argv);

    /* parse arguments */
    /* main program */

    MPI_Finalize();    /* see below */
    return 0;
}

```

The Fortran version takes only IERROR.

Conforming implementations of MPI are required to allow applications to pass NULL for both the `argc` and `argv` arguments of `main` in C.

Failures may disrupt the execution of the program before or during MPI initialization. A high-quality implementation shall not deadlock during MPI initialization, even in the presence of failures. Except for functions with the `MPI_T_` prefix, failures in MPI operations prior to or during MPI initialization are reported by invoking the initial error handler. Users can use the "mpi_initial_errhandler" info key during the launch of MPI processes (e.g., `MPI_COMM_SPAWN` / `MPI_COMM_SPAWN_MULTIPLE`, or `mpixec`) to set a nonfatal initial error handler before MPI initialization. When the initial error handler is set to `MPI_ERRORS_ABORT`, raising an error before or during initialization aborts the local MPI process (i.e., it is similar to calling `MPI_ABORT` on `MPI_COMM_SELF`). An implementation may not always be capable of determining, before MPI initialization, what constitutes the local MPI process, or the set of connected processes. In this case, errors before initialization may cause a different set of MPI processes to abort than specified. During MPI initialization, the initial error handler is associated with `MPI_COMM_WORLD`, `MPI_COMM_SELF`, and the communicator returned by `MPI_COMM_GET_PARENT` (if any).

Advice to implementors. Some failures may leave MPI in an undefined state, or raise an error before the error handling capabilities are fully operational, in which cases the implementation may be incapable of providing the desired error handling behavior. Of note, in some implementations, the notion of an MPI process is not clearly established in the early stages of MPI initialization (for example, when the implementation considers threads that called `MPI_INIT` as independent MPI processes); in this case, before MPI is initialized, the `MPI_ERRORS_ABORT` error handler may abort what would have become multiple MPI processes.

When a failure occurs during MPI initialization, the implementation may decide to return `MPI_SUCCESS` from the MPI initialization function instead of raising an error. It is recommended that an implementation masks an initialization error only when it expects that later MPI calls will result in well-specified behavior (i.e., barring additional failures, either the outcome of any call will be correct, or the call will raise an appropriate error). For example, it may be difficult for an implementation to avoid unspecified behavior when the group of `MPI_COMM_WORLD` does not contain the same set of MPI processes at all members of the communicator, or if the communicator returned from `MPI_COMM_GET_PARENT` was not initialized correctly. (*End of advice to implementors.*)

After MPI is initialized, the application can access information about the execution environment by querying the predefined info object `MPI_INFO_ENV`. The following keys are predefined for this object, corresponding to the arguments of `MPI_COMM_SPAWN` or of `mpiexec`:

"command": Name of program executed.

"argv": Space separated arguments to command.

"maxprocs": Maximum number of MPI processes to start.

"mpi_initial_errhandler": Name of the initial errhandler.

"mpi_memory_alloc_kinds": Memory allocation kinds supported by the MPI library (see Section 11.4.3).

"soft": Allowed values for number of processors.

"host": Hostname.

"arch": Architecture name.

"wdir": Working directory of the MPI process.

"file": Value is the name of a file in which additional information is specified.

"thread_level": Requested level of thread support, if requested before the program started execution.

Note that all values are strings. Thus, the maximum number of processes is represented by a string such as "1024" and the requested level is represented by a string such as "MPI_THREAD_SINGLE".

Advice to users. If one of the "argv" arguments contains a space, there is no way to tell from the value of the "argv" info key whether a space is part of the argument or is separating different arguments. (*End of advice to users.*)

The info object `MPI_INFO_ENV` need not contain a (key,value) pair for each of these predefined keys; the set of (key,value) pairs provided is implementation-dependent. Implementations may provide additional, implementation specific, (key,value) pairs.

In cases where the MPI processes were started with `MPI_COMM_SPAWN_MULTIPLE` or, equivalently, with a startup mechanism that supports multiple process specifications, then the values stored in the info object `MPI_INFO_ENV` at a process are those values that affect the local MPI process.

Example 11.2. If MPI is started with a call to

```
mpiexec -n 5 -arch x86_64 ocean : -n 10 -arch power9 atmos
```

Then the first 5 processes will have in their `MPI_INFO_ENV` object the pairs (command, ocean), (maxprocs, 5), and (arch, x86_64). The next 10 processes will have in `MPI_INFO_ENV` (command, atmos), (maxprocs, 10), and (arch, power9)

Advice to users. The values passed in `MPI_INFO_ENV` are the values of the arguments passed to the mechanism that started the MPI execution—not the actual value provided. Thus, the value associated with "maxprocs" is the number of MPI processes requested; it can be larger than the actual number of processes obtained, if the `soft` option was used. (*End of advice to users.*)

Advice to implementors. High-quality implementations will provide a (key,value) pair for each parameter that can be passed to the command that starts an MPI program. (*End of advice to implementors.*)

The following function may be used to initialize MPI, and to initialize the MPI thread environment, instead of `MPI_INIT`.

`MPI_INIT_THREAD(required, provided)`

IN	required	desired level of thread support (integer)
OUT	provided	provided level of thread support (integer)

C binding

`int MPI_Init_thread(int *argc, char ***argv, int required, int *provided)`

Fortran 2008 binding

`MPI_Init_thread(required, provided, ierror)`
 INTEGER, INTENT(IN) :: required
 INTEGER, INTENT(OUT) :: provided
 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

`MPI_INIT_THREAD(REQUIRED, PROVIDED, IERROR)`
 INTEGER REQUIRED, PROVIDED, IERROR

This call initializes MPI in the same way that a call to `MPI_INIT` would. In addition, it initializes the thread environment. The argument `required` is used to specify the desired level of thread support. The possible values are listed in increasing order of thread support.

MPI_THREAD_SINGLE: Only one thread will execute.

MPI_THREAD_FUNNELED: The process may be multithreaded, but the application must ensure that only the main thread makes MPI calls (for the definition of main thread, see `MPI_IS_THREAD_MAIN` on page 487).

MPI_THREAD_SERIALIZED: The process may be multithreaded, and multiple threads may make MPI calls, but only one at a time: MPI calls are not made concurrently from two distinct threads (all MPI calls are "serialized").

MPI_THREAD_MULTIPLE: Multiple threads may call MPI, with no restrictions.

These values are monotonic; i.e., `MPI_THREAD_SINGLE < MPI_THREAD_FUNNELED < MPI_THREAD_SERIALIZED < MPI_THREAD_MULTIPLE`.

Different processes in `MPI_COMM_WORLD` may require different levels of thread support.

1 The call returns in provided information about the actual level of thread support that
2 will be provided by MPI. It can be one of the four values listed above.

3 The level(s) of thread support that can be provided by MPI_INIT_THREAD will depend
4 on the implementation, and may depend on information provided by the user before the
5 program started to execute (e.g., with arguments to mpiexec). If possible, the call will
6 return provided = required. Failing this, the call will return the least supported level such
7 that provided > required (thus providing a stronger level of support than required by the
8 user). Finally, if the user requirement cannot be satisfied, then the call will return in
9 provided the highest supported level.

10 A **thread compliant** MPI implementation will be able to return provided =
11 MPI_THREAD_MULTIPLE. Such an implementation may always return provided =
12 MPI_THREAD_MULTIPLE, irrespective of the value of required.

13 An MPI library that is not thread compliant must always return provided =
14 MPI_THREAD_SINGLE, even if MPI_INIT_THREAD is called on a multithreaded process.
15 The library should also return correct values for the MPI calls that can be executed before
16 initialization, even if multiple threads have been spawned.

17
18 *Rationale.* Such code is erroneous, but if the MPI initialization is performed by a
19 library, the error cannot be detected until MPI_INIT_THREAD is called. The require-
20 ments in the previous paragraph ensure that the error can be properly detected. (*End*
21 *of rationale.*)

22
23 A call to MPI_INIT has the same effect as a call to MPI_INIT_THREAD with a required
24 = MPI_THREAD_SINGLE.

25 Vendors may provide (implementation dependent) means to specify the level(s) of
26 thread support available when the MPI program is started, e.g., with arguments to mpiexec.
27 This will affect the outcome of calls to MPI_INIT and MPI_INIT_THREAD. Suppose, for ex-
28 ample, that an MPI program has been started so that only MPI_THREAD_MULTIPLE is avail-
29 able. Then MPI_INIT_THREAD will return provided = MPI_THREAD_MULTIPLE, irrespective
30 of the value of required; a call to MPI_INIT will also initialize the MPI thread support level
31 to MPI_THREAD_MULTIPLE. Suppose, instead, that an MPI program has been started so
32 that all four levels of thread support are available. Then, a call to MPI_INIT_THREAD will
33 return provided = required; alternatively, a call to MPI_INIT will initialize the MPI thread
34 support level to MPI_THREAD_SINGLE.

35
36 *Rationale.* Various optimizations are possible when MPI code is executed single-
37 threaded, or is executed on multiple threads, but not concurrently: mutual exclusion
38 code may be omitted. Furthermore, if only one thread executes, then the MPI library
39 can use library functions that are not thread safe, without risking conflicts with user
40 threads. Also, the model of one communication thread, multiple computation threads
41 fits many applications well, e.g., if the process code is a sequential Fortran/C program
42 with MPI calls that has been parallelized by a compiler for execution on an SMP node,
43 in a cluster of SMPs, then the process computation is multithreaded, but MPI calls
44 will likely execute on a single thread.

45 The design accommodates a static specification of the thread support level, for en-
46 vironments that require static binding of libraries, and for compatibility for current
47 multithreaded MPI codes. (*End of rationale.*)

Advice to implementors. If `provided` is not `MPI_THREAD_SINGLE` then the MPI library should not invoke C or Fortran library calls that are not thread safe, e.g., in an environment where `malloc` is not thread safe, then `malloc` should not be used by the MPI library.

Some implementors may want to use different MPI libraries for different levels of thread support. They can do so using dynamic linking and selecting which library will be linked when `MPI_INIT_THREAD` is invoked. If this is not possible, then optimizations for lower levels of thread support will occur only when the level of thread support required is specified at link time.

Note that `required` need not be the same value on all processes of `MPI_COMM_WORLD`. (*End of advice to implementors.*)

As with `MPI_INIT`, discussed in Section 11.2.1, the version for ISO C accepts the `argc` and `argv` that are provided by the arguments to `main` or `NULL` for both arguments.

The following function can be used to query the current level of thread support.

`MPI_QUERY_THREAD(provided)`

OUT `provided` provided level of thread support (integer)

C binding

`int MPI_Query_thread(int *provided)`

Fortran 2008 binding

`MPI_Query_thread(provided, ierror)`
 INTEGER, INTENT(OUT) :: `provided`
 INTEGER, OPTIONAL, INTENT(OUT) :: `ierror`

Fortran binding

`MPI_QUERY_THREAD(PROVIDED, IERROR)`
 INTEGER PROVIDED, IERROR

The call returns in `provided` the current level of thread support, which will be the value returned in `provided` by `MPI_INIT_THREAD`, if MPI was initialized by a call to `MPI_INIT_THREAD`. This function is only applicable when using the World Model to initialize MPI. In the case of applications using both the World Model and the Sessions Model, this function only returns the thread support level returned in `provided` by `MPI_INIT_THREAD`.

`MPI_IS_THREAD_MAIN(flag)`

OUT `flag` true if calling thread is main thread, false otherwise
 (logical)

C binding

`int MPI_Is_thread_main(int *flag)`

Fortran 2008 binding

`MPI_Is_thread_main(flag, ierror)`
 LOGICAL, INTENT(OUT) :: `flag`

1 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

2 **Fortran binding**

3 MPI_IS_THREAD_MAIN(FLAG, IERROR)

4 LOGICAL FLAG

5 INTEGER IERROR

6
7 This function can be called by a thread to determine if it is the main thread (the thread
8 that called MPI_INIT or MPI_INIT_THREAD). This function is only applicable when using
9 the World Model to initialize MPI. In the case of applications using both the World Model
10 and the Sessions Model, the behavior of this procedure is the same as if the application
11 were only using the World Model.

12 All routines listed in this section must be supported by all MPI implementations.

13
14 *Rationale.* MPI libraries are required to provide these calls even if they do not
15 support threads, so that portable code that contains invocations to these functions
16 can link correctly. MPI_INIT continues to be supported so as to provide compatibility
17 with current MPI codes. (*End of rationale.*)

18
19 *Advice to users.* It is possible to spawn threads before MPI is initialized, but
20 MPI_COMM_WORLD and MPI_COMM_SELF cannot be used until the World Model is
21 active, i.e., until MPI_INIT_THREAD is invoked by one thread (which, thereby, be-
22 comes the main thread). In particular, it is possible to enter the MPI execution with
23 a multithreaded process.

24 In the World Model, the level of thread support provided is a global property of the
25 MPI process that can be specified only once, when MPI is initialized on that process (or
26 before). Portable third party libraries have to be written so as to accommodate any
27 provided level of thread support. Otherwise, their usage will be restricted to specific
28 level(s) of thread support. If such a library can run only with specific level(s) of thread
29 support, e.g., only with MPI_THREAD_MULTIPLE, then MPI_QUERY_THREAD can be
30 used to check whether the user initialized MPI to the correct level of thread support.
31 (*End of advice to users.*)

32
33 11.2.2 Finalizing MPI

34
35
36 MPI_FINALIZE()

37
38 **C binding**

39 int MPI_Finalize(void)

40 **Fortran 2008 binding**

41 MPI_Finalize(ierror)

42 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

43 **Fortran binding**

44 MPI_FINALIZE(IERROR)

45 INTEGER IERROR

46
47 This routine cleans up all MPI state associated with the World Model. If an MPI
48 program that initializes the World Model terminates normally (i.e., not due to a call to

MPI_ABORT or an unrecoverable error) then each process must call MPI_FINALIZE before it exits.

Before an MPI process invokes MPI_FINALIZE, the process must perform all MPI calls needed to complete its involvement in MPI communications associated with the World Model. It must locally complete all MPI operations that it initiated and must execute matching calls needed to complete MPI communications initiated by other processes. For example, if the process executed a nonblocking send, it must eventually call MPI_WAIT, MPI_TEST, MPI_REQUEST_FREE, or any derived function; if the process is the target of a send, then it must post the matching receive; if it is part of a group executing a collective operation, then it must have completed its participation in the operation. This means that before calling MPI_FINALIZE, all message handles associated with the World Model must be received (with MPI_MRECV or derived procedures) and all request handles associated with the World Model must be freed in the case of nonblocking operations, and must be inactive or freed in the case of persistent operations (i.e., by calling one of the procedures MPI_{TEST|WAIT}{|ANY|SOME|ALL} or MPI_REQUEST_FREE).

The call to MPI_FINALIZE does not clean up MPI state associated with objects created using MPI_SESSION_INIT and other Sessions Model methods, nor objects created using the communicator returned by MPI_COMM_GET_PARENT. See Sections 11.3 and 11.8.

The call to MPI_FINALIZE does not free objects created by MPI calls; these objects are freed using MPI_XXX_FREE, MPI_COMM_DISCONNECT, or MPI_FILE_CLOSE calls.

Once MPI_FINALIZE returns, no MPI procedure may be called in the World Model (not even MPI_INIT, or freeing objects created within the World Model), except for those listed in Section 11.4.1.

MPI_FINALIZE is collective over all connected processes. If no processes were spawned, accepted or connected then this means over MPI_COMM_WORLD; otherwise it is collective over the union of all processes that have been and continue to be connected, as explained in Section 11.10.4.

The following examples illustrate these rules.

Example 11.3. The following code is correct

Process 0	Process 1
MPI_Init();	MPI_Init();
MPI_Send(dest=1);	MPI_Recv(src=0);
MPI_Finalize();	MPI_Finalize();

Example 11.4. Without a matching receive, the program is erroneous

Process 0	Process 1
MPI_Init();	MPI_Init();
MPI_Send(dest=1);	
MPI_Finalize();	MPI_Finalize();

Example 11.5. This program is correct: Process 0 calls MPI_Finalize after it has executed the MPI calls that complete the send operation. Likewise, process 1 executes the MPI call that completes the matching receive operation before it calls MPI_Finalize.

Process 0	Process 1

```

1      MPI_Init();
2      MPI_Isend(dest=1);
3      MPI_Request_free();
4      MPI_Finalize();
5      exit();
6
7      MPI_Init();
8      MPI_Recv(src=0);
9      MPI_Finalize();
10     exit();

```

Example 11.6. This program is correct. The attached buffer is a resource allocated by the user, not by MPI; it is available to the user after MPI is finalized.

Process 0	Process 1
MPI_Init();	MPI_Init();
buffer = malloc(1000000);	MPI_Recv(src=0);
MPI_Buffer_attach();	MPI_Finalize();
MPI_Send(dest=1);	exit();
MPI_Finalize();	
free(buffer);	
exit();	

Example 11.7. This program is correct. The cancel operation must succeed, since the send cannot complete normally. The wait operation, after the call to MPI_Cancel, is local—no matching MPI call is required on process 1. Cancelling a send request by calling MPI_CANCEL is deprecated.

Process 0	Process 1
MPI_Issend(dest=1);	MPI_Finalize();
MPI_Cancel();	
MPI_Wait();	
MPI_Finalize();	

Advice to implementors. Even though a process has executed all MPI calls needed to complete the communications it is involved with, such communication may not yet be completed from the viewpoint of the underlying MPI system. For example, a blocking send may have returned, even though the data is still buffered at the sender in an MPI buffer; an MPI process may receive a cancel request for a message it has completed receiving. The MPI implementation must ensure that a process has completed any involvement in MPI communication before MPI_FINALIZE returns. Thus, if a process exits after the call to MPI_FINALIZE, this will not cause an ongoing communication to fail. The MPI implementation should also complete freeing all objects marked for deletion by MPI calls that freed them. See also Section 2.9 on *progress*. (*End of advice to implementors.*)

Failures may disrupt MPI operations during and after MPI finalization. A high quality implementation shall not deadlock in MPI finalization, even in the presence of failures. The normal rules for MPI error handling continue to apply. After MPI_COMM_SELF has been “freed” (see Section 11.2.4), errors that are not associated with a communicator, window, or file raise the initial error handler (set during the launch operation, see 11.8.4).

Although it is not required that all processes return from MPI_FINALIZE, it is required that, when it has not failed or aborted, at least the MPI process that was assigned rank 0 in MPI_COMM_WORLD returns, so that users can know that the MPI portion of the computation

is over. In addition, in a POSIX environment, users may desire to supply an exit code for each process that returns from `MPI_FINALIZE`.

Note that a failure may terminate the MPI process that was assigned rank 0 in `MPI_COMM_WORLD`, in which case it is possible that no MPI process returns from `MPI_FINALIZE`.

Advice to users. Applications that handle errors are encouraged to implement all rank-specific code before the call to `MPI_FINALIZE`. In Example 11.8, the process with rank 0 in `MPI_COMM_WORLD` may have been terminated before, during, or after the call to `MPI_FINALIZE`, possibly leading to the code after `MPI_FINALIZE` never being executed. (*End of advice to users.*)

Example 11.8. The following illustrates the use of requiring that at least one process return and that it be known that process 0 is one of the processes that return. One wants code like the following to work no matter how many processes return.

```
...
MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
...
MPI_Finalize();
if (myrank == 0) {
    resultfile = fopen("outfile", "w");
    dump_results(resultfile);
    fclose(resultfile);
}
exit(0);
```

11.2.3 Determining Whether MPI Has Been Initialized When Using the World Model

One of the goals of MPI is to allow for layered libraries. A library using the World Model needs to know if MPI has been initialized using either of `MPI_INIT` or `MPI_INIT_THREAD`. In MPI the function `MPI_INITIALIZED` is provided to tell if MPI had been initialized using the World Model. In the World Model, once MPI has been finalized it cannot be restarted. A library needs to be able to determine this to act accordingly. To achieve this, the function `MPI_FINALIZED` is needed.

`MPI_INITIALIZED(flag)`

OUT flag

Flag is true if `MPI_INIT` or `MPI_INIT_THREAD` has been called and false otherwise (logical)

C binding

`int MPI_Initialized(int *flag)`

Fortran 2008 binding

`MPI_Initialized(flag, ierror)`

LOGICAL, INTENT(OUT) :: flag

INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

```
MPI_INITIALIZED(FLAG, IERROR)
```

```
    LOGICAL FLAG
```

```
    INTEGER IERROR
```

This routine may be used to determine whether `MPI_INIT` or `MPI_INIT_THREAD` has been called. `MPI_INITIALIZED` returns true if the calling process has called either of these MPI procedures. Whether `MPI_FINALIZE` has been called does not affect the behavior of `MPI_INITIALIZED`. This function must always be thread-safe, as defined in Section 11.6. This function returns false for applications using the Sessions Model exclusively.

```
MPI_FINALIZED(flag)
```

```
    OUT    flag
```

true if `MPI_FINALIZE` has been called and false otherwise. (logical)

C binding

```
int MPI_Finalized(int *flag)
```

Fortran 2008 binding

```
MPI_Finalized(flag, ierror)
```

```
    LOGICAL, INTENT(OUT) :: flag
```

```
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FINALIZED(FLAG, IERROR)
```

```
    LOGICAL FLAG
```

```
    INTEGER IERROR
```

This routine returns true if `MPI_FINALIZE` has completed. It is valid to call `MPI_FINALIZED` before `MPI_INIT` and after `MPI_FINALIZE`. This function must always be thread-safe, as defined in Section 11.6.

11.2.4 Allowing User Functions at MPI Finalization

In the context of the World Model, there are times in which it would be convenient to have actions happen when an MPI process finalizes MPI. For example, a routine may do initializations that are useful until the MPI job (or that part of the job that is being terminated in the case of dynamically created processes) finalizes MPI. This can be accomplished in MPI by attaching an attribute to `MPI_COMM_SELF` with a callback function. When `MPI_FINALIZE` is called, it will first execute the equivalent of an `MPI_COMM_FREE` on `MPI_COMM_SELF`. This will cause the delete callback function to be executed on all keys associated with `MPI_COMM_SELF`, in the reverse order that they were set on `MPI_COMM_SELF`. If no key has been attached to `MPI_COMM_SELF`, then no callback is invoked. The “freeing” of `MPI_COMM_SELF` occurs before any other parts of MPI are affected. Thus, for example, calling `MPI_FINALIZED` will return false in any of these callback functions. Once done with `MPI_COMM_SELF`, the order and rest of the actions taken by `MPI_FINALIZE` is not specified.

Advice to implementors. Since attributes can be added from any supported language, the MPI implementation needs to remember the creating language so the correct

callback is made. Implementations that use the attribute delete callback on MPI_COMM_SELF internally should register their internal callbacks before returning from MPI_INIT / MPI_INIT_THREAD, so that libraries or applications will not have portions of the MPI implementation shut down before the application-level callbacks are made. (*End of advice to implementors.*)

11.3 The Sessions Model

There are a number of limitations with the World Model described in the preceding section. Among these are the following: MPI cannot be initialized from different application components without *a priori* knowledge or coordination; MPI cannot be initialized more than once; and MPI cannot be reinitialized after MPI_FINALIZE has been called. This section describes an alternative approach to MPI initialization—the Sessions Model. With this approach, an MPI application, or components of the application, can instantiate MPI resources for the specific communication needs of this component. MPI_COMM_WORLD is not valid for use as a communicator. MPI_INFO_ENV is not valid for use as an info object when only using the Sessions Model. As described in Section 11.2.1, MPI must be initialized using the World Model to use this info object. Note that an application may employ both the Sessions Model and World Model concurrently (see Section 11.1).


In the Sessions Model, MPI resources can be allocated and freed multiple times in an MPI process.

As shown in Figure 11.1, when using the Sessions Model, an MPI process instantiates an **MPI Session handle**, which can be used to query the runtime system about characteristics of the job within which the process is running, as well as other system resources. Using this information, the MPI process can then create an MPI Group based on application requirements and available resources, which in turn can be used to create an MPI Communicator, Window, or File. By judicious creation of communicators, an application only needs to allocate MPI resources based on its communication requirements. Although there are existing MPI interfaces for creating communicators that can, in principle, allow for resource optimizations within an MPI implementation, this can only be done following initialization of MPI.

For multithreaded applications, the Sessions Model provides fine-grain control of the thread support level for MPI objects. It is possible to specify different thread support levels when creating different *MPI Session handles*. Thus different components of an application can use different thread support levels.

The Sessions Model introduces a concept of isolation. MPI objects derived from different *MPI Session handles* shall not be intermixed with each other in a single MPI procedure call. MPI objects derived from the Sessions Model shall not be intermixed in a single MPI procedure call with MPI objects derived from the World Model. MPI objects derived from the Sessions Model shall not be intermixed in a single MPI procedure call with MPI objects derived from the communicator obtained from a call to MPI_COMM_GET_PARENT or MPI_COMM_JOIN.

This restriction does not apply to generalized requests (Section 13.2) as such requests are not associated directly with communicators or other MPI objects. Note however, the Sessions Model does not otherwise change the semantics or behavior of MPI objects.



26 Figure 11.1: Steps to creating an MPI Communicator from an MPI Session handle.

27
28
29 11.3.1 Session Creation and Destruction Methods

30
31
32 `MPI_SESSION_INIT(info, errhandler, session)`

33 IN info info object to specify thread support level and MPI
34 implementation specific resources (handle)
35 IN errhandler error handler to invoke in the event that an error is
36 encountered during this function call (handle)
37
38 OUT session new session (handle)
39

40 **C binding**

41 `int MPI_Session_init(MPI_Info info, MPI_Errhandler errhandler,`
42 `MPI_Session *session)`

43 **Fortran 2008 binding**

44 `MPI_Session_init(info, errhandler, session, ierror)`
45 `TYPE(MPI_Info), INTENT(IN) :: info`
46 `TYPE(MPI_Errhandler), INTENT(IN) :: errhandler`
47 `TYPE(MPI_Session), INTENT(OUT) :: session`
48

INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

MPI_SESSION_INIT(INFO, ERRHANDLER, SESSION, IERROR)
 INTEGER INFO, ERRHANDLER, SESSION, IERROR

The info argument is used to request MPI functionality requirements and possible MPI implementation specific capabilities. The following info keys are predefined:

"thread_level" used to request the thread support level required for MPI objects derived from the Session. Allowed values are "MPI_THREAD_SINGLE", "MPI_THREAD_FUNNELED", "MPI_THREAD_SERIALIZED", and "MPI_THREAD_MULTIPLE". Note that the thread support value is specified by a string rather than the integer values supplied to MPI_INIT_THREAD. The thread support level actually provided by the MPI implementation can be determined via a subsequent call to MPI_SESSION_GET_INFO to return the info object associated with the Session. The default thread support level is MPI implementation dependent.

"mpi_memory_alloc_kinds" used to request support for memory allocation kinds to be used by the calling MPI process on MPI objects derived from the Session. See Section 11.4.3. A value for this info key can also be supplied as an argument to an MPI startup mechanism as described in Section 11.5.

The errhandler argument specifies an error handler to invoke in the event that the Session instantiation call encounters an error. The error handler shall be either a pre-defined error handler (see 9.3) or one created using MPI_SESSION_CREATE_ERRHANDLER. Session instantiation is intended to be a lightweight operation. An MPI process may instantiate multiple Sessions. MPI_SESSION_INIT is always thread safe; multiple threads within an application may invoke it concurrently.

Advice to users. Requesting "MPI_THREAD_SINGLE" thread support level is generally not recommended, because this will conflict with other components of an application requesting higher levels of thread support. (*End of advice to users.*)

Advice to implementors. Owing to the restrictions of the MPI_THREAD_SINGLE thread support level, implementors are discouraged from making this the default thread support level for Sessions. (*End of advice to implementors.*)

MPI_SESSION_FINALIZE(session)

INOUT session session to be finalized (handle)

C binding

int MPI_Session_finalize(MPI_Session *session)

Fortran 2008 binding

MPI_Session_finalize(session, ierror)
 TYPE(MPI_Session), INTENT(INOUT) :: session
 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

```
MPI_SESSION_FINALIZE(SESSION, IERROR)
    INTEGER SESSION, IERROR
```

This routine cleans up all MPI state associated with the supplied session. Every instantiated Session must be finalized using `MPI_SESSION_FINALIZE`. The handle `session` is set to `MPI_SESSION_NULL` by the call.

Before an MPI process invokes `MPI_SESSION_FINALIZE`, the process must perform all MPI calls needed to complete its involvement in MPI communications: it must locally complete all MPI operations that it initiated and it must execute matching calls needed to complete MPI communications initiated by other processes. This means that before calling `MPI_SESSION_FINALIZE`, all message handles associated with this session must be received (with `MPI_MRECV` or derived procedures) and all request handles associated with this session must be freed in the case of nonblocking operations, and must be inactive or freed in the case of persistent operations (i.e., by calling one of the procedures `MPI_{TEST|WAIT}{[ANY|SOME|ALL]}` or `MPI_REQUEST_FREE`).

The call to `MPI_SESSION_FINALIZE` does not free objects created by MPI calls; these objects are freed using `MPI_XXX_FREE`, `MPI_COMM_DISCONNECT`, or `MPI_FILE_CLOSE` calls.

Once `MPI_SESSION_FINALIZE` returns, no MPI procedure may be called in the Sessions Model that are related to this session (not even freeing objects that are derived from this session), except for those listed in Section 11.4.1.

Advice to users. Opaque objects and their handles may bind internal resources. Therefore, it is highly recommended to explicitly free the handles associated with this session before finalizing it. Such associated handles can be group, communicator, window, file, message, and request handles, whereas datatype, operation (e.g., for reductions), error handler, and info handles exist independently of the World Model or a session in the Sessions Model. In addition, if attributes are cached on such an opaque object (see Section 7.7), then the delete callback functions are only invoked when the object is explicitly freed (or disconnected). (*End of advice to users.*)

Most handles that exist independently from the World Model or a session in the Sessions Model, e.g., datatype handles, can be created only while MPI is initialized. For example, a datatype handle that was created when one particular session existed can be used in any other session (or in the World Model), even if the second session was initialized after the first session had already been finalized and no other session existed in between. See Section 11.4.1 for handle creation procedures that do not require that MPI is initialized.

`MPI_SESSION_FINALIZE` may be synchronizing on any or all of the groups associated with communicators, windows, or files derived from the session and not disconnected, freed, or closed, respectively, before the call to the `MPI_SESSION_FINALIZE` procedure.

`MPI_SESSION_FINALIZE` behaves as if all such synchronizations occur concurrently. As `MPI_COMM_FREE` may mark a communicator for freeing later, `MPI_SESSION_FINALIZE` may be synchronizing on the group associated with a communicator that is only freed (with `MPI_COMM_FREE`) rather than disconnected (with `MPI_COMM_DISCONNECT`).

Rationale. This rule is similar to the rule that `MPI_FINALIZE` is collective (see 11.2.2), but does not require that `MPI_SESSION_FINALIZE` be collective over all connected MPI processes. It also allows for cases where some MPI processes may have derived a

set of communicators using a different number of session handles. See Example 11.9. (*End of rationale.*)

Advice to implementors. This rule also allows for the completion of communications the MPI process is involved with that may not yet be completed from the viewpoint of the underlying MPI system. See Section 2.9 on *progress* and the advice to implementors at the end of Section 11.2.2. (*End of advice to implementors.*)

Advice to implementors. An MPI implementation should be able to implement the semantics of MPI_SESSION_FINALIZE as a *local* procedure, provided an application frees all MPI windows, closes all MPI files, and uses MPI_COMM_DISCONNECT to free all MPI communicators associated with a session prior to invoking MPI_SESSION_FINALIZE on the corresponding session handle. (*End of advice to implementors.*)

Example 11.9. Three MPI processes are connected with 2 communicators (indicated by the = symbols), derived from one session handle in process X but from two separate session handles in both process Y and Z.

process-X	process-Y	process-Z	Remarks
			sesX, sesYA, ses YB, sesZA and sesZB are session handles.
(sesX)=====	(sesYA)=====	(sesZA)	communicator_1 and
(sesX)=====	(sesYB)=====	(sesZB)	communicator_2 are derived from them.
SF(sesX)	SF(sesYA)	SF(sesZA)	SF = MPI_SESSION_FINALIZE
	SF(sesYB)	SF(sesZB)	

Process X has only to finalize its one session handle, whereas the other two MPI processes have to call MPI_SESSION_FINALIZE twice in the same sequence with respect to the communicators derived from the session handles. Specifically, both process Y and process Z shall call MPI_SESSION_FINALIZE for the session from which communicator_1 was derived before calling the MPI_SESSION_FINALIZE for the session from which communicator_2 was derived, or vice versa (i.e., both shall finalize the session for communicator_2 first then finalize the session for communicator_1). The call SF(ses) in process X may not return until both SF(ses*A) and SF(ses*B) are called in processes Y and Z.

11.3.2 Processes Sets

Process sets are the mechanism for MPI applications to query the runtime. Process sets are identified by process set names. Process set names have a *Uniform Resource Identifier* (URI) format. Two process set names are mandated: "mpi://WORLD" and "mpi://SELF". Additional process set names may be defined, for example, "mpix://UNIVERSE" and "hwloc://L3Cache" may be defined by the MPI implementation. The "mpi://" namespace is reserved for exclusive use by the MPI standard. Figure 11.2 depicts process sets that the runtime could associate with an instance of an MPI job. In this example, the two mandated process sets are defined, in addition to optional, implementation specific ones.

Mechanisms for defining process sets and how system resources are assigned to these sets is considered to be implementation dependent.




Figure 11.2: Examples of process sets. Illustrated are the two mandated process sets—`"mpi://WORLD"` and `"mpi://SELF"`—along with several optional ones that a runtime could define. In this example, `MPI_SESSION_GET_NUM_PSETS` would return five at each MPI process.

A process set caches key/value tuples that are accessible to the application via an `MPI_Info` object. The `"mpi_size"` key is mandatory for all process sets.

11.3.3 Runtime Query Functions

`MPI_SESSION_GET_NUM_PSETS(session, info, npset_names)`

IN	<code>session</code>	session (handle)
IN	<code>info</code>	info object (handle)
OUT	<code>npset_names</code>	number of available process sets (non-negative integer)

C binding

```
int MPI_Session_get_num_psets(MPI_Session session, MPI_Info info,
                             int *npset_names)
```

Fortran 2008 binding

```
MPI_Session_get_num_psets(session, info, npset_names, ierror)
  TYPE(MPI_Session), INTENT(IN) :: session
  TYPE(MPI_Info), INTENT(IN) :: info
  INTEGER, INTENT(OUT) :: npset_names
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SESSION_GET_NUM_PSETS(SESSION, INFO, NPSET_NAMES, IERROR)
    INTEGER SESSION, INFO, NPSET_NAMES, IERROR
```

This function is used to query the runtime for the number of available process sets in which the calling MPI process is a member. An MPI implementation is allowed to increase the number of available process sets during the execution of an MPI application when new process sets become available. However, MPI implementations are not allowed to change the index of a particular process set name, or to change the name of the process set at a particular index, or to delete a process set name once it has been added. When a process set becomes invalid, for example, when some processes become unreachable due to failures in the communication system, subsequent usage of the process set name should raise an error. For example, creating an MPI_Group from such a process set might succeed because it is a local operation, but creating an MPI_Comm from that group and attempting collective communication should raise an error.

Advice to implementors. It is anticipated that an MPI implementation may be relying on an external runtime system to provide process sets. Such runtime systems may have the ability to dynamically create process sets during the course of application execution. Requiring the number of process sets returned by MPI_SESSION_GET_NUM_PSETS to be constant over the course of application execution would prevent an application from taking advantage of such capabilities. (*End of advice to implementors.*)

```
MPI_SESSION_GET_NTH_PSET(session, info, n, pset_len, pset_name)
```

IN	session	session (handle)
IN	info	info object (handle)
IN	n	index of the desired process set name (integer)
INOUT	pset_len	length of the pset_name argument (integer)
OUT	pset_name	name of the nth process set (string)

C binding

```
int MPI_Session_get_nth_pset(MPI_Session session, MPI_Info info, int n,
    int *pset_len, char *pset_name)
```

Fortran 2008 binding

```
MPI_Session_get_nth_pset(session, info, n, pset_len, pset_name, ierror)
    TYPE(MPI_Session), INTENT(IN) :: session
    TYPE(MPI_Info), INTENT(IN) :: info
    INTEGER, INTENT(IN) :: n
    INTEGER, INTENT(INOUT) :: pset_len
    CHARACTER(LEN=*), INTENT(OUT) :: pset_name
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SESSION_GET_NTH_PSET(SESSION, INFO, N, PSET_LEN, PSET_NAME, IERROR)
```

```

1     INTEGER SESSION, INFO, N, PSET_LEN, IERROR
2     CHARACTER*(*) PSET_NAME

```

This function returns the name of the *n*th process set in the supplied `pset_name` buffer. `pset_len` is the size of the buffer needed to store the *n*th process set name. If the `pset_len` passed into the function is less than the actual buffer size needed for the process set name, then the string value returned in `pset_name` is truncated. If `pset_len` is set to 0, `pset_name` is not changed. On return, the value of `pset_len` will be set to the required buffer size to hold the process set name. In C, `pset_len` includes the required space for the null terminator. In C, this function returns a null terminated string in all cases where the `pset_len` input value is greater than 0.

If two MPI processes get the same process set name, then the intersection of the two process sets shall either be the empty set or identical to the union of the two process sets.

After a successful call to `MPI_SESSION_GET_NTH_PSET`, subsequent calls to routines that query information about the same process set name and same session handle must return the same information. An MPI implementation is not allowed to alter any of the returned process set names.

Process set names have an implementation-defined maximum length of `MPI_MAX_PSET_NAME_LEN` characters. `MPI_MAX_PSET_NAME_LEN` shall have a value of at least 63.

Advice to users. `MPI_MAX_PSET_NAME_LEN` might be very large, so it might not be wise to declare a string of that size. Users are encouraged to use `MPI_SESSION_GET_NTH_PSET` both for obtaining the length of a `pset_name` and the process set name. (*End of advice to users.*)

```

28 MPI_SESSION_GET_INFO(session, info_used)

```

```

29     IN      session          session (handle)
30     OUT    info_used        see explanation below (handle)

```

C binding

```

34 int MPI_Session_get_info(MPI_Session session, MPI_Info *info_used)

```

Fortran 2008 binding

```

36 MPI_Session_get_info(session, info_used, ierror)
37     TYPE(MPI_Session), INTENT(IN) :: session
38     TYPE(MPI_Info), INTENT(OUT) :: info_used
39     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

42 MPI_SESSION_GET_INFO(SESSION, INFO_USED, IERROR)
43     INTEGER SESSION, INFO_USED, IERROR

```

`MPI_SESSION_GET_INFO` returns a new info object containing the hints of the MPI Session associated with `session`. The current setting of all hints related to this MPI Session is returned in `info_used`. An MPI implementation is required to return all hints that are supported by the implementation and have default values specified; any user-supplied hints that were not ignored by the implementation; and any additional hints that were set by

the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no key/value pair. The user is responsible for freeing info_used via MPI_INFO_FREE.

MPI_SESSION_GET_PSET_INFO(session, pset_name, info)

IN	session	session (handle)
IN	pset_name	name of process set (string)
OUT	info	info object containing information about the given process set (handle)

C binding

```
int MPI_Session_get_pset_info(MPI_Session session, const char *pset_name,
                             MPI_Info *info)
```

Fortran 2008 binding

```
MPI_Session_get_pset_info(session, pset_name, info, ierror)
  TYPE(MPI_Session), INTENT(IN) :: session
  CHARACTER(LEN=*), INTENT(IN) :: pset_name
  TYPE(MPI_Info), INTENT(OUT) :: info
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SESSION_GET_PSET_INFO(SESSION, PSET_NAME, INFO, IERROR)
  INTEGER SESSION, INFO, IERROR
  CHARACTER*(*) PSET_NAME
```

This function is used to query properties of a specific process set. The returned info object can be queried with existing MPI info object query functions. One key/value pair must be defined, "mpi_size". The value of the "mpi_size" key specifies the number of MPI processes in the process set. The user is responsible for freeing the returned MPI_Info object.

11.3.4 Sessions Model Examples

This section presents several examples of how to use MPI Sessions to create MPI Groups and MPI Communicators.

Example 11.10. Simple example illustrating creation of an MPI communicator using the Sessions Model.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include "mpi.h"

static MPI_Session lib_shandle = MPI_SESSION_NULL;
static MPI_Comm lib_comm = MPI_COMM_NULL;

int library_foo_init(void)
{
  int rc, flag, valuelen;
```

```
1  int ret = 0;
2  const char pset_name[] = "mpi://WORLD";
3  const char mt_key[] = "thread_level";
4  const char mt_value[] = "MPI_THREAD_MULTIPLE";
5  char out_value[100]; /* large enough */
6  MPI_Group wgroup = MPI_GROUP_NULL;
7  MPI_Info sinfo = MPI_INFO_NULL;
8  MPI_Info tinfo = MPI_INFO_NULL;
9
10 MPI_Info_create(&sinfo);
11 MPI_Info_set(sinfo, mt_key, mt_value);
12 rc = MPI_Session_init(sinfo, MPI_ERRORS_RETURN,
13                      &lib_shandle);
14 if (rc != MPI_SUCCESS) {
15     ret = -1;
16     goto fn_exit;
17 }
18 /*
19  * check we got thread support level foo library needs
20  */
21 rc = MPI_Session_get_info(lib_shandle, &tinfo);
22 if (rc != MPI_SUCCESS) {
23     ret = -1;
24     goto fn_exit;
25 }
26 valuelen = sizeof(out_value);
27 MPI_Info_get_string(tinfo, mt_key, &valuelen,
28                    out_value, &flag);
29 if (0 == flag) {
30     printf("Could not find key %s\n", mt_key);
31     ret = -1;
32     goto fn_exit;
33 }
34 if (strcmp(out_value, mt_value)) {
35     printf("Did not get thread multiple support, got %s\n",
36           out_value);
37     ret = -1;
38     goto fn_exit;
39 }
40 /*
41  * create a group from the WORLD process set
42  */
43 rc = MPI_Group_from_session_pset(lib_shandle,
44                                  pset_name,
45                                  &wgroup);
46 if (rc != MPI_SUCCESS) {
47     ret = -1;
48     goto fn_exit;
49 }
```



```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

```

/*
 * get a communicator
 */
rc = MPI_Comm_create_from_group(wgroup,
                                "org.mpi-forum.mpi-v4_0.example-ex11_10",
                                MPI_INFO_NULL,
                                MPI_ERRORS_RETURN,
                                &lib_comm);

if (rc != MPI_SUCCESS) {
    ret = -1;
    goto fn_exit;
}

/*
 * free group, library doesn't need it.
 */

fn_exit:
    MPI_Group_free(&wgroup);

    if (sinfo != MPI_INFO_NULL) {
        MPI_Info_free(&sinfo);
    }

    if (tinfo != MPI_INFO_NULL) {
        MPI_Info_free(&tinfo);
    }

    if (ret != 0) {
        MPI_Session_finalize(&lib_shandle);
    }

    return ret;
}

```

Example 11.10 shows how the predefined "mpi://WORLD" process set can be used to first create a local MPI group and then subsequently to create an MPI communicator from this group.

Example 11.11. This example illustrates the use of Process Set query functions to select a Process Set to use for MPI Group creation.

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include "mpi.h"

int main(int argc, char *argv[])
{
    int i, n_psets, psetlen, rc, ret;
    int valuelen;
    int flag = 0;
    char *pset_name = NULL;

```

```

1  char *info_val = NULL;
2  MPI_Session shandle = MPI_SESSION_NULL;
3  MPI_Info sinfo = MPI_INFO_NULL;
4  MPI_Group pgroup = MPI_GROUP_NULL;
5
6  if (argc < 2) {
7      fprintf(stderr, "A process set name fragment is required\n");
8      return EXIT_FAILURE;
9  }
10
11 rc = MPI_Session_init(MPI_INFO_NULL, MPI_ERRORS_RETURN, &shandle);
12 if (rc != MPI_SUCCESS) {
13     fprintf(stderr, "Could not initialize session, bailing out\n");
14     return EXIT_FAILURE;
15 }
16
17 MPI_Session_get_num_psets(shandle, MPI_INFO_NULL, &n_psets);
18
19 for (i=0, pset_name=NULL; i<n_psets; i++) {
20     psetlen = 0;
21     MPI_Session_get_nth_pset(shandle, MPI_INFO_NULL, i,
22                             &psetlen, NULL);
23     pset_name = (char *)malloc(sizeof(char) * psetlen);
24     MPI_Session_get_nth_pset(shandle, MPI_INFO_NULL, i,
25                             &psetlen, pset_name);
26     if (strstr(pset_name, argv[1]) != NULL) break;
27
28     free(pset_name);
29     pset_name = NULL;
30 }
31
32 /*
33  * get instance of an info object for this Session
34  */
35
36 MPI_Session_get_pset_info(shandle, pset_name, &sinfo);
37 valuelen = 0;
38 MPI_Info_get_string(sinfo, "mpi_size", &valuelen, NULL, &flag);
39 if (flag) {
40     info_val = (char *)malloc(valuelen);
41     MPI_Info_get_string(sinfo, "mpi_size", &valuelen, info_val, &flag);
42     free(info_val);
43 }
44
45 /*
46  * create a group from the process set
47  */
48
49 rc = MPI_Group_from_session_pset(shandle, pset_name,
50                                 &pgroup);
51 ret = (rc == MPI_SUCCESS) ? 0 : EXIT_FAILURE;
52
53 free(pset_name);
54 MPI_Group_free(&pgroup);
55 MPI_Info_free(&sinfo);
56 MPI_Session_finalize(&shandle);
57
58

```

```

    fprintf(stderr, "Test completed ret = %d\n", ret);
    return ret;
}

```

Example 11.11 illustrates several aspects of the Sessions Model. First, the default error handler can be specified when instantiating a Session instance. Second, there must be at least two process sets associated with a Session. Third, the example illustrates use of the Sessions info object and the one required key: "mpi_size".

Example 11.12. A Fortran 2008 example illustrating how to obtain information about available process sets, create an MPI Group from a process set, and subsequently create an MPI Communicator.

```

PROGRAM MAIN
  USE mpi_f08
  IMPLICIT NONE
  INTEGER :: pset_len, ierror, n_psets
  CHARACTER(LEN=:), ALLOCATABLE :: pset_name
  TYPE(MPI_Session) :: shandle
  TYPE(MPI_Group) :: pgroup
  TYPE(MPI_Comm) :: pcomm

  CALL MPI_Session_init(MPI_INFO_NULL, MPI_ERRORS_RETURN, &
    shandle, ierror)
  IF (ierror .NE. MPI_SUCCESS) THEN
    WRITE(*,*) "MPI_Session_init failed"
    ERROR STOP
  END IF

  CALL MPI_Session_get_num_psets(shandle, MPI_INFO_NULL, n_psets)
  IF (n_psets .LT. 2) THEN
    WRITE(*,*) "MPI_Session_get_num_psets didn't return at least 2 psets"
    ERROR STOP
  END IF

  !
  ! Just get the second pset's length and name
  ! Note that index values are zero-based, even in Fortran
  !
  pset_len = 0
  CALL MPI_Session_get_nth_pset(shandle, MPI_INFO_NULL, 1, &
    pset_len, pset_name)
  ALLOCATE(CHARACTER(LEN=pset_len)::pset_name)
  CALL MPI_Session_get_nth_pset(shandle, MPI_INFO_NULL, 1, &
    pset_len, pset_name)

  !
  ! create a group from the pset
  !
  CALL MPI_Group_from_session_pset(shandle, pset_name, pgroup)

  !
  ! free the buffer used for the pset name
  !
  DEALLOCATE(pset_name)

```

```

1  !
2  !   create a MPI communicator from the group
3  !
4  CALL MPI_Comm_create_from_group(pgroup, "session_example",    &
5                                MPI_INFO_NULL,                &
6                                MPI_ERRORS_RETURN,            &
7                                pcomm)
8
9  CALL MPI_Barrier(pcomm, ierror)
10 IF (ierror .NE. MPI_SUCCESS) THEN
11     WRITE(*,*) "Barrier call on communicator failed"
12     ERROR STOP
13 END IF
14
15 CALL MPI_Comm_free(pcomm)
16 CALL MPI_Group_free(pgroup)
17 CALL MPI_Session_finalize(shandle, ierror)
18
19 END PROGRAM MAIN

```

Note in this example that the call to `MPI_SESSION_FINALIZE` may block in order to ensure that the calling MPI process has completed its involvement in the preceding `MPI_BARRIER` operation. If `MPI_COMM_DISCONNECT` had been used instead of `MPI_COMM_FREE`, the example would have blocked in `MPI_COMM_DISCONNECT` rather than `MPI_SESSION_FINALIZE`.

11.4 Common Elements of Both Process Models

11.4.1 MPI Functionality that is Always Available

Some MPI functions may be invoked at any time, including prior to calling `MPI_INIT` or `MPI_SESSION_INIT`, and following MPI finalization, independent of whether the World Model, Sessions Model, or both are used. These functions can be called concurrently by multiple threads within an MPI Process. Table 11.1 lists the applicable MPI functions.

In addition to the functions listed in Table 11.1, any function with the prefix `MPI_T_` (within the constraints for functions with this prefix listed in Section 15.3.4) may also be called prior to MPI initialization and after MPI finalization.

11.4.2 Aborting MPI Processes

`MPI_ABORT(comm, errorcode)`

IN	comm	communicator of MPI processes to abort (handle)
IN	errorcode	error code to return to invoking environment (integer)

C binding

`int MPI_Abort(MPI_Comm comm, int errorcode)`

Table 11.1: List of MPI Functions that can be called at any time within an MPI program, including prior to MPI initialization and following MPI finalization

MPI_INITIALIZED
MPI_FINALIZED
MPI_GET_VERSION
MPI_GET_LIBRARY_VERSION
MPI_INFO_CREATE
MPI_INFO_CREATE_ENV
MPI_INFO_SET
MPI_INFO_DELETE
MPI_INFO_GET_STRING
MPI_INFO_GET_NKEYS
MPI_INFO_GET_NTHKEY
MPI_INFO_DUP
MPI_INFO_FREE
MPI_INFO_F2C
MPI_INFO_C2F
MPI_SESSION_CREATE_ERRHANDLER
MPI_SESSION_CALL_ERRHANDLER
MPI_ERRHANDLER_FREE
MPI_ERRHANDLER_F2C
MPI_ERRHANDLER_C2F
MPI_ERROR_STRING
MPI_ERROR_CLASS
MPI_ADD_ERROR_CLASS
MPI_REMOVE_ERROR_CLASS
MPI_ADD_ERROR_CODE
MPI_REMOVE_ERROR_CODE
MPI_ADD_ERROR_STRING
MPI_REMOVE_ERROR_STRING

Fortran 2008 binding

```
MPI_Abort(comm, errorcode, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, INTENT(IN) :: errorcode
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_ABORT(COMM, ERRORCODE, IERROR)
  INTEGER COMM, ERRORCODE, IERROR
```

This routine makes a “best attempt” to abort all MPI processes in the group of `comm`. This function does not require that the invoking environment take any action with the error code. However, a Unix or POSIX environment should handle this as a return `errorcode` from the main program.

1 It may not be possible for an MPI implementation to abort only the processes rep-
2 resented by `comm` if this is a subset of the processes. In this case, the MPI implemen-
3 tation should attempt to abort all the connected processes but should not abort any un-
4 connected processes. When using the World Model, and if no processes were spawned,
5 accepted, or connected then this has the effect of aborting all the processes associated with
6 `MPI_COMM_WORLD`. In the case of the Sessions Model, if an MPI process has instantiated
7 multiple sessions, the union of the process sets in these sessions are considered connected
8 processes. Thus invoking `MPI_ABORT` on a communicator derived from one of these ses-
9 sions will result in all MPI processes in this union being aborted.

10
11 *Advice to implementors.* After aborting a subset of processes, a high quality im-
12 plementation should be able to provide error handling for communicators, windows,
13 and files involving both aborted and nonaborted processes. As an example, if the
14 user changes the error handler for `MPI_COMM_WORLD` to `MPI_ERRORS_RETURN` or a
15 custom error handler, when a subset of `MPI_COMM_WORLD` is aborted, the remaining
16 processes in `MPI_COMM_WORLD` should be able to continue communicating with each
17 other and receive an appropriate error code when attempting communication with
18 an aborted process (e.g., an error of class `MPI_ERR_PROC_ABORTED`). A high quality
19 implementation should support equivalent behavior for communicators derived from
20 sessions. (*End of advice to implementors.*)

21
22 *Advice to users.* Whether the `errorcode` is returned from the executable or from the
23 MPI process startup mechanism (e.g., `mpiexec`), is an aspect of quality of the MPI
24 library but not mandatory. (*End of advice to users.*)

25
26 *Advice to implementors.* Where possible, a high-quality implementation will try
27 to return the `errorcode` from the MPI process startup mechanism (e.g. `mpiexec` or
28 singleton `init`). (*End of advice to implementors.*)

29 11.4.3 Memory Allocation Info

30
31 Computing systems contain memory with different properties, including differences in per-
32 formance, persistence, access permissions, or access mode. These distinct memories are
33 generally allocated using distinct mechanisms and are referred to as memory allocation
34 kinds that are named according to the method of allocation. The following info keys can
35 be used to request or query the memory allocation kinds supported by the MPI library and
36 to assert application usage of memory allocation kinds with respect to specific MPI objects,
37 as shown in Example 11.13.

38
39 **"`mpi_memory_alloc_kinds`" (string, default: "`mpi,system`"):** A comma separated list of
40 memory allocation kinds. When set on the input info object in a call to
41 `MPI_SESSION_INIT`, `MPI_COMM_SPAWN`, or `MPI_COMM_SPAWN_MULTIPLE`, or
42 when supplied as an argument to an MPI startup mechanism, this info key requests
43 support for the specified memory allocation kinds.

44 When returned by MPI, this info key indicates the memory allocation kinds supported
45 by the MPI library on the given session, MPI object, or objects derived from the World
46 Model. This info key does not affect the kind of memory allocated by MPI, e.g., in
47 a call to `MPI_ALLOC_MEM` or `MPI_WIN_ALLOCATE`. A value corresponding to the
48 empty string represents no memory allocation kinds.

"mpi_assert_memory_alloc_kinds" (string, not set by default): A comma separated list of memory allocation kinds that the calling MPI process will use with the given MPI object. A value corresponding to the empty string represents no memory allocation kinds.

The "mpi_memory_alloc_kinds" info key is used both for requesting and querying support for memory allocation kinds from the MPI library.

When supplied to MPI_SESSION_INIT, this info key requests support for memory allocation kinds for all objects that will be derived from the new session. This info hint can also be supplied through an argument to an MPI startup mechanism. In the Sessions Model, this behaves as though the "mpi_memory_alloc_kinds" info key with the given value was supplied in the info argument in calls to MPI_SESSION_INIT. A value of "mpi_memory_alloc_kinds" supplied in the info argument to MPI_SESSION_INIT takes precedence over a value supplied as an argument to an MPI startup mechanism.

In the World Model, an info hint passed to an MPI startup mechanism requests support for memory allocation kinds for all objects derived from the World Model. This info hint can also be supplied to MPI_COMM_SPAWN or MPI_COMM_SPAWN_MULTIPLE in the World Model. This requests support for memory allocation kinds for all objects derived from the World Model in the spawned MPI process or MPI processes.

When returned by MPI_SESSION_GET_INFO, this info key indicates the memory allocation kinds supported by the MPI library on the given session. When returned in MPI_INFO_ENV, this info key indicates the memory allocation kinds supported by the MPI library for all objects derived from the World Model.

If "mpi_memory_alloc_kinds" was supplied during session creation, then the value of the corresponding key in the info object returned by MPI_SESSION_GET_INFO must include all requested memory allocation kinds that are supported. The substrings that indicate support for these memory allocation kinds must be identical to those supplied by the user. MPI may also return additional memory allocation kinds that were not requested by the user. The order of the memory allocation kinds returned through this info key is undefined.

Rationale. MPI libraries may have implementation-specific mechanisms (e.g., environment variables) that control the supported memory allocation kinds. Allowing implementations to return additional memory allocation kinds provides for compatibility with such mechanisms. (*End of rationale.*)

The "mpi_memory_alloc_kinds" info key must also be contained in the info object returned by MPI_COMM_GET_INFO, MPI_WIN_GET_INFO, and MPI_FILE_GET_INFO. If the communicator, window, or file is derived from the World Model, the value of this info key must be identical to the value of this info key in MPI_INFO_ENV unless the user has asserted that support for memory allocation kinds can be restricted by setting "mpi_assert_memory_alloc_kinds" on that communicator, window, or file. If the communicator, window, or file is derived from the Sessions Model, the value of this info key must be identical to the value of this info key in the info object returned by MPI_SESSION_GET_INFO for that session unless the user has asserted that support for memory allocation kinds can be restricted by setting "mpi_assert_memory_alloc_kinds" on that communicator, window, or file.

When the user sets the "mpi_assert_memory_alloc_kinds" info key on the input info object for communicator creation, including via MPI_COMM_SPAWN or MPI_COMM_SPAWN_MULTIPLE, window creation, or file creation the implementation may

assume that the memory for all communication buffers passed to MPI operations performed by the calling MPI process on the newly created MPI object will use only the memory allocation kinds listed in the value string. If the MPI library does not support one or more of the allocation kinds associated with the "mpi_assert_memory_alloc_kinds" info key, it will ignore this info key. When an MPI library recognizes this info key, the value returned when querying this info key (e.g., through a call to MPI_COMM_GET_INFO) must be identical to the value supplied by the user. It is erroneous to pass a communication buffer with an unsupported memory allocation kind to an MPI routine.

Memory allocation kind strings are comma separated lists that follow the rules specified in Section 10. Each element in the list is a memory allocation kind that is formatted as the name of the kind, followed by an optional colon separated list of restrictors. Whitespace is not permitted within the list of restrictors. For example, "kind_a:restrictor_1,kind_b:restrictor_1:restrictor_2,...".

Within a memory allocation kind string, a given kind may be listed more than once with different restrictors, e.g., "kind_a:restrictor_1,kind_a:restrictor_2". A given kind may also be listed more than once with fewer restrictors, e.g., "kind_a,kind_a:restrictor_1". A memory allocation kind with no restrictors indicates an unrestricted memory allocation kind. Each instance of a kind in the memory allocation kind string indicates a separate and potentially overlapping memory allocation kind. The following memory allocation kinds and restrictors are defined by MPI. This list may be extended by MPI side documents and implementations.

- "system": Memory allocated by standard operating system allocators. When support for the "system" memory allocation kind is requested by the user, it must be provided by the MPI library.
- "mpi": Memory allocated by the MPI library. When support for the "mpi" memory allocation kind is requested by the user, it must be provided by the MPI library.

Restrictors for the "mpi" memory allocation kind:

- "alloc_mem": Memory allocated by a call to MPI_ALLOC_MEM
- "win_allocate": Memory allocated by a call to MPI_WIN_ALLOCATE
- "win_allocate_shared": Memory allocated by a call to MPI_WIN_ALLOCATE_SHARED

Example 11.13. This example demonstrates the usage of memory allocation kinds info keys with the Sessions Model. It shows how support for additional memory allocation kinds can be requested, how supported memory allocation kinds can be queried, how to parse the list of supported memory allocation kinds, and how to assert that a subset of supported memory allocation kinds are used with operations on a specific communicator.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <mpi.h>

int main(int argc, char *argv[])
{
    int gpu_aware = 0, len = 0, flag = 0;
    MPI_Info info;
```



```

1 MPI_Session session;
2 MPI_Group wgroup;
3 MPI_Comm system_comm, gpu_comm = MPI_COMM_NULL;
4
5 MPI_Info_create(&info);
6 MPI_Info_set(info, "mpi_memory_alloc_kinds", "system,gpu:device");
7 MPI_Session_init(info, MPI_ERRORS_ARE_FATAL, &session);
8 MPI_Info_free(&info);
9
10 MPI_Session_get_info(session, &info);
11 MPI_Info_get_string(info, "mpi_memory_alloc_kinds",
12                     &len, NULL, &flag);
13
14 if (flag) {
15     char *val, *valptr, *kind;
16
17     val = valptr = (char *) malloc(len);
18     if (NULL == val) return 1;
19
20     MPI_Info_get_string(info, "mpi_memory_alloc_kinds",
21                         &len, val, &flag);
22
23     while ((kind = strsep(&val, ",")) != NULL) {
24         if (strcasecmp(kind, "gpu:device") == 0) {
25             gpu_aware = 1;
26             break;
27         }
28     }
29     free(valptr);
30 }
31
32 MPI_Info_free(&info);
33
34 MPI_Group_from_session_pset(session, "mpi://WORLD" , &wgroup);
35
36 // Create a communicator for operations on system memory
37 MPI_Info_create(&info);
38 MPI_Info_set(info, "mpi_assert_memory_alloc_kinds", "system");
39 MPI_Comm_create_from_group(wgroup,
40                             "org.mpi-forum.example.mem-alloc-kind-usage.system",
41                             info, MPI_ERRORS_ABORT, &system_comm);
42
43 MPI_Info_free(&info);
44
45 // Check if all processes have GPU support
46 MPI_Allreduce(MPI_IN_PLACE, &gpu_aware, 1, MPI_INT, MPI_LAND,
47               system_comm);
48
49 // Create a communicator for operations that use GPU buffers.
50 // Note, the "gpu" memory allocation kind is provided as an example
51 // and is not one of the memory allocation kinds defined by the MPI
52 // standard.
53 if (gpu_aware) {

```

```

1     MPI_Info_create(&info);
2     MPI_Info_set(info, "mpi_assert_memory_alloc_kinds",
3                   "gpu:device");
4     MPI_Comm_create_from_group(wgroup,
5                                "org.mpi-forum.example.mem-alloc-kind-usage.gpu",
6                                info, MPI_ERRORS_ABORT, &gpu_comm);
7     MPI_Info_free(&info);
8   }
9   else {
10    printf("Warning: GPU alloc kind not supported\n");
11  }
12
13  MPI_Group_free(&wgroup);
14
15  // Perform communication using gpu_comm if it's available.
16  // Otherwise, copy data to a system buffer and use system_comm.
17
18  if (gpu_comm != MPI_COMM_NULL) MPI_Comm_disconnect(&gpu_comm);
19  MPI_Comm_disconnect(&system_comm);
20
21  MPI_Session_finalize(&session);
22
23  return 0;
24 }

```

11.5 Portable MPI Process Startup

A number of implementations of MPI provide a startup command for MPI programs that is of the form

```
mpirun <mpirun arguments> <program> <program arguments>
```

Separating the command to start the program from the program itself provides flexibility, particularly for network and heterogeneous implementations. For example, the startup script need not run on one of the machines that will be executing the MPI program itself.

Having a standard startup mechanism also extends the portability of MPI programs one step further, to the command lines and scripts that manage them. For example, a validation suite script that runs hundreds of programs can be a portable script if it is written using such a standard startup mechanism. In order that the “standard” command not be confused with existing practice, which is not standard and not portable among implementations, instead of `mpirun` MPI specifies `mpiexec`.

While a standardized startup mechanism improves the usability of MPI, the range of environments is so diverse (e.g., there may not even be a command line interface) that MPI cannot mandate such a mechanism. Instead, MPI specifies an `mpiexec` startup command and recommends, but does not require it, as advice to implementors. However, if an implementation does provide a command called `mpiexec`, it must be of the form described below.

It is suggested that

```
mpiexec -n <numprocs> <program>
```

be at least one way to start `<program>` with an initial set of `<numprocs>` processes, which will be accessible as the process set named "mpi://WORLD" in the Sessions Model and/or used to form the group associated with the built-in communicator, MPI_COMM_WORLD in the World Model. Other arguments to `mpiexec` may be implementation-dependent.

Advice to implementors. Implementors, if they do provide a special startup command for MPI programs, are advised to give it the following form. The syntax is chosen in order that `mpiexec` be able to be viewed as a command-line version of MPI_COMM_SPAWN (See Section 11.8.4).

Analogous to MPI_COMM_SPAWN, we have

```
mpiexec -n                <maxprocs>
      -soft                <      >
      -host                <      >
      -arch                <      >
      -wdir                <      >
      -path                <      >
      -file                <      >
      -initial-errhandler <      >
      -memory-alloc-kinds <      >
      ...
      <command line>
```

for the case where a single command line for the application program and its arguments will suffice. See Section 11.8.4 for the meanings of these arguments. For the case corresponding to MPI_COMM_SPAWN_MULTIPLE there are two possible formats:

Form A:

```
mpiexec { <above arguments> } : { ... } : { ... } : ... : { ... }
```

As with MPI_COMM_SPAWN, all the arguments are optional. (Even the `-n x` argument is optional; the default is implementation dependent. It might be 1, it might be taken from an environment variable, or it might be specified at compile time.) The names and meanings of the arguments are taken from the keys in the info argument to MPI_COMM_SPAWN. There may be other, implementation-dependent arguments as well.

Note that Form A, though convenient to type, prevents colons from being program arguments. Therefore an alternate, file-based form is allowed:

Form B:

```
mpiexec -configfile <filename>
```

where the lines of `<filename>` are of the form separated by the colons in Form A. Lines beginning with '#' are comments, and lines may be continued by terminating the partial line with '\'.

Example 11.14. Start 16 instances of `myprog` on the current or default machine:

```
mpiexec -n 16 myprog
```

Example 11.15. Start 10 instances of `myprog` on the machine called `ferrari`:

```
mpiexec -n 10 -host ferrari myprog
```

Example 11.16. Start 3 instances of the same program `myprog` with different command-line arguments:

```
mpiexec myprog infile1 : myprog infile2 : myprog infile3
```

Example 11.17. Start 5 instances of the `ocean` program on `x86_64` hosts and 10 instances of the `atmos` program on Power9 hosts (Form B):

```
mpiexec -n 5 -arch x86_64 ocean : -n 10 -arch power9 atmos
```

It is assumed that the implementation in this case has a method for choosing hosts of the appropriate type. Their ranks in `MPI_COMM_WORLD` are in the order specified.

Example 11.18. Start the `ocean` program on five Suns and the `atmos` program on 10 RS/6000's (Form B):

```
mpiexec -configfile myfile
```

where `myfile` contains

```
-n 5 -arch sun ocean
-n 10 -arch rs6000 atmos
```

(End of advice to implementors.)

11.6 MPI and Threads

This section specifies the interaction between MPI calls and threads. Although thread compliance is not required, the standard specifies how threads are to work if they are provided. The section lists minimal requirements for **thread compliant** MPI implementations and defines functions that can be used for initializing the thread environment. MPI may be implemented in environments where threads are not supported or perform poorly. Therefore, MPI implementations are not required to be thread compliant as defined in this section. Regardless of whether or not the MPI implementation is thread compliant, a subset of MPI functions must always be thread safe. A complete list of such MPI functions is given in Table 11.1. When a thread is executing one of these routines, if another concurrently running thread also makes an MPI call, the outcome will be as if the calls executed in some order.

This section generally assumes a thread package similar to POSIX threads [45], but the syntax and semantics of thread calls are not specified here—these are beyond the scope of this document.

11.6.1 General

In a thread-compliant implementation, an MPI process is a process that may be multithreaded. Each thread can issue MPI calls; however, threads are not separately addressable: the rank argument in a send or receive call identifies an MPI process, not a thread. A message sent to an MPI process can be received by any thread in this MPI process.

Rationale. This model corresponds to the POSIX model of interprocess communication: the fact that a process is multithreaded, rather than single-threaded, does not affect the external interface of this process. MPI implementations in which MPI ‘processes’ are POSIX threads inside a single POSIX process are not thread-compliant by this definition (indeed, their “processes” are single-threaded). (*End of rationale.*)

Advice to users. It is the user’s responsibility to prevent races when threads within the same application post conflicting communication calls. The user can make sure that two threads in the same process will not issue conflicting communication calls by using distinct communicators at each thread. (*End of advice to users.*)

The two main requirements for a thread-compliant implementation are listed below.

1. All MPI calls are **thread-safe**, i.e., two concurrently running threads may make MPI calls and the outcome will be as if the calls executed in some order, even if their execution is interleaved.
2. Blocking MPI calls will block the calling thread only, allowing another thread to execute, if available. The calling thread will be blocked until the event on which it is waiting occurs. Once the blocked communication is enabled and can proceed, then the call will complete and the thread will be marked runnable, within a finite time. A blocked thread will not prevent *progress* of other runnable threads on the same process, and will not prevent them from executing MPI calls.

Example 11.19. Process 0 consists of two threads. The first thread executes a blocking send call `MPI_Send(buff1, count, type, 0, 0, comm)`, whereas the second thread executes a blocking receive call `MPI_Recv(buff2, count, type, 0, 0, comm, &status)`, i.e., the first thread sends a message that is received by the second thread. This communication should always succeed. According to the first requirement, the execution will correspond to some interleaving of the two calls. According to the second requirement, a call can only block the calling thread and cannot prevent progress of the other thread. If the send call went ahead of the receive call, then the sending thread may block, but this will not prevent the receiving thread from executing. Thus, the receive call will occur. Once both calls occur, the communication is enabled and both calls will complete. On the other hand, a single-threaded process that posts a send, followed by a matching receive, may deadlock. The progress requirement for multithreaded implementations is stronger, as a blocked call cannot prevent progress in other threads.

Advice to implementors. MPI calls can be made thread-safe by executing only one at a time, e.g., by protecting MPI code with one process-global lock. However, blocked operations cannot hold the lock, as this would prevent progress of other threads in the process. The lock is held only for the duration of an atomic, locally-completing

1 suboperation such as posting a send or completing a send, and is released in between.
2 Finer locks can provide more concurrency, at the expense of higher locking overheads.
3 Concurrency can also be achieved by having some of the MPI protocol executed by
4 separate server threads. (*End of advice to implementors.*)
5

6 11.6.2 Clarifications 7 8

9 **Initialization and Completion.** When using the World Model, the call to `MPI_FINALIZE`
10 should occur on the same thread that initialized MPI. We call this thread the **main thread**.
11 The call should occur only after all process threads have completed their MPI calls, and
12 have no *pending* communication or I/O operations.

13 *Rationale.* This constraint simplifies implementation. (*End of rationale.*)
14
15

16 **Threads and the Sessions Model.** The Sessions Model provides a finer-grain approach
17 to controlling the interaction between MPI calls and threads. When using this model, the
18 desired level of thread support is specified at Session initialization time. See Section 11.3.
19 Thus it is possible for communicators and other MPI objects derived from one Session to
20 provide a different level of thread support than those created from another Session for which
21 a different level of thread support was requested. Depending on the level of thread support
22 requested at Session initialization time, different threads in a MPI process can make concur-
23 rent calls to MPI when using MPI objects derived from different *session handles*. Note that
24 the requested and provided level of thread support when creating a Session may influence
25 the granted level of thread support in a subsequent invocation of `MPI_SESSION_INIT`. Like-
26 wise, if the application at some point calls `MPI_INIT_THREAD`, the requested and granted
27 level of thread support may influence the granted level of thread support for subsequent
28 calls to `MPI_SESSION_INIT`. Similarly, if the application calls `MPI_INIT_THREAD` after a
29 call to `MPI_SESSION_INIT`, the level of thread support returned from `MPI_INIT_THREAD`
30 may be similarly influenced by the requested level of thread support in the prior call to
31 `MPI_SESSION_INIT`.

32 In addition, if an MPI application is only using the Sessions Model, the provided thread
33 support level returned by `MPI_QUERY_THREAD` is the same as that returned prior to
34 invocation of `MPI_INIT_THREAD` or `MPI_INIT`. If the application also used the World Model
35 in some component of the application, `MPI_QUERY_THREAD` will return the level of thread
36 support returned by the original call to `MPI_INIT_THREAD`.
37

38 **Multiple threads completing the same request.** A program in which two threads
39 block, waiting on the same request, is erroneous. Similarly, the same request cannot appear
40 in the array of requests of two concurrent `MPI_{WAIT|TEST}_{ANY|SOME|ALL}` calls. In
41 MPI, a request can only be completed once. Any combination of wait or test that violates
42 this rule is erroneous.
43

44 *Rationale.* This restriction is consistent with the view that a multithreaded execution
45 corresponds to an interleaving of the MPI calls. In a single threaded implementation,
46 once a wait is posted on a request the request handle will be nullified before it is
47 possible to post a second wait on the same handle. With threads, an
48 `MPI_WAIT_{ANY|SOME|ALL}` may be blocked without having nullified its request(s)

so it becomes the user's responsibility to avoid using the same request in an `MPI_WAIT` on another thread. This constraint also simplifies implementation, as only one thread will be blocked on any communication or I/O event. (*End of rationale.*)

Probe. A receive call that uses source and tag values returned by a preceding call to `MPI_PROBE` or `MPI_IPROBE` will receive the message matched by the probe call only if there was no other matching receive after the probe and before that receive. In a multi-threaded environment, it is up to the user to enforce this condition using suitable mutual exclusion logic. This can be enforced by making sure that each communicator is used by only one thread on each process. Alternatively, `MPI_MPROBE` or `MPI_IMPROBE` can be used.

Collective calls. Matching of collective calls on a communicator, window, or file handle is done according to the order in which the calls are issued at each process. If concurrent threads issue such calls on the same communicator, window or file handle, it is up to the user to make sure the calls are correctly ordered, using interthread synchronization.

Advice to users. With three concurrent threads in each MPI process of a communicator `comm`, it is allowed that thread A in each MPI process calls a collective operation on `comm`, thread B calls a file operation on an existing file handle that was formerly opened on `comm`, and thread C invokes one-sided operations on an existing window handle that was also formerly created on `comm`. (*End of advice to users.*)

Rationale. As specified in `MPI_FILE_OPEN` and `MPI_WIN_CREATE`, a file handle and a window handle inherit only the group of processes of the underlying communicator, but not the communicator itself. Accesses to communicators, window handles and file handles cannot affect one another. (*End of rationale.*)

Advice to implementors. If the implementation of file or window operations internally uses MPI communication then a duplicated communicator may be cached on the file or window object. (*End of advice to implementors.*)

Error handlers. An error handler does not necessarily execute in the context of the thread that made the error-raising MPI call; the error handler may be executed by a thread that is distinct from the thread that will return the error code.

Rationale. The MPI implementation may be multithreaded, so that part of the communication protocol may execute on a thread that is distinct from the thread that made the MPI call. The design allows the error handler to be executed on the thread where the error is raised. (*End of rationale.*)

Interaction with signals and cancellations. The outcome is undefined if a thread that executes an MPI call is cancelled (by another thread), or if a thread catches a signal while executing an MPI call. However, a thread of an MPI process may terminate, and may catch signals or be cancelled by another thread when not executing MPI calls.

1 *Rationale.* Few C library functions are signal safe, and many have cancellation
2 points—points at which the thread executing them may be cancelled. The above
3 restriction simplifies implementation (no need for the MPI library to be “async-cancel-
4 safe” or “async-signal-safe”). (*End of rationale.*)

5
6 *Advice to users.* Users can catch signals in separate, non-MPI threads (e.g., by
7 masking signals on MPI calling threads, and unmasking them in one or more non-MPI
8 threads). A good programming practice is to have a distinct thread blocked in a
9 call to `sigwait` for each user expected signal that may occur. Users must not catch
10 signals used by the MPI implementation; as each MPI implementation is required to
11 document the signals used internally, users can avoid these signals. (*End of advice to*
12 *users.*)

13 *Advice to implementors.* The MPI library should not invoke library calls that are
14 not thread safe, if multiple threads execute. (*End of advice to implementors.*)

17 11.7 The Dynamic Process Model

18
19 The dynamic process model allows for the creation and cooperative termination of processes
20 after an MPI application has started. It provides a mechanism to establish communication
21 between the newly created processes and the existing MPI application. It also provides a
22 mechanism to establish communication between two existing MPI applications, even when
23 one did not “start” the other.

24 The MPI procedures described in this section require the World Model, meaning that
25 `MPI_INIT` or `MPI_INIT_THREAD` has been used to initialize MPI.

27 11.7.1 Starting Processes

28 MPI applications may start new processes through an interface to an external process man-
29 ager.

30 `MPI_COMM_SPAWN` starts MPI processes and establishes communication with them,
31 returning an inter-communicator. `MPI_COMM_SPAWN_MULTIPLE` starts several different
32 binaries (or the same binary with different arguments), placing them in the same
33 `MPI_COMM_WORLD` and returning an inter-communicator.

34 MPI uses the group abstraction to represent processes. A process is identified by a
35 (group, rank) pair.

37 11.7.2 The Runtime Environment

38
39 The `MPI_COMM_SPAWN` and `MPI_COMM_SPAWN_MULTIPLE` routines provide an inter-
40 face between MPI and the *runtime environment* of an MPI application. The difficulty is
41 that there is an enormous range of runtime environments and application requirements, and
42 MPI must not be tailored to any particular one.

43 MPI assumes, implicitly, the existence of an environment in which an application runs.
44 It does not provide “operating system” services, such as a general ability to query what
45 processes are running, to kill arbitrary processes, to find out properties of the runtime
46 environment (how many processors, how much memory, etc.). Complex interaction of an
47 MPI application with its runtime environment should be done through an environment-
48 specific API.

At some low level, MPI must be able to interact with the runtime system, but the interaction is not visible at the application level and the details of the interaction are not specified by the MPI standard.

In many cases, it is impossible to keep environment-specific information out of the MPI interface without seriously compromising MPI functionality. To permit applications to take advantage of environment-specific functionality, many MPI routines take an `info` argument that allows an application to specify environment-specific information. There is a tradeoff between functionality and portability: applications that make use of environment-specific `info` are not portable.

MPI does not require the existence of an underlying “virtual machine” model, in which there is a consistent global view of an MPI application and an implicit “operating system” managing resources and processes. For instance, MPI processes spawned by one MPI process may not be visible to another; additional hosts added to the runtime environment by one MPI process may not be visible in another MPI process; MPI processes spawned by different processes may not be automatically distributed over available resources.

Interaction between MPI and the runtime environment is limited to the following areas:

- A process may start new processes with `MPI_COMM_SPAWN` and `MPI_COMM_SPAWN_MULTIPLE`.
- When a process spawns a child process, it may optionally use an `info` argument to tell the runtime environment where or how to start the process. This extra information may be opaque to MPI.
- An attribute `MPI_UNIVERSE_SIZE` (See Section 11.10.1) on `MPI_COMM_WORLD` tells a program how “large” the initial runtime environment is, namely how many processes can usefully be started in all. One can subtract the size of `MPI_COMM_WORLD` from this value to find out how many processes might usefully be started in addition to those already running.

11.8 Process Manager Interface

11.8.1 Processes in MPI

A process is represented in MPI by a (group, rank) pair. A (group, rank) pair specifies a unique process but a process does not determine a unique (group, rank) pair, since a process may belong to several groups.

11.8.2 Starting Processes and Establishing Communication

The following routine starts a number of MPI processes and establishes communication with them, returning an inter-communicator.

Advice to users. It is possible in MPI to start an SPMD or MPMD application with a fixed number of processes after initialization by first starting one process and having that process start its siblings with `MPI_COMM_SPAWN`. This practice is discouraged primarily for reasons of performance. If possible, it is preferable to start all processes at once, as a single MPI application. (*End of advice to users.*)

```

1 MPI_COMM Spawn(command, argv, maxprocs, info, root, comm, intercomm,
2     array_of_errcodes)
3
4     IN     command     name of program to be spawned (string, significant
5                       only at root)
6
7     IN     argv        arguments to command (array of strings, significant
8                       only at root)
9
10    IN     maxprocs    maximum number of processes to start (integer,
11                       significant only at root)
12
13    IN     info        a set of key-value pairs telling the runtime system
14                       where and how to start the processes (handle,
15                       significant only at root)
16
17    IN     comm        intra-communicator containing group of spawning
18                       processes (handle)
19
20    OUT    intercomm    inter-communicator between original group and the
21                       newly spawned group (handle)
22
23    OUT    array_of_errcodes one code per process (array of integers)

```

C binding

```

24 int MPI_Comm_spawn(const char *command, char *argv[], int maxprocs,
25     MPI_Info info, int root, MPI_Comm comm, MPI_Comm *intercomm,
26     int array_of_errcodes[])

```

Fortran 2008 binding

```

27 MPI_Comm_spawn(command, argv, maxprocs, info, root, comm, intercomm,
28     array_of_errcodes, ierror)
29
30 CHARACTER(LEN=*) , INTENT(IN) :: command, argv(*)
31 INTEGER, INTENT(IN) :: maxprocs, root
32 TYPE(MPI_Info), INTENT(IN) :: info
33 TYPE(MPI_Comm), INTENT(IN) :: comm
34 TYPE(MPI_Comm), INTENT(OUT) :: intercomm
35 INTEGER :: array_of_errcodes(*)
36 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

37
38 MPI_COMM_SPAWN(COMMAND, ARGV, MAXPROCS, INFO, ROOT, COMM, INTERCOMM,
39     ARRAY_OF_ERRCODES, IERROR)
40
41 CHARACTER*(*) COMMAND, ARGV(*)
42 INTEGER MAXPROCS, INFO, ROOT, COMM, INTERCOMM, ARRAY_OF_ERRCODES(*), IERROR

```

MPI_COMM_SPAWN tries to start maxprocs identical copies of the MPI program specified by command, establishing communication with them and returning an inter-communicator. The spawned processes are referred to as children. The children have their own MPI_COMM_WORLD, which is separate from that of the parents. MPI_COMM_SPAWN is collective over comm, and also may not return until MPI_INIT has been called in the children. Similarly, MPI_INIT in the children may not return until all parents have called

MPI_COMM_SPAWN. In this sense, MPI_COMM_SPAWN in the parents and MPI_INIT in the children form a collective operation over the union of parent and child processes. The inter-communicator returned by MPI_COMM_SPAWN contains the parent processes in the local group and the child processes in the remote group. The ordering of processes in the local and remote groups is the same as the ordering of the group of the comm in the parents and of MPI_COMM_WORLD of the children, respectively. This inter-communicator can be obtained in the children through the function MPI_COMM_GET_PARENT.

Advice to users. An implementation may automatically establish communication before MPI_INIT is called by the children. Thus, completion of MPI_COMM_SPAWN in the parent does not necessarily mean that MPI_INIT has been called in the children (although the returned inter-communicator can be used immediately). (*End of advice to users.*)

The arguments are:

command: The `command` argument is a string containing the name of a program to be spawned. The string is null-terminated in C. In Fortran, leading and trailing spaces are stripped. MPI does not specify how to find the executable or how the working directory is determined. These rules are implementation-dependent and should be appropriate for the runtime environment.

Advice to implementors. The implementation should use a natural rule for finding executables and determining working directories. For instance, a homogeneous system with a global file system might look first in the working directory of the spawning process, or might search the directories in a PATH environment variable as do Unix shells. An implementation should document its rules for finding executables and determining working directories, and a high-quality implementation should give the user some control over these rules. (*End of advice to implementors.*)

If the program named in `command` does not call MPI_INIT, but instead forks a process that calls MPI_INIT, the results are undefined. Implementations may allow this case to work but are not required to.

Advice to users. MPI does not say what happens if the program you start is a shell script and that shell script starts a program that calls MPI_INIT. Though some implementations may allow you to do this, they may also have restrictions, such as requiring that arguments supplied to the shell script be supplied to the program, or requiring that certain parts of the environment not be changed. (*End of advice to users.*)

argv: `argv` is an array of strings containing arguments that are passed to the program. The first element of `argv` is the first argument passed to `command`, not, as is conventional in some contexts, the command itself. The argument list is terminated by NULL in C and an empty string in Fortran. In Fortran, leading and trailing spaces are always stripped, so that a string consisting of all spaces is considered an empty string. The constant MPI_ARGV_NULL may be used in C and Fortran to indicate an empty argument list. In C this constant is the same as NULL.

Example 11.20. Examples of `argv` in C and Fortran

To run the program “ocean” with arguments “-gridfile” and “ocean1.grd” in C:

```
char command[] = "ocean";
char *argv[] = {"-gridfile", "ocean1.grd", NULL};
MPI_Comm_spawn(command, argv, ...);
```

or, if not everything is known at compile time:

```
char *command;
char **argv;
command = "ocean";
argv=(char **)malloc(3 * sizeof(char *));
argv[0] = "-gridfile";
argv[1] = "ocean1.grd";
argv[2] = NULL;
MPI_Comm_spawn(command, argv, ...);
```

In Fortran:

```
CHARACTER*25 command, argv(3)
command = 'ocean'
argv(1) = '-gridfile'
argv(2) = 'ocean1.grd'
argv(3) = ' '
call MPI_COMM_SPAWN(command, argv, ...)
```

Arguments are supplied to the program if this is allowed by the operating system. In C, the `MPI_COMM_SPAWN` argument `argv` differs from the `argv` argument of `main` in two respects. First, it is shifted by one element. Specifically, `argv[0]` of `main` is provided by the implementation and conventionally contains the name of the program (given by `command`). `argv[1]` of `main` corresponds to `argv[0]` in `MPI_COMM_SPAWN`, `argv[2]` of `main` to `argv[1]` of `MPI_COMM_SPAWN`, etc. Passing an `argv` of `MPI_ARGV_NULL` to `MPI_COMM_SPAWN` results in `main` receiving `argc` of 1 and an `argv` whose element 0 is (conventionally) the name of the program. Second, `argv` of `MPI_COMM_SPAWN` must be null-terminated, so that its length can be determined.

If a Fortran implementation supplies routines that allow a program to obtain its arguments, the arguments may be available through that mechanism. In C, if the operating system does not support arguments appearing in `argv` of `main()`, the MPI implementation may add the arguments to the `argv` that is passed to `MPI_INIT`.

maxprocs: MPI tries to spawn `maxprocs` processes. If it is unable to spawn `maxprocs` processes, it raises an error of class `MPI_ERR_SPAWN`.

An implementation may allow the `info` argument to change the default behavior, such that if the implementation is unable to spawn all `maxprocs` processes, it may spawn a smaller number of processes instead of raising an error. In principle, the `info` argument may specify an arbitrary set $\{m_i : 0 \leq m_i \leq \text{maxprocs}\}$ of allowed values for the number of processes spawned. The set $\{m_i\}$ does not necessarily include the value `maxprocs`. If an implementation is able to spawn one of these allowed numbers of processes, `MPI_COMM_SPAWN` returns successfully and the number of spawned processes, m , is given by the size of the remote group of `intercomm`. If m is less than `maxproc`, reasons why the other processes were not spawned are given in

array_of_errcodes as described below. If it is not possible to spawn one of the allowed numbers of processes, MPI_COMM_SPAWN raises an error of class MPI_ERR_SPAWN.

A spawn call with the default behavior is called *hard*. A spawn call for which fewer than maxprocs processes may be returned is called "soft". See Section 11.8.4 for more information on the "soft" key for info.

Advice to users. By default, requests are hard and MPI errors are fatal. This means that by default there will be a fatal error if MPI cannot spawn all the requested processes. If you want the behavior "spawn as many processes as possible, up to N ," you should do a soft spawn, where the set of allowed values $\{m_i\}$ is $\{0, \dots, N\}$. However, this is not completely portable, as implementations are not required to support soft spawning. (*End of advice to users.*)

info: The info argument to all of the routines in this chapter is an opaque handle of type MPI_Info in C and Fortran with the mpi_f08 module and INTEGER in Fortran with the mpi module or the include file mpif.h (deprecated). It is a container for a number of user-specified (key,value) pairs. key and value are strings (null-terminated char* in C, character*(*) in Fortran). Routines to create and manipulate the info argument are described in Chapter 10.

For the SPAWN calls, info provides additional (and possibly implementation-dependent) instructions to MPI and the runtime system on how to start processes. An application may pass MPI_INFO_NULL in C or Fortran. Portable programs not requiring detailed control over process locations should use MPI_INFO_NULL.

MPI does not specify the content of the info argument, except to reserve a number of special key values (see Section 11.8.4). The info argument is quite flexible and could even be used, for example, to specify the executable and its command-line arguments. In this case the command argument to MPI_COMM_SPAWN could be empty. The ability to do this follows from the fact that MPI does not specify how an executable is found, and the info argument can tell the runtime system where to "find" the executable "" (empty string). Of course, a program that does this will not be portable across MPI implementations.

root: All arguments before the root argument are examined only on the process whose rank in comm is equal to root. The value of these arguments on other processes is ignored.

array_of_errcodes: The array_of_errcodes is an array of length maxprocs in which MPI reports the status of each process that MPI was requested to start. If all maxprocs processes were spawned, array_of_errcodes is filled in with the value MPI_SUCCESS. If only m ($0 \leq m < \text{maxprocs}$) processes are spawned, m of the entries will contain MPI_SUCCESS and the rest will contain an implementation-specific error code indicating the reason MPI could not start the process. MPI does not specify which entries correspond to failed processes. An implementation may, for instance, fill in error codes in one-to-one correspondence with a detailed specification in the info argument. These error codes all belong to the error class MPI_ERR_SPAWN if there was no error in the argument list. In C or Fortran, an application may pass MPI_ERRCODES_IGNORE if it is not interested in the error codes.

Advice to implementors. MPI_ERRCODES_IGNORE in Fortran is a special type of constant, like MPI_BOTTOM. See the discussion in Section 2.5.4. (*End of advice to implementors.*)

1 MPI_COMM_GET_PARENT(parent)

2 OUT parent the parent communicator (handle)

4 **C binding**

5 int MPI_Comm_get_parent(MPI_Comm *parent)

7 **Fortran 2008 binding**

8 MPI_Comm_get_parent(parent, ierror)

9 TYPE(MPI_Comm), INTENT(OUT) :: parent

10 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

11 **Fortran binding**

12 MPI_COMM_GET_PARENT(PARENT, IERROR)

13 INTEGER PARENT, IERROR

14
15 If a process was started with MPI_COMM_SPAWN or MPI_COMM_SPAWN_MULTIPLE,
16 MPI_COMM_GET_PARENT returns the “parent” inter-communicator of the current pro-
17 cess. This parent inter-communicator is created implicitly inside of MPI_INIT and is the
18 same inter-communicator returned by SPAWN in the parents.

19 If the process was not spawned, MPI_COMM_GET_PARENT returns MPI_COMM_NULL.

20 After the parent communicator is freed or disconnected, MPI_COMM_GET_PARENT
21 returns MPI_COMM_NULL.

22 *Advice to users.* MPI_COMM_GET_PARENT returns a handle to a single inter-
23 communicator. Calling MPI_COMM_GET_PARENT a second time returns a handle
24 to the same inter-communicator. Freeing the handle with MPI_COMM_DISCONNECT
25 or MPI_COMM_FREE will cause other references to the inter-communicator to become
26 invalid (dangling). Note that calling MPI_COMM_FREE on the parent communicator
27 is not useful. (*End of advice to users.*)

28
29 *Rationale.* The desire of the Forum was to create a constant
30 MPI_COMM_PARENT similar to MPI_COMM_WORLD. Unfortunately such a constant
31 cannot be used (syntactically) as an argument to MPI_COMM_DISCONNECT, which
32 is explicitly allowed. (*End of rationale.*)

34 11.8.3 Starting Multiple Executables and Establishing Communication

35 While MPI_COMM_SPAWN is sufficient for most cases, it does not allow the spawning of
36 multiple binaries, or of the same binary with multiple sets of arguments. The following
37 routine spawns multiple binaries or the same binary with multiple sets of arguments, estab-
38 lishing communication with them and placing them in the same MPI_COMM_WORLD.

41 MPI_COMM_SPAWN_MULTIPLE(count, array_of_commands, array_of_argv,
42 array_of_maxprocs, array_of_info, root, comm, intercomm,
43 array_of_errcodes)

44
45 IN count number of commands (positive integer, significant
46 only at root)

47 IN array_of_commands programs to be executed (array of strings, significant
48 only at root)

IN	array_of_argv	arguments for commands (array of array of strings, significant only at root)	1 2
IN	array_of_maxprocs	maximum number of processes to start for each command (array of integers, significant only at root)	3 4 5
IN	array_of_info	info objects telling the runtime system where and how to start processes (array of handles, significant only at root)	6 7 8
IN	root	rank of process in which previous arguments are examined (integer)	9 10
IN	comm	intra-communicator containing group of spawning processes (handle)	11 12 13
OUT	intercomm	inter-communicator between original group and the newly spawned group (handle)	14 15
OUT	array_of_errcodes	one error code per process (array of integers)	16 17
C binding			18
<pre>int MPI_Comm_spawn_multiple(int count, char *array_of_commands[], char **array_of_argv[], const int array_of_maxprocs[], const MPI_Info array_of_info[], int root, MPI_Comm comm, MPI_Comm *intercomm, int array_of_errcodes[])</pre>			19 20 21 22 23
Fortran 2008 binding			24
<pre>MPI_Comm_spawn_multiple(count, array_of_commands, array_of_argv, array_of_maxprocs, array_of_info, root, comm, intercomm, array_of_errcodes, ierror) INTEGER, INTENT(IN) :: count, array_of_maxprocs(*), root CHARACTER(LEN=*) , INTENT(IN) :: array_of_commands(*), array_of_argv(count, *) TYPE(MPI_Info), INTENT(IN) :: array_of_info(*) TYPE(MPI_Comm), INTENT(IN) :: comm TYPE(MPI_Comm), INTENT(OUT) :: intercomm INTEGER :: array_of_errcodes(*) INTEGER, OPTIONAL, INTENT(OUT) :: ierror</pre>			25 26 27 28 29 30 31 32 33 34 35
Fortran binding			36
<pre>MPI_COMM_SPAWN_MULTIPLE(COUNT, ARRAY_OF_COMMANDS, ARRAY_OF_ARGV, ARRAY_OF_MAXPROCS, ARRAY_OF_INFO, ROOT, COMM, INTERCOMM, ARRAY_OF_ERRCODES, IERROR) INTEGER COUNT, ARRAY_OF_MAXPROCS(*), ARRAY_OF_INFO(*), ROOT, COMM, INTERCOMM, ARRAY_OF_ERRCODES(*), IERROR CHARACTER*(*) ARRAY_OF_COMMANDS(*), ARRAY_OF_ARGV(COUNT, *)</pre>			37 38 39 40 41 42 43
<p>MPI_COMM_SPAWN_MULTIPLE is identical to MPI_COMM_SPAWN except that there are multiple executable specifications. The first argument, count, gives the number of specifications. Each of the next four arguments are simply arrays of the corresponding arguments in MPI_COMM_SPAWN. For the Fortran version of array_of_argv, the element array_of_argv(i,j) is the j-th argument to command number i.</p>			44 45 46 47 48

Rationale. This may seem backwards to Fortran programmers who are familiar with Fortran's column-major ordering. However, it is necessary to do it this way to allow `MPI_COMM_SPAWN` to sort out arguments. Note that the leading dimension of `array_of_argv` *must* be the same as `count`. Also note that Fortran rules for sequence association allow a different value in the first dimension; in this case, the sequence of array elements is interpreted by `MPI_COMM_SPAWN_MULTIPLE` as if the sequence is stored in an array defined with the first dimension set to `count`. This Fortran feature allows an implementor to define `MPI_ARGVS_NULL` (see below) with fixed dimensions, e.g., (1,1), or only with one dimension, e.g., (1). (*End of rationale.*)

Advice to users. The argument `count` is interpreted by MPI only at the root, as is `array_of_argv`. Since the leading dimension of `array_of_argv` is `count`, a nonpositive value of `count` at a nonroot node could theoretically cause a runtime bounds check error, even though `array_of_argv` should be ignored by the subroutine. If this happens, you should explicitly supply a reasonable value of `count` on the nonroot nodes. (*End of advice to users.*)

In any language, an application may use the constant `MPI_ARGVS_NULL` (which is likely to be `(char ***)0` in C) to specify that no arguments should be passed to any commands. The effect of setting individual elements of `array_of_argv` to `MPI_ARGV_NULL` is not defined. To specify arguments for some commands but not others, the commands without arguments should have a corresponding `argv` whose first element is null (`(char *)0` in C and empty string in Fortran). In Fortran at nonroot processes, the `count` argument must be set to a value that is consistent with the provided `array_of_argv` although the content of these arguments has no meaning for this operation.

All of the spawned processes have the same `MPI_COMM_WORLD`. Their ranks in `MPI_COMM_WORLD` correspond directly to the order in which the commands are specified in `MPI_COMM_SPAWN_MULTIPLE`. Assume that m_1 MPI processes are generated by the first command, m_2 by the second, etc. The MPI processes corresponding to the first command have ranks $0, 1, \dots, m_1 - 1$ in `MPI_COMM_WORLD`. The MPI processes in the second command have ranks $m_1, m_1 + 1, \dots, m_1 + m_2 - 1$ in `MPI_COMM_WORLD`. The MPI processes in the third have ranks $m_1 + m_2, m_1 + m_2 + 1, \dots, m_1 + m_2 + m_3 - 1$ in `MPI_COMM_WORLD`, etc.

Advice to users. Calling `MPI_COMM_SPAWN` multiple times would create many sets of children with different `MPI_COMM_WORLD`s whereas `MPI_COMM_SPAWN_MULTIPLE` creates children with a single `MPI_COMM_WORLD`, so the two methods are not completely equivalent. There are also two performance-related reasons why, if you need to spawn multiple executables, you may want to use `MPI_COMM_SPAWN_MULTIPLE` instead of calling `MPI_COMM_SPAWN` several times. First, spawning several things at once may be faster than spawning them sequentially. Second, in some implementations, communication between processes spawned at the same time may be faster than communication between processes spawned separately. (*End of advice to users.*)

The `array_of_errcodes` argument is a 1-dimensional array of size $\sum_{i=1}^{\text{count}} n_i$, where n_i is the i -th element of `array_of_maxprocs`. Command number i corresponds to the n_i contiguous slots in this array from element $\sum_{j=1}^{i-1} n_j$ to $\left[\sum_{j=1}^i n_j\right] - 1$. Error codes are treated the same as with `MPI_COMM_SPAWN`.

Example 11.21. Examples of `array_of_argv` in C and Fortran

To run the program “ocean” with arguments “-gridfile” and “ocean1.grd” and the program “atmos” with argument “atmos.grd” in C:

```
char *array_of_commands[2] = {"ocean", "atmos"};
char **array_of_argv[2];
char *argv0[] = {"-gridfile", "ocean1.grd", (char *)0};
char *argv1[] = {"atmos.grd", (char *)0};
array_of_argv[0] = argv0;
array_of_argv[1] = argv1;
MPI_Comm_spawn_multiple(2, array_of_commands, array_of_argv, ...);
```

Here is how you do it in Fortran:

```
CHARACTER*25 commands(2), array_of_argv(2, 3)
commands(1) = 'ocean'
array_of_argv(1, 1) = '-gridfile'
array_of_argv(1, 2) = 'ocean1.grd'
array_of_argv(1, 3) = ' '

commands(2) = 'atmos'
array_of_argv(2, 1) = 'atmos.grd'
array_of_argv(2, 2) = ' '

call MPI_COMM_SPAWN_MULTIPLE(2, commands, array_of_argv, ...)
```

11.8.4 Reserved Keys

The following keys are reserved. An implementation is not required to interpret these keys, but if it does interpret the key, it must provide the functionality described.

"host": Value is a hostname. The format of the hostname is determined by the implementation.

"arch": Value is an architecture name. Valid architecture names and what they mean are determined by the implementation.

"wdir": Value is the name of a directory on a machine on which the spawned process(es) execute(s). This directory is made the working directory of the executing process(es). The format of the directory name is determined by the implementation.

"path": Value is a directory or set of directories where the implementation should look for the executable. The format of "path" is determined by the implementation.

"file": Value is the name of a file in which additional information is specified. The format of the filename and internal format of the file are determined by the implementation.

"mpi_initial_errhandler": Value is the name of an errhandler that will be set as the initial error handler. The "mpi_initial_errhandler" key can take the case insensitive values "mpi_errors_are_fatal", "mpi_errors_abort", and "mpi_errors_return" representing the pre-defined MPI error handlers (MPI_ERRORS_ARE_FATAL—the default, MPI_ERRORS_ABORT, and MPI_ERRORS_RETURN, respectively). Other, nonstandard

values may be supported by the implementation, which should document the resultant behavior.

"mpi_memory_alloc_kinds" (string, default: "mpi,system"): Value is a comma separated list of memory allocation kinds. Support for these memory allocation kinds is requested from the MPI library (see Section 11.4.3).

"soft": Value specifies a set of numbers that are allowed values for the number of processes that MPI_COMM_SPAWN (et al.) may create. The format of the value is a comma-separated list of Fortran-90 triplets each of which specifies a set of integers and that together specify the set formed by the union of these sets. Negative values in this set and values greater than `maxprocs` are ignored. MPI will spawn the largest number of processes it can, consistent with some number in the set. The order in which triplets are given is not significant.

By Fortran-90 triplets, we mean:

1. `a` means a
2. `a:b` means $a, a + 1, a + 2, \dots, b$
3. `a:b:c` means $a, a + c, a + 2c, \dots, a + ck$, where for $c > 0$, k is the largest integer for which $a + ck \leq b$ and for $c < 0$, k is the largest integer for which $a + ck \geq b$. If $b > a$ then c must be positive. If $b < a$ then c must be negative.

Examples:

1. `a:b` gives a range between a and b
2. `0:N` gives full "soft" functionality
3. `1,2,4,8,16,32,64,128,256,512,1024,2048,4096` allows a power-of-two number of processes.
4. `2:10000:2` allows an even number of processes up to a maximum of 10,000 processes.
5. `2:10:2,7` allows 2, 4, 6, 7, 8, or 10 processes.

11.8.5 Spawn Example

Example 11.22. Manager-worker Example Using MPI_COMM_SPAWN

```

/* manager */
#include <stdio.h>
#include "mpi.h"
int main(int argc, char *argv[])
{
    int world_size, universe_size, *universe_sizep, flag;
    MPI_Comm everyone;          /* inter-communicator */
    char worker_program[100];

    MPI_Init(&argc, &argv);
    MPI_Comm_size(MPI_COMM_WORLD, &world_size);

    if (world_size != 1)    error("Top heavy with management");

```

```

1  MPI_Comm_get_attr(MPI_COMM_WORLD, MPI_UNIVERSE_SIZE,
2  &universe_sizep, &flag);
3
4  if (!flag) {
5      printf("This MPI does not support UNIVERSE_SIZE. How many\n\
6  processes total?");
7      scanf("%d", &universe_size);
8  } else universe_size = *universe_sizep;
9
10 MPI_Comm_spawn(MPI_COMM_WORLD, MPI_ARGV_NULL, universe_size,
11 MPI_INFO_NULL, 0, MPI_COMM_SELF, &everyone,
12 MPI_ERRCODES_IGNORE);
13
14 /*
15  * Now spawn the workers. Note that there is a run-time determination
16  * of what type of worker to spawn, and presumably this calculation
17  * must be done at run time and cannot be calculated before starting
18  * the program. If everything is known when the application is
19  * first started, it is generally better to start them all at once
20  * in a single MPI_COMM_WORLD.
21  */
22
23 choose_worker_program(worker_program);
24 MPI_Comm_spawn(worker_program, MPI_ARGV_NULL, universe_size-1,
25 MPI_INFO_NULL, 0, MPI_COMM_SELF, &everyone,
26 MPI_ERRCODES_IGNORE);
27
28 /*
29  * Parallel code here. The communicator "everyone" can be used
30  * to communicate with the spawned processes, which have ranks 0,..
31  * MPI_UNIVERSE_SIZE-1 in the remote group of the inter-communicator
32  * "everyone".
33  */
34
35 MPI_Finalize();
36 return 0;
37 }
38
39 /* worker */
40
41 #include "mpi.h"
42 int main(int argc, char *argv[])
43 {
44     int size;
45     MPI_Comm parent;
46     MPI_Init(&argc, &argv);
47     MPI_Comm_get_parent(&parent);
48     if (parent == MPI_COMM_NULL) error("No parent!");
49     MPI_Comm_remote_size(parent, &size);
50     if (size != 1) error("Something's wrong with the parent");
51
52     /*
53      * Parallel code here.
54      * The manager is represented as the process with rank 0 in (the
55      * remote group of) the parent communicator. If the workers need
56      * to communicate among themselves, they can use MPI_COMM_WORLD.
57      */
58
59 }

```

```
1 MPI_Finalize();
2 return 0;
3 }
4
```

11.9 Establishing Communication

This section provides functions that establish communication between two sets of MPI processes that do not share a communicator.

Some situations in which these functions are useful are:

1. Two parts of an application that are started independently need to communicate.
2. A visualization tool wants to attach to a running process.
3. A server wants to accept connections from multiple clients. Both clients and server may be parallel programs.

In each of these situations, MPI must establish communication channels where none existed before, and there is no parent/child relationship. The routines described in this section establish communication between the two sets of processes by creating an MPI inter-communicator, where the two groups of the inter-communicator are the original sets of processes.

Establishing contact between two groups of processes that do not share an existing communicator is a collective but asymmetric process. One group of processes indicates its willingness to accept connections from other groups of processes. We will call this group the (parallel) *server*, even if this is not a client/server type of application. The other group connects to the server; we will call it the (parallel) *client*.

Advice to users. While the names *client* and *server* are used throughout this section, MPI does not guarantee the traditional robustness of client/server systems. The functionality described in this section is intended to allow two cooperating parts of the same application to communicate with one another. For instance, a client that gets a segmentation fault and dies, or one that does not participate in a collective operation may cause a server to crash or hang. (*End of advice to users.*)

11.9.1 Names, Addresses, Ports, and All That

Almost all of the complexity in MPI client/server routines addresses the question “how does the client find out how to contact the server?” The difficulty, of course, is that there is no existing communication channel between them, yet they must somehow agree on a rendezvous point where they will establish communication.

Agreeing on a rendezvous point always involves a third party. The third party may itself provide the rendezvous point or may communicate rendezvous information from server to client. Complicating matters might be the fact that it is not important to the client which particular server it contacts, only that it be able to get in touch with one that can handle its request.

Ideally, MPI can accommodate a wide variety of run-time systems while retaining the ability to write simple, portable code. The following should be compatible with MPI:

- The server resides at a well-known internet address host:port. 1
- The server prints out an address to the terminal; the user gives this address to the client program. 2
- The server places the address information on a nameserver, where it can be retrieved with an agreed-upon name. 3
- The server to which the client connects is actually a broker, acting as a middleman between the client and the real server. 4

MPI does not require a nameserver, so not all implementations will be able to support all of the above scenarios. However, MPI provides an optional nameserver interface, and is compatible with external name servers. 5

A `port_name` is a *system-supplied* string that encodes a low-level network address at which a server can be contacted. Typically this is an IP address and a port number, but an implementation is free to use any protocol. The server establishes a `port_name` with the `MPI_OPEN_PORT` routine. It accepts a connection to a given port with `MPI_COMM_ACCEPT`. A client uses `port_name` to connect to the server. 6

By itself, the `port_name` mechanism is completely portable, but it may be clumsy to use because of the necessity to communicate `port_name` to the client. It would be more convenient if a server could specify that it be known by an *application-supplied* `service_name` so that the client could connect to that `service_name` without knowing the `port_name`. 7

An MPI implementation may allow the server to publish a (`port_name`, `service_name`) pair with `MPI_PUBLISH_NAME` and the client to retrieve the port name from the service name with `MPI_LOOKUP_NAME`. This allows three levels of portability, with increasing levels of functionality. 8

1. Applications that do not rely on the ability to publish names are the most portable. Typically the `port_name` must be transferred “by hand” from server to client. 9
2. Applications that use the `MPI_PUBLISH_NAME` mechanism are completely portable among implementations that provide this service. To be portable among all implementations, these applications should have a fall-back mechanism that can be used when names are not published. 10
3. Applications may ignore MPI’s name publishing functionality and use their own mechanism (possibly system-supplied) to publish names. This allows arbitrary flexibility but is not portable. 11

11.9.2 Server Routines 12

A server makes itself available with two routines. First it must call `MPI_OPEN_PORT` to establish a `port` at which it may be contacted. Secondly it must call `MPI_COMM_ACCEPT` to accept connections from clients. 13

1 MPI_OPEN_PORT(info, port_name)

2 IN info implementation-specific information on how to
3 establish an address (handle)
4
5 OUT port_name newly established port (string)

6 7 **C binding**

8 int MPI_Open_port(MPI_Info info, char *port_name)

9 10 **Fortran 2008 binding**

11 MPI_Open_port(info, port_name, ierror)
12 TYPE(MPI_Info), INTENT(IN) :: info
13 CHARACTER(LEN=MPI_MAX_PORT_NAME), INTENT(OUT) :: port_name
14 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

15 **Fortran binding**

16 MPI_OPEN_PORT(INFO, PORT_NAME, IERROR)
17 INTEGER INFO, IERROR
18 CHARACTER*(*) PORT_NAME

19 This function establishes a network address, encoded in the `port_name` string, at which
20 the server will be able to accept connections from clients. `port_name` is supplied by the
21 system, possibly using information in the `info` argument.

22 MPI copies a system-supplied port name into `port_name`. `port_name` identifies the newly
23 opened port and can be used by a client to contact the server. The maximum size of the
24 string that may be supplied by the system is `MPI_MAX_PORT_NAME`.

25
26 *Advice to users.* The system copies the port name into `port_name`. The application
27 must pass a buffer of sufficient size to hold this value. (*End of advice to users.*)

28
29 `port_name` is essentially a network address. It is unique within the communication
30 universe to which it belongs (determined by the implementation), and may be used by any
31 client within that communication universe. For instance, if it is an internet (host:port)
32 address, it will be unique on the internet. If it is a low level switch address on an IBM SP,
33 it will be unique to that SP.

34 *Advice to implementors.* These examples are not meant to constrain implementa-
35 tions. A `port_name` could, for instance, contain a user name or the name of a batch
36 job, as long as it is unique within some well-defined communication domain. The
37 larger the communication domain, the more useful MPI's client/server functionality
38 will be. (*End of advice to implementors.*)

39
40 The precise form of the address is implementation-defined. For instance, an internet address
41 may be a host name or IP address, or anything that the implementation can decode into
42 an IP address. A port name may be reused after it is freed with `MPI_CLOSE_PORT` and
43 released by the system.

44
45 *Advice to implementors.* Since the user may type in `port_name` by hand, it is useful
46 to choose a form that is easily readable and does not have embedded spaces. (*End of*
47 *advice to implementors.*)

info may be used to tell the implementation how to establish the address. It may, and usually will, be MPI_INFO_NULL in order to get the implementation defaults.

MPI_CLOSE_PORT(port_name)

IN port_name a port (string)

C binding

```
int MPI_Close_port(const char *port_name)
```

Fortran 2008 binding

```
MPI_Close_port(port_name, ierror)
  CHARACTER(LEN=*), INTENT(IN) :: port_name
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_CLOSE_PORT(PORT_NAME, IERROR)
  CHARACTER*(*) PORT_NAME
  INTEGER IERROR
```

This function releases the network address represented by port_name.

MPI_COMM_ACCEPT(port_name, info, root, comm, newcomm)

IN port_name port name (string, significant only at root)

IN info implementation-dependent information (handle, significant only at root)

IN root rank of root in comm (integer)

IN comm intra-communicator over which call is collective (handle)

OUT newcomm inter-communicator with client as remote group (handle)

C binding

```
int MPI_Comm_accept(const char *port_name, MPI_Info info, int root,
  MPI_Comm comm, MPI_Comm *newcomm)
```

Fortran 2008 binding

```
MPI_Comm_accept(port_name, info, root, comm, newcomm, ierror)
  CHARACTER(LEN=*), INTENT(IN) :: port_name
  TYPE(MPI_Info), INTENT(IN) :: info
  INTEGER, INTENT(IN) :: root
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Comm), INTENT(OUT) :: newcomm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_ACCEPT(PORT_NAME, INFO, ROOT, COMM, NEWCOMM, IERROR)
  CHARACTER*(*) PORT_NAME
```

1 INTEGER INFO, ROOT, COMM, NEWCOMM, IERROR

2 MPI_COMM_ACCEPT establishes communication with a client. It is collective over the
3 calling communicator. It returns an inter-communicator that allows communication with
4 the client.

5 The `port_name` must have been established through a call to `MPI_OPEN_PORT`.

6 `info` can be used to provide directives that may influence the behavior of the `ACCEPT`
7 call.
8

9 11.9.3 Client Routines

10 There is only one routine on the client side.

11 MPI_COMM_CONNECT(`port_name`, `info`, `root`, `comm`, `newcomm`)

12	IN	<code>port_name</code>	network address (string, significant only at root)
13	IN	<code>info</code>	implementation-dependent information (handle, 14 significant only at root)
15	IN	<code>root</code>	rank of root in <code>comm</code> (integer)
16	IN	<code>comm</code>	intra-communicator over which call is collective 17 (handle)
17	OUT	<code>newcomm</code>	inter-communicator with server as remote group 18 (handle)

19 C binding

20 int MPI_Comm_connect(const char *port_name, MPI_Info info, int root,
21 MPI_Comm comm, MPI_Comm *newcomm)

22 Fortran 2008 binding

23 MPI_Comm_connect(`port_name`, `info`, `root`, `comm`, `newcomm`, `ierror`)
24 CHARACTER(LEN=*) , INTENT(IN) :: `port_name`
25 TYPE(MPI_Info) , INTENT(IN) :: `info`
26 INTEGER , INTENT(IN) :: `root`
27 TYPE(MPI_Comm) , INTENT(IN) :: `comm`
28 TYPE(MPI_Comm) , INTENT(OUT) :: `newcomm`
29 INTEGER , OPTIONAL , INTENT(OUT) :: `ierror`

30 Fortran binding

31 MPI_COMM_CONNECT(PORT_NAME, INFO, ROOT, COMM, NEWCOMM, IERROR)
32 CHARACTER*(*) PORT_NAME
33 INTEGER INFO, ROOT, COMM, NEWCOMM, IERROR

34 This routine establishes communication with a server specified by `port_name`. It is
35 collective over the calling communicator and returns an inter-communicator in which the
36 remote group participated in an `MPI_COMM_ACCEPT`.

37 If the named port does not exist (or has been closed), `MPI_COMM_CONNECT` raises
38 an error of class `MPI_ERR_PORT`.

39 If the port exists, but does not have a pending `MPI_COMM_ACCEPT`, the connection
40 attempt will eventually time out after an implementation-defined time, or succeed when

the server calls `MPI_COMM_ACCEPT`. In the case of a time out, `MPI_COMM_CONNECT` raises an error of class `MPI_ERR_PORT`.

Advice to implementors. The time out period may be arbitrarily short or long. However, a high-quality implementation will try to queue connection attempts so that a server can handle simultaneous requests from several clients. A high-quality implementation may also provide a mechanism, through the `info` arguments to `MPI_OPEN_PORT`, `MPI_COMM_ACCEPT`, and/or `MPI_COMM_CONNECT`, for the user to control timeout and queuing behavior. (*End of advice to implementors.*)

MPI provides no guarantee of fairness in servicing connection attempts. That is, connection attempts are not necessarily satisfied in the order they were initiated and competition from other connection attempts may prevent a particular connection attempt from being satisfied.

`port_name` is the address of the server. It must be the same as the name returned by `MPI_OPEN_PORT` on the server. Some freedom is allowed here. If there are equivalent forms of `port_name`, an implementation may accept them as well. For instance, if `port_name` is `(hostname:port)`, an implementation may accept `(ip_address:port)` as well.

11.9.4 Name Publishing

The routines in this section provide a mechanism for publishing names. A `(service_name, port_name)` pair is published by the server, and may be retrieved by a client using the `service_name` only. An MPI implementation defines the *scope* of the `service_name`, that is, the domain over which the `service_name` can be retrieved. If the domain is the empty set, that is, if no client can retrieve the information, then we say that name publishing is not supported. Implementations should document how the scope is determined. High-quality implementations will give some control to users through the `info` arguments to name publishing functions. Examples are given in the descriptions of individual functions.

`MPI_PUBLISH_NAME(service_name, info, port_name)`

IN	<code>service_name</code>	a service name to associate with the port (string)
IN	<code>info</code>	implementation-specific information (handle)
IN	<code>port_name</code>	a port name (string)

C binding

```
int MPI_Publish_name(const char *service_name, MPI_Info info,
                    const char *port_name)
```

Fortran 2008 binding

```
MPI_Publish_name(service_name, info, port_name, ierror)
  CHARACTER(LEN=*) , INTENT(IN) :: service_name, port_name
  TYPE(MPI_Info), INTENT(IN) :: info
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_PUBLISH_NAME(SERVICE_NAME, INFO, PORT_NAME, IERROR)
  CHARACTER*(*) SERVICE_NAME, PORT_NAME
```

1 INTEGER INFO, IERROR

2 This routine publishes the pair (port_name, service_name) so that an application may
3 retrieve a system-supplied port_name using a well-known service_name.

4 The implementation must define the *scope* of a published service name, that is, the
5 domain over which the service name is unique, and conversely, the domain over which the
6 (port_name, service_name) pair may be retrieved. For instance, a service name may be
7 unique to a job (where job is defined by a distributed operating system or batch scheduler),
8 unique to a machine, or unique to a Kerberos realm. The scope may depend on the info
9 argument to MPI_PUBLISH_NAME.

10 MPI permits publishing more than one service_name for a single port_name. On the
11 other hand, if service_name has already been published within the scope determined by info,
12 the behavior of MPI_PUBLISH_NAME is undefined. An MPI implementation may, through
13 a mechanism in the info argument to MPI_PUBLISH_NAME, provide a way to allow multiple
14 servers with the same service in the same scope. In this case, an implementation-defined
15 policy will determine which of several port names is returned by MPI_LOOKUP_NAME.

16 Note that while service_name has a limited scope, determined by the implementation,
17 port_name always has global scope within the communication universe used by the imple-
18 mentation (i.e., it is globally unique).

19 port_name should be the name of a port established by MPI_OPEN_PORT and not yet
20 released by MPI_CLOSE_PORT. If it is not, the result is undefined.

21
22 *Advice to implementors.* In some cases, an MPI implementation may use a name
23 service that a user can also access directly. In this case, a name published by MPI
24 could easily conflict with a name published by a user. In order to avoid such conflicts,
25 MPI implementations should mangle service names so that they are unlikely to conflict
26 with user code that makes use of the same service. Such name mangling will of course
27 be completely transparent to the user.

28 The following situation is problematic but unavoidable, if we want to allow implemen-
29 tations to use nameservers. Suppose there are multiple instances of “ocean” running
30 on a machine. If the scope of a service name is confined to a job, then multiple
31 oceans can coexist. If an implementation provides site-wide scope, however, multiple
32 instances are not possible as all calls to MPI_PUBLISH_NAME after the first may fail.
33 There is no universal solution to this.

34 To handle these situations, a high-quality implementation should make it possible to
35 limit the domain over which names are published. (*End of advice to implementors.*)

36
37
38
39
40 MPI_UNPUBLISH_NAME(service_name, info, port_name)

41 IN service_name a service name (string)
42 IN info implementation-specific information (handle)
43 IN port_name a port name (string)

44 **C binding**

45
46 int MPI_Unpublish_name(const char *service_name, MPI_Info info,
47 const char *port_name)
48

Fortran 2008 binding

```

MPI_Unpublish_name(service_name, info, port_name, ierror)
  CHARACTER(LEN=*), INTENT(IN) :: service_name, port_name
  TYPE(MPI_Info), INTENT(IN) :: info
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_UNPUBLISH_NAME(SERVICE_NAME, INFO, PORT_NAME, IERROR)
  CHARACTER*(*) SERVICE_NAME, PORT_NAME
  INTEGER INFO, IERROR

```

This routine unpublishes a service name that has been previously published. Attempting to unpublish a name that has not been published or has already been unpublished is erroneous and is indicated by the error class `MPI_ERR_SERVICE`.

All published names must be unpublished before the corresponding port is closed and before the publishing process exits. The behavior of `MPI_UNPUBLISH_NAME` is implementation dependent when a process tries to unpublish a name that it did not publish.

If the `info` argument was used with `MPI_PUBLISH_NAME` to tell the implementation how to publish names, the implementation may require that `info` passed to `MPI_UNPUBLISH_NAME` contain information to tell the implementation how to unpublish a name.

```

MPI_LOOKUP_NAME(service_name, info, port_name)

```

IN	service_name	a service name (string)
IN	info	implementation-specific information (handle)
OUT	port_name	a port name (string)

C binding

```

int MPI_Lookup_name(const char *service_name, MPI_Info info, char *port_name)

```

Fortran 2008 binding

```

MPI_Lookup_name(service_name, info, port_name, ierror)
  CHARACTER(LEN=*), INTENT(IN) :: service_name
  TYPE(MPI_Info), INTENT(IN) :: info
  CHARACTER(LEN=MPI_MAX_PORT_NAME), INTENT(OUT) :: port_name
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_LOOKUP_NAME(SERVICE_NAME, INFO, PORT_NAME, IERROR)
  CHARACTER*(*) SERVICE_NAME, PORT_NAME
  INTEGER INFO, IERROR

```

This function retrieves a `port_name` published by `MPI_PUBLISH_NAME` with `service_name`. If `service_name` has not been published, it raises an error in the error class `MPI_ERR_NAME`. The application must supply a `port_name` buffer large enough to hold the largest possible port name (see discussion above under `MPI_OPEN_PORT`).

If an implementation allows multiple entries with the same `service_name` within the same scope, a particular `port_name` is chosen in a way determined by the implementation.

If the `info` argument was used with `MPI_PUBLISH_NAME` to tell the implementation how to publish names, a similar `info` argument may be required for `MPI_LOOKUP_NAME`.

11.9.5 Reserved Key Values

The following key values are reserved. An implementation is not required to interpret these key values, but if it does interpret the key value, it must provide the functionality described.

"ip_port": Value contains IP port number at which to establish a `port`. (Reserved for `MPI_OPEN_PORT` only).

"ip_address": Value contains IP address at which to establish a `port`. If the address is not a valid IP address of the host on which the `MPI_OPEN_PORT` call is made, the results are undefined. (Reserved for `MPI_OPEN_PORT` only).

11.9.6 Client/Server Examples

Example 11.23. Simplest Example—Completely Portable.

The following example shows the simplest way to use the client/server interface. It does not use service names at all.

On the server side:

```
char myport[MPI_MAX_PORT_NAME];
MPI_Comm intercomm;
/* ... */
MPI_Open_port(MPI_INFO_NULL, myport);
printf("port name is: %s\n", myport);

MPI_Comm_accept(myport, MPI_INFO_NULL, 0, MPI_COMM_SELF, &intercomm);
/* do something with intercomm */
```

The server prints out the port name to the terminal and the user must type it in when starting up the client (assuming the MPI implementation supports `stdin` such that this works). On the client side:

```
MPI_Comm intercomm;
char name[MPI_MAX_PORT_NAME];
printf("enter port name: ");
gets(name);
MPI_Comm_connect(name, MPI_INFO_NULL, 0, MPI_COMM_SELF, &intercomm);
```

Example 11.24. Ocean/Atmosphere—Relies on Name Publishing

In this example, the “ocean” application is the “server” side of a coupled ocean-atmosphere climate model. It assumes that the MPI implementation publishes names.

```
char port_name[MPI_MAX_PORT_NAME];
MPI_Comm intercomm;
/* ... */
MPI_Open_port(MPI_INFO_NULL, port_name);
MPI_Publish_name("ocean", MPI_INFO_NULL, port_name);

MPI_Comm_accept(port_name, MPI_INFO_NULL, 0, MPI_COMM_SELF,
```

```

        &intercomm);
/* do something with intercomm */
MPI_Unpublish_name("ocean", MPI_INFO_NULL, port_name);

```

On the client side:

```

MPI_Lookup_name("ocean", MPI_INFO_NULL, port_name);
MPI_Comm_connect(port_name, MPI_INFO_NULL, 0, MPI_COMM_SELF,
                 &intercomm);

```

Example 11.25. Simple Client-Server Example

This is a simple example; the server accepts only a single connection at a time and serves that connection until the client requests to be disconnected. The server is a single process. Here is the server. It accepts a single connection and then processes data until it receives a message with tag 1. A message with tag 0 tells the server to exit.

```

#include "mpi.h"
int main(int argc, char *argv[])
{
    MPI_Comm client;
    MPI_Status status;
    char port_name[MPI_MAX_PORT_NAME];
    double buf[MAX_DATA];
    int size, again;

    MPI_Init(&argc, &argv);
    MPI_Comm_size(MPI_COMM_WORLD, &size);
    if (size != 1) error(FATAL, "Server too big");
    MPI_Open_port(MPI_INFO_NULL, port_name);
    printf("server available at %s\n", port_name);
    while (1) {
        MPI_Comm_accept(port_name, MPI_INFO_NULL, 0, MPI_COMM_WORLD,
                       &client);

        again = 1;
        while (again) {
            MPI_Recv(buf, MAX_DATA, MPI_DOUBLE,
                    MPI_ANY_SOURCE, MPI_ANY_TAG, client, &status);
            switch (status.MPI_TAG) {
                case 0: MPI_Comm_free(&client);
                       MPI_Close_port(port_name);
                       MPI_Finalize();
                       return 0;
                case 1: MPI_Comm_disconnect(&client);
                       again = 0;
                       break;
                case 2: /* do something */
                       ...
                default:
                       /* Unexpected message type */
                       MPI_Abort(MPI_COMM_WORLD, 1);
            }
        }
    }
}

```

```

1     }
2   }
3
4   Here is the client.
5
6   #include "mpi.h"
7   int main(int argc, char *argv[])
8   {
9     MPI_Comm server;
10    int done = 0;
11    double buf[MAX_DATA];
12    char port_name[MPI_MAX_PORT_NAME];
13
14    MPI_Init(&argc, &argv);
15    strcpy(port_name, argv[1]); /* assume server's name is cmd-line arg */
16
17    MPI_Comm_connect(port_name, MPI_INFO_NULL, 0, MPI_COMM_WORLD,
18                    &server);
19
20    while (!done) {
21      tag = 2; /* Action to perform */
22      MPI_Send(buf, n, MPI_DOUBLE, 0, tag, server);
23      /* etc */
24    }
25    MPI_Send(buf, 0, MPI_DOUBLE, 0, 1, server);
26    MPI_Comm_disconnect(&server);
27    MPI_Finalize();
28    return 0;
29  }

```

11.10 Other Functionality

11.10.1 Universe Size

Many “dynamic” MPI applications are expected to exist in a static runtime environment, in which resources have been allocated before the application is run. When running one of these quasi-static applications, the user (or possibly a batch system) will usually specify a number of processes to start and a total number of processes that are expected. An application simply needs to know how many slots there are, i.e., how many processes it should spawn.

MPI provides an attribute on `MPI_COMM_WORLD`, `MPI_UNIVERSE_SIZE`, that allows the application to obtain this information in a portable manner. This attribute indicates the total number of processes that are expected. In Fortran, the attribute is the integer value. In C, the attribute is a pointer to the integer value. An application typically subtracts the size of `MPI_COMM_WORLD` from `MPI_UNIVERSE_SIZE` to find out how many processes it should spawn. `MPI_UNIVERSE_SIZE` is initialized in `MPI_INIT` and is not changed by MPI. If defined, it has the same value on all processes of `MPI_COMM_WORLD`. `MPI_UNIVERSE_SIZE` is determined by the application startup mechanism in a way not specified by MPI. (The size of `MPI_COMM_WORLD` is another example of such a parameter.)

Possibilities for how `MPI_UNIVERSE_SIZE` might be set include:

- A `-universe_size` argument to a program that starts MPI processes.

- Automatic interaction with a batch scheduler to figure out how many processors have been allocated to an application.
- An environment variable set by the user.
- Extra information passed to MPI_COMM_SPAWN through the info argument.

An implementation must document how MPI_UNIVERSE_SIZE is set. An implementation may not support the ability to set MPI_UNIVERSE_SIZE, in which case the attribute MPI_UNIVERSE_SIZE is not set.

MPI_UNIVERSE_SIZE is a recommendation, not necessarily a hard limit. For instance, some implementations may allow an application to spawn 50 processes per processor, if they are requested. However, it is likely that the user only wants to spawn one process per processor.

MPI_UNIVERSE_SIZE is assumed to have been specified when an application was started, and is in essence a portable mechanism to allow the user to pass to the application (through the MPI process startup mechanism, such as `mpirun`) a piece of critical runtime information. Note that no interaction with the runtime environment is required. If the runtime environment changes size while an application is running, MPI_UNIVERSE_SIZE is not updated, and the application must find out about the change through direct communication with the runtime system.

11.10.2 Singleton MPI Initialization

A high-quality implementation will allow any process (including those not started with a “parallel application” mechanism) to become an MPI process by calling MPI_INIT, MPI_INIT_THREAD, or MPI_SESSION_INIT. Such a process can then connect to other MPI processes using the MPI_COMM_ACCEPT and MPI_COMM_CONNECT routines, or spawn other MPI processes. MPI does not mandate this behavior, but strongly encourages it where technically feasible.

Advice to implementors. Special coordination is required to start MPI processes belonging to the same MPI_COMM_WORLD in the case of the World Model, or the same “`mpi://WORLD`” process set in the Sessions Model. The processes must be started at the “same” time, they must have a mechanism to establish communication, etc. Either the user or the operating system must take special steps beyond simply starting processes.

Considering the World Model, when an application enters MPI_INIT, clearly it must be able to determine if these special steps were taken. If a process enters MPI_INIT and determines that no special steps were taken (i.e., it has not been given the information to form an MPI_COMM_WORLD with other processes) it succeeds and forms a singleton MPI program, that is, one in which MPI_COMM_WORLD has size 1.

In some implementations, MPI may not be able to function without an “MPI environment.” For example, MPI may require that daemons be running or MPI may not be able to work at all on the front-end of an MPP. In this case, an MPI implementation may either

1. Create the environment (e.g., start a daemon) or
2. Raise an error if it cannot create the environment and the environment has not been started independently.

1 A high-quality implementation will try to create a singleton MPI process and not raise
 2 an error. (*End of advice to implementors.*)

4 11.10.3 MPI_APPNUM

5 There is a predefined attribute MPI_APPNUM of MPI_COMM_WORLD. In Fortran, the at-
 6 tribute is an integer value. In C, the attribute is a pointer to an integer value. If a process
 7 was spawned with MPI_COMM Spawn_MULTIPLE, MPI_APPNUM is the command number
 8 that generated the current process. Numbering starts from zero. If a process was spawned
 9 with MPI_COMM_SPAWN, it will have MPI_APPNUM equal to zero.

10 Additionally, if the process was not started by a spawn call, but by an implementation-
 11 specific startup mechanism that can handle multiple process specifications, MPI_APPNUM
 12 should be set to the number of the corresponding process specification. In particular, if it
 13 is started with

```
14 mpiexec spec0 [: spec1 : spec2 : ...]
```

15 MPI_APPNUM should be set to the number of the corresponding specification.

16 If an application was not spawned with MPI_COMM_SPAWN or
 17 MPI_COMM_SPAWN_MULTIPLE, and MPI_APPNUM does not make sense in the context of
 18 the implementation-specific startup mechanism, MPI_APPNUM is not set.

19 MPI implementations may optionally provide a mechanism to override the value of
 20 MPI_APPNUM through the info argument. MPI reserves the following key for all SPAWN
 21 calls.

22 **"appnum"**: Value contains an integer that overrides the default value for MPI_APPNUM in
 23 the child.

24 *Rationale.* When a single application is started, it is able to figure out how many pro-
 25 cesses there are by looking at the size of MPI_COMM_WORLD. An application consisting
 26 of multiple SPMD sub-applications has no way to find out how many sub-applications
 27 there are and to which sub-application the process belongs. While there are ways to
 28 figure it out in special cases, there is no general mechanism. MPI_APPNUM provides
 29 such a general mechanism. (*End of rationale.*)

30 11.10.4 Releasing Connections

31 Before a client and a server connect, they are independent MPI applications. An error in
 32 one does not affect the other. After establishing a connection with MPI_COMM_CONNECT
 33 and MPI_COMM_ACCEPT, an error in one may affect the other. It is desirable for a client
 34 and a server to be able to disconnect, so that an error in one will not affect the other.
 35 Similarly, it might be desirable for a parent and child to disconnect, so that errors in the
 36 child do not affect the parent, or vice-versa.

- 37 • Two processes are **connected** if there is a communication path (direct or indirect)
 38 between them. More precisely:
 - 39 1. Two processes are connected if
 - 40 a) they both belong to the same communicator (inter- or intra-, including
 41 MPI_COMM_WORLD) *or*

- b) they have previously belonged to a communicator that was freed with `MPI_COMM_FREE` instead of `MPI_COMM_DISCONNECT` *or*
- c) they both belong to the group of the same window or file handle.

2. If A is connected to B and B to C, then A is connected to C.

- Two processes are **disconnected** (also **independent**) if they are not connected.
- By the above definitions, connectivity is a transitive property, and divides the universe of MPI processes into disconnected (independent) sets (equivalence classes) of processes.
- Processes that are connected, but do not share the same `MPI_COMM_WORLD`, may become disconnected (independent) if the communication path between them is broken by using `MPI_COMM_DISCONNECT`.

The following additional rules apply to MPI routines in other chapters:

- `MPI_FINALIZE` is collective over a set of connected processes.
- `MPI_ABORT` does not abort independent processes. It may abort all processes in the caller's `MPI_COMM_WORLD` (ignoring its `comm` argument). Additionally, it may abort connected processes as well, though it makes a “best attempt” to abort only the processes in `comm`.
- If a process terminates without calling `MPI_FINALIZE`, independent processes are not affected but the effect on connected processes is not defined.

Advice to implementors. In practice, it may be difficult to distinguish between an MPI process failure and an erroneous program that terminates without calling an MPI finalization function: an implementation that defines semantics for process failure management may have to exhibit the behavior defined for MPI process failures with such erroneous programs. A high quality implementation should exhibit a different behavior for erroneous programs and MPI process failures. (*End of advice to implementors.*)

`MPI_COMM_DISCONNECT(comm)`

INOUT `comm` communicator (handle)

C binding

```
int MPI_Comm_disconnect(MPI_Comm *comm)
```

Fortran 2008 binding

```
MPI_Comm_disconnect(comm, ierror)
  TYPE(MPI_Comm), INTENT(INOUT) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_DISCONNECT(COMM, IERROR)
  INTEGER COMM, IERROR
```

This function waits for all *decoupled MPI activities* on `comm` to complete internally, deallocates the communicator object, and sets the handle to `MPI_COMM_NULL`. It is a collective operation.

It may not be called with the communicator `MPI_COMM_WORLD` or `MPI_COMM_SELF`.

`MPI_COMM_DISCONNECT` may be called only if all communication is complete and matched, so that buffered data can be delivered to its destination. This requirement is the same as for `MPI_FINALIZE`. This means that before calling `MPI_COMM_DISCONNECT`, all request handles associated with `comm` must be freed in the case of nonblocking operations, and must be inactive or freed in the case of persistent operations (i.e., by calling one of the procedures `MPI_{TEST|WAIT}{|ANY|SOME|ALL}` or `MPI_REQUEST_FREE`).

`MPI_COMM_DISCONNECT` has the same effect as `MPI_COMM_FREE`, except that it waits for *decoupled MPI activities* on `comm` to finish internally, disallows any further use of derived inactive persistent requests, and enables the guarantee about the behavior of disconnected processes. The *decoupled MPI activities* also include any communication that is needed to complete a nonblocking or persistent operation on `comm` that was freed with `MPI_REQUEST_FREE`. After calling `MPI_COMM_DISCONNECT`, freeing or starting an inactive persistent request handle for a communication operation on `comm` is erroneous.

Advice to users. To disconnect two processes you may need to call `MPI_COMM_DISCONNECT`, `MPI_WIN_FREE`, and `MPI_FILE_CLOSE` to remove all communication paths between the two processes. Note that it may be necessary to disconnect several communicators (or to free several windows or files) before two processes are completely independent. (*End of advice to users.*)

Rationale. It would be nice to be able to use `MPI_COMM_FREE` instead, but that procedure explicitly does not wait for *decoupled MPI activities* to complete, and it does not disallow freeing or starting of related inactive (but not yet freed) persistent request handles. (*End of rationale.*)

11.10.5 Another Way to Establish MPI Communication

`MPI_COMM_JOIN(fd, intercomm)`

IN	<code>fd</code>	socket file descriptor
OUT	<code>intercomm</code>	new inter-communicator (handle)

C binding

```
int MPI_Comm_join(int fd, MPI_Comm *intercomm)
```

Fortran 2008 binding

```
MPI_Comm_join(fd, intercomm, ierror)
  INTEGER, INTENT(IN) :: fd
  TYPE(MPI_Comm), INTENT(OUT) :: intercomm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMM_JOIN(FD, INTERCOMM, IERROR)
  INTEGER FD, INTERCOMM, IERROR
```

MPI_COMM_JOIN is intended for MPI implementations that exist in an environment supporting the Berkeley Socket interface [52, 57]. Implementations that exist in an environment not supporting Berkeley Sockets should provide the entry point for MPI_COMM_JOIN and should return MPI_COMM_NULL.

This call creates an inter-communicator from the union of two MPI processes that are connected by a socket. MPI_COMM_JOIN should normally succeed if the local and remote processes have access to the same implementation-defined MPI communication universe.

Advice to users. An MPI implementation may require a specific communication medium for MPI communication, such as a shared memory segment or a special switch. In this case, it may not be possible for two processes to successfully join even if there is a socket connecting them and they are using the same MPI implementation. (*End of advice to users.*)

Advice to implementors. A high-quality implementation will attempt to establish communication over a slow medium if its preferred one is not available. If implementations do not do this, they must document why they cannot do MPI communication over the medium used by the socket (especially if the socket is a TCP connection). (*End of advice to implementors.*)

fd is a file descriptor representing a socket of type SOCK_STREAM (a two-way reliable byte-stream connection). Nonblocking I/O and asynchronous notification via SIGIO must not be enabled for the socket. The socket must be in a connected state. The socket must be quiescent when MPI_COMM_JOIN is called (see below). It is the responsibility of the application to create the socket using standard socket API calls.

MPI_COMM_JOIN must be called by the process at each end of the socket. It does not return until both processes have called MPI_COMM_JOIN. The two processes are referred to as the local and remote processes.

MPI only uses the socket to bootstrap the creation of the inter-communicator. Upon return from MPI_COMM_JOIN, the file descriptor will be open and quiescent (see below).

If MPI is unable to create an inter-communicator, but is able to leave the socket in its original state, with no pending communication, it succeeds and sets intercomm to MPI_COMM_NULL.

The socket must be quiescent before MPI_COMM_JOIN is called and after MPI_COMM_JOIN returns. More specifically, on entry to MPI_COMM_JOIN, a read on the socket will not read any data that was written to the socket before the remote process called MPI_COMM_JOIN. On exit from MPI_COMM_JOIN, a read will not read any data that was written to the socket before the remote process returned from MPI_COMM_JOIN. It is the responsibility of the application to ensure the first condition, and the responsibility of the MPI implementation to ensure the second. In a multithreaded application, the application must ensure that one thread does not access the socket while another is calling MPI_COMM_JOIN, or call MPI_COMM_JOIN concurrently.

Advice to implementors. MPI is free to use any available communication path(s) for MPI messages in the new communicator; the socket is only used for the initial handshaking. (*End of advice to implementors.*)

MPI_COMM_JOIN uses non-MPI communication to do its work. The interaction of non-MPI communication with pending MPI communication is not defined. Therefore, the

1 result of calling `MPI_COMM_JOIN` on two connected processes (see Section [11.10.4](#) for the
2 definition of connected) is undefined.

3 The returned communicator may be used to establish MPI communication with addi-
4 tional processes, through the usual MPI communicator creation mechanisms.
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

DRAFT

Chapter 12

One-Sided Communications

12.1 Introduction

Remote Memory Access (RMA) extends the communication mechanisms of MPI by allowing one MPI process to specify all communication parameters, both for the sending side and for the receiving side. This mode of communication facilitates the coding of some applications with dynamically changing data access patterns where the data distribution is fixed or slowly changing. In such a case, each MPI process can compute what data it needs to access or to update at other MPI processes. However, the programmer may not be able to easily determine which data in an MPI process may need to be accessed or to be updated by operations initiated by a different MPI process, and may not even know which MPI processes may perform such updates. Thus, the transfer parameters are all available only on one side. Regular send/receive communication requires matching operations by sender and receiver. In order to issue the matching operations, an application needs to distribute the transfer parameters. This distribution may require all MPI processes to participate in a time-consuming global computation, or to poll for potential communication requests to receive and upon which to act periodically. The use of RMA communication operations avoids the need for global computations or explicit polling. A generic example of this nature is the execution of an assignment of the form $A = B(\text{map})$, where `map` is a permutation vector, and `A`, `B`, and `map` are distributed in the same manner.

Message-passing communication achieves two effects: *communication* of data from sender to receiver and *synchronization* of sender with receiver. The RMA design separates these two functions. The following communication calls are provided:

- Remote write: `MPI_PUT`, `MPI_RPUT`
- Remote read: `MPI_GET`, `MPI_RGET`
- Remote update: `MPI_ACCUMULATE`, `MPI_RACCUMULATE`
- Remote read and update: `MPI_GET_ACCUMULATE`, `MPI_RGET_ACCUMULATE`, and `MPI_FETCH_AND_OP`
- Remote atomic swap: `MPI_COMPARE_AND_SWAP`

This chapter refers to an operations set that includes all remote update, remote read and update, and remote atomic swap operations as “accumulate” operations.

MPI supports two fundamentally different *memory models*: *separate* and *unified*. The separate model makes no assumption about memory consistency and is highly portable. This model is similar to that of weakly coherent memory systems: the user must impose correct ordering of memory accesses through synchronization calls. The unified model can

1 exploit cache-coherent hardware and hardware-accelerated, one-sided operations that are
 2 commonly available in high-performance systems. The two different models are discussed
 3 in detail in Section 12.4. Both models support several synchronization calls to support
 4 different synchronization styles.

5 The design of the RMA functions allows implementors to take advantage of fast or
 6 asynchronous communication mechanisms provided by various platforms, such as coherent
 7 or noncoherent shared memory, DMA engines, hardware-supported put/get operations, and
 8 communication coprocessors. The most frequently used RMA communication mechanisms
 9 can be layered on top of message-passing. However, certain RMA functions might need
 10 support for asynchronous communication agents in software (handlers, threads, etc.) in a
 11 distributed memory environment.

12 We shall denote by **origin** or *origin process* the MPI process that calls an RMA proce-
 13 dure, and by **target** or *target process* the MPI process whose memory is accessed. Thus, in
 14 a put operation, source = origin and destination = target; in a get operation, source = target
 15 and destination = origin.

17 12.2 Initialization

19 MPI provides the following window initialization functions: MPI_WIN_CREATE,
 20 MPI_WIN_ALLOCATE, MPI_WIN_ALLOCATE_SHARED, and
 21 MPI_WIN_CREATE_DYNAMIC, which are collective over the group of an intra-communica-
 22 tor. MPI_WIN_CREATE allows each MPI process to specify a “window” in its memory that
 23 is made available for accesses by other MPI processes. The call returns an opaque object
 24 that represents the group of MPI processes that own and access the set of windows, and the
 25 attributes of each window, as specified by the initialization call. MPI_WIN_ALLOCATE and
 26 MPI_WIN_ALLOCATE_SHARED differ from MPI_WIN_CREATE in that the user does not
 27 pass allocated memory; instead MPI_WIN_ALLOCATE and MPI_WIN_ALLOCATE_SHARED
 28 return a pointer to memory allocated by the MPI implementation.

29 MPI_WIN_ALLOCATE_SHARED differs from MPI_WIN_ALLOCATE in that the allocated
 30 memory is guaranteed to be accessible from all MPI processes in the window’s group with
 31 direct load/store accesses. Some restrictions may apply to the specified communicator.
 32 MPI_WIN_CREATE_DYNAMIC creates a window that allows the user to dynamically control
 33 which memory is exposed by the window.
 34

35 12.2.1 Window Creation

36 MPI_WIN_CREATE(base, size, disp_unit, info, comm, win)

37	IN	base	initial address of window (choice)
38	IN	size	size of window in bytes (non-negative integer)
39	IN	disp_unit	local unit size for displacements, in bytes (positive integer)
40	IN	info	info argument (handle)
41	IN	comm	intra-communicator (handle)
42	OUT	win	window object (handle)

C binding

```
int MPI_Win_create(void *base, MPI_Aint size, int disp_unit, MPI_Info info,
                  MPI_Comm comm, MPI_Win *win)
```

```
int MPI_Win_create_c(void *base, MPI_Aint size, MPI_Aint disp_unit,
                    MPI_Info info, MPI_Comm comm, MPI_Win *win)
```

Fortran 2008 binding

```
MPI_Win_create(base, size, disp_unit, info, comm, win, ierror)
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
INTEGER, INTENT(IN) :: disp_unit
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Win), INTENT(OUT) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Win_create(base, size, disp_unit, info, comm, win, ierror) !(_c)
```

```
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size, disp_unit
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Win), INTENT(OUT) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_CREATE(BASE, SIZE, DISP_UNIT, INFO, COMM, WIN, IERROR)
```

```
<type> BASE(*)
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE
INTEGER DISP_UNIT, INFO, COMM, WIN, IERROR
```

This procedure is collective over the group of `comm`. It returns a handle to a window that can be used by the MPI processes in this group to perform RMA operations. Each MPI process specifies a window of existing memory that it exposes to RMA accesses by any MPI processes in the group of `comm`. The window consists of `size` bytes, starting at address `base`. In C, `base` is the starting address of a memory region. In Fortran, one can pass the first element of a memory region or a whole array, which must be ‘simply contiguous’ (for ‘simply contiguous,’ see also Section 19.1.12). An MPI process may elect to expose no memory by specifying `size = 0`.

The displacement unit argument is provided to facilitate address arithmetic in RMA operations: the target displacement argument of an RMA operation is scaled by the factor `disp_unit` specified by the target process, at window creation.

Rationale. The window size is specified using an address-sized integer, rather than a basic integer type, to allow windows that span more memory than can be described with a basic integer type. (*End of rationale.*)

Advice to users. Common choices for `disp_unit` are 1 (no scaling), and (in C syntax) `sizeof(type)`, for a window that consists of an array of elements of type `type`. The latter choice will allow one to use array indices in RMA calls, and have those scaled correctly to byte displacements, even in a heterogeneous environment. (*End of advice to users.*)

1 The `info` argument provides optimization hints to the runtime about the expected usage
2 pattern of the window. The following info keys are predefined:

3
4 **"no_locks"** (boolean, default: **"false"**): if set to true, then the implementation may as-
5 sume that passive target synchronization (i.e., `MPI_WIN_LOCK`,
6 `MPI_WIN_LOCK_ALL`) will not be used on the given window. This implies that this
7 window is not used for 3-party communication, and RMA can be implemented with
8 no (less) asynchronous agent activity at this process.

9
10 **"accumulate_ordering"** (string, default **"rar,raw,war,waw"**): controls the ordering of accu-
11 mulate operations at the target. See Section 12.7.2 for details.

12
13 **"accumulate_ops"** (string, default: **"same_op_no_op"**): if set to **"same_op"**, the implemen-
14 tation will assume that all concurrent accumulate calls to the same target address
15 will use the same operator. If set to **"same_op_no_op"**, then the implementation will
16 assume that all concurrent accumulate calls to the same target address will use the
17 same operator or `MPI_NO_OP`. This can eliminate the need to protect access for certain
18 operators where the hardware can guarantee atomicity.

19
20 **"mpi_accumulate_granularity"** (integer, default **0**): provides a hint to implementations
21 about the desired synchronization granularity for accumulate operations, i.e., the size
22 of memory ranges in bytes for which the implementation should acquire a synchroni-
23 zation primitive to ensure atomicity of updates. If the specified granularity is not
24 divisible by the size of the type used in an accumulate operation, it should be treated
25 as if it was the next multiple of the element size. For example, a granularity of 1 byte
26 should be treated as 8 in an accumulate operation using `MPI_UINT64_T`. By default,
27 this info key is set to 0, which leaves the choice of synchronization granularity to the
28 implementation. If specified, all MPI processes in the group of a window must supply
29 the same value.

30 *Advice to users.* Small synchronization granularities may provide improved
31 latencies for accumulate operations with few elements and potentially increase
32 concurrency of updates, at the cost of lower throughput. For example, a value
33 matching the size of a type involved in an accumulate operation may enable
34 implementations to use atomic memory operations instead of mutual exclusion
35 devices. Larger synchronization granularities may yield higher throughput of ac-
36 cumulate operation with large numbers of elements due to lower synchronization
37 costs, potentially at the expense of higher latency for accumulate operations with
38 few elements, e.g., if atomic memory operations are not employed. By dividing
39 larger accumulate operations into smaller segments, concurrent accumulate op-
40 erations to the same window memory may update different segments in parallel.
41 (*End of advice to users.*)

42 *Advice to implementors.* Implementations are encouraged to avoid mutual
43 exclusion devices in cases where the granularity is small enough to warrant the use
44 of atomic memory operations. For larger granularities, implementations should
45 use this info value as a hint to partition the window memory into zones of mutual
46 exclusion to enable segmentation of large accumulate operations. (*End of advice*
47 *to implementors.*)

"same_size" (boolean, default: **"false"**): if set to true, then the implementation may assume that the argument `size` is identical on all MPI processes, and that all MPI processes have provided this info key with the same value.

"same_disp_unit" (boolean, default: **"false"**): if set to true, then the implementation may assume that the argument `disp_unit` is identical on all MPI processes, and that all MPI processes have provided this info key with the same value.

"mpi_assert_memory_alloc_kinds" (string, not set by default): If set, the implementation may assume that the memory for all communication buffers passed to MPI operations performed by the calling MPI process on the given window will use only the memory allocation kinds listed in the value string. See Section 11.4.3. This info hint also applies to the window buffer provided in a call to `MPI_WIN_CREATE` or `MPI_WIN_ATTACH`. It does not apply to the memory allocated in a call to `MPI_WIN_ALLOCATE` or `MPI_WIN_ALLOCATE_SHARED`.

Advice to users. The info query mechanism described in Section 12.2.7 can be used to query the specified info arguments for windows that have been passed to a library. It is recommended that libraries check attached info keys for each passed window. (*End of advice to users.*)

The various MPI processes in the group of `comm` may specify completely different target windows, in location, size, displacement units, and info arguments. As long as all the get, put and accumulate accesses to a particular MPI process fit their specific target window this should pose no problem. The same area in memory may appear in multiple windows, each associated with a different window object. However, concurrent communications to distinct, overlapping windows may lead to undefined results.

Implementations may make the memory provided by the user available for load/store accesses by MPI processes in the same *shared memory domain*. A communicator of such processes can be constructed as described in Section 7.4.2 using `MPI_COMM_SPLIT_TYPE`. Pointers to access a *shared memory segment* can be queried using `MPI_WIN_SHARED_QUERY`.

Rationale. The reason for specifying the memory that may be accessed from another MPI process in an RMA operation is to permit the programmer to specify what memory can be a target of RMA operations and for the implementation to enforce that specification. For example, with this definition, a server MPI process can safely allow a client MPI process to use RMA operations, knowing that (under the assumption that the MPI implementation does enforce the specified limits on the exposed memory) an error in the client cannot affect any memory other than what was explicitly exposed. (*End of rationale.*)

Advice to users. A window can be created in any part of the MPI process memory. However, on some systems, the performance of windows in memory allocated by `MPI_ALLOC_MEM` (Section 9.2) will be better. Also, on some systems, performance is improved when window boundaries are aligned at "natural" boundaries (word, double-word, cache line, page frame, etc.). (*End of advice to users.*)

Advice to implementors. In cases where RMA operations use different mechanisms in different memory areas (e.g., load/store accesses in a *shared memory segment*, and

an asynchronous handler in private memory), the `MPI_WIN_CREATE` call needs to figure out which type of memory is used for the window. To do so, MPI maintains, internally, the list of memory segments allocated by `MPI_ALLOC_MEM`, or by other, implementation-specific, mechanisms, together with information on the type of memory segment allocated. When a call to `MPI_WIN_CREATE` occurs, then MPI checks which segment contains each window, and decides, accordingly, which mechanism to use for RMA operations.

Vendors may provide additional, implementation-specific mechanisms to allocate or to specify memory regions that are preferable for use in one-sided communication. In particular, such mechanisms can be used to place static variables into such preferred regions.

Implementors should document any performance impact of window alignment. (*End of advice to implementors.*)

12.2.2 Window That Allocates Memory

`MPI_WIN_ALLOCATE(size, disp_unit, info, comm, baseptr, win)`

IN	size	size of window in bytes (non-negative integer)
IN	disp_unit	local unit size for displacements, in bytes (positive integer)
IN	info	info argument (handle)
IN	comm	intra-communicator (handle)
OUT	baseptr	initial address of window (choice)
OUT	win	window object (handle)

C binding

```
int MPI_Win_allocate(MPI_Aint size, int disp_unit, MPI_Info info,
                    MPI_Comm comm, void *baseptr, MPI_Win *win)
```

```
int MPI_Win_allocate_c(MPI_Aint size, MPI_Aint disp_unit, MPI_Info info,
                      MPI_Comm comm, void *baseptr, MPI_Win *win)
```

Fortran 2008 binding

```
MPI_Win_allocate(size, disp_unit, info, comm, baseptr, win, ierror)
  USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
  INTEGER, INTENT(IN) :: disp_unit
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(C_PTR), INTENT(OUT) :: baseptr
  TYPE(MPI_Win), INTENT(OUT) :: win
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Win_allocate(size, disp_unit, info, comm, baseptr, win, ierror) !(_c)
  USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
```

```

INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size, disp_unit      1
TYPE(MPI_Info), INTENT(IN) :: info                                2
TYPE(MPI_Comm), INTENT(IN) :: comm                                3
TYPE(C_PTR), INTENT(OUT) :: baseptr                               4
TYPE(MPI_Win), INTENT(OUT) :: win                                5
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                          6

```

Fortran binding

```

MPI_WIN_ALLOCATE(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, WIN, IERROR)  8
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR                       9
INTEGER DISP_UNIT, INFO, COMM, WIN, IERROR                         10

```

This procedure is collective over the group of `comm`. On each MPI process, it allocates memory of at least `size` bytes and returns a pointer to it along with a handle to a new window that can be used by all MPI processes in the group of `comm` to perform RMA operations. The returned memory consists of `size` bytes local to each MPI process, starting at address `baseptr` and is associated with the window as if the user called `MPI_WIN_CREATE` on existing memory. The size argument may be different at each MPI process and `size = 0` is valid; however, a library might allocate and expose more memory in order to create a fast, globally symmetric allocation. The discussion of and rationales for `MPI_ALLOC_MEM` and `MPI_FREE_MEM` in Section 9.2 also apply to `MPI_WIN_ALLOCATE`; in particular, see the rationale in Section 9.2 for an explanation of the type used for `baseptr`.

Implementations may make allocated memory available for load/store accesses by MPI processes in the same *shared memory domain*. A communicator of such processes can be constructed as described in Section 7.4.2 using `MPI_COMM_SPLIT_TYPE`. Pointers to access a *shared memory segment* can be queried using `MPI_WIN_SHARED_QUERY`. If *shared memory* is available it is not guaranteed to be *contiguous* (see Section 12.2.3).

If the Fortran compiler provides `TYPE(C_PTR)`, then the following generic interface must be provided in the `mpi` module and should be provided in the (deprecated) `mpif.h` include file through overloading, i.e., with the same routine name as the routine with `INTEGER(KIND=MPI_ADDRESS_KIND) BASEPTR`, but with a different specific procedure name:

```

INTERFACE MPI_WIN_ALLOCATE                                          32
  SUBROUTINE MPI_WIN_ALLOCATE(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, &  33
    WIN, IERROR)                                                    34
    IMPORT :: MPI_ADDRESS_KIND                                        35
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR                  36
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR                 37
  END SUBROUTINE                                                    38
  SUBROUTINE MPI_WIN_ALLOCATE_CPTR(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, &  39
    WIN, IERROR)                                                    40
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR                  41
    IMPORT :: MPI_ADDRESS_KIND                                        42
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR                  43
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE                          44
    TYPE(C_PTR) :: BASEPTR                                          45
  END SUBROUTINE                                                    46
END INTERFACE                                                        47

```

1 The base procedure name of this overloaded function is `MPI_WIN_ALLOCATE_CPTR`.
 2 The implied specific procedure names are described in Section 19.1.5.

3
 4 *Rationale.* By allocating (potentially aligned) memory instead of allowing the user
 5 to pass in an arbitrary buffer, this call can improve the performance for systems with
 6 remote direct memory access. This also permits the collective allocation of memory
 7 and supports what is sometimes called the “symmetric allocation” model that can be
 8 more scalable (for example, the implementation can arrange to return an address for
 9 the allocated memory that is the same on all MPI processes). (*End of rationale.*)

10
 11 The `info` argument can be used to specify hints similar to the `info` argument for
 12 `MPI_WIN_CREATE` and `MPI_ALLOC_MEM`.

13 The default memory alignment requirements and the “`mpi_minimum_memory_alignment`”
 14 info key described for `MPI_ALLOC_MEM` in Section 9.2 apply to all MPI processes with
 15 nonzero `size` argument.

16 12.2.3 Window That Allocates Shared Memory

17
 18
 19
 20 `MPI_WIN_ALLOCATE_SHARED(size, disp_unit, info, comm, baseptr, win)`

21	IN	size	size of local window in bytes (non-negative integer)
22	IN	disp_unit	local unit size for displacements, in bytes (positive integer)
23			
24			
25	IN	info	info argument (handle)
26	IN	comm	intra-communicator (handle)
27	OUT	baseptr	address of local allocated window segment (choice)
28			
29	OUT	win	window object (handle)

30 **C binding**

31
 32 `int MPI_Win_allocate_shared(MPI_Aint size, int disp_unit, MPI_Info info,`
 33 `MPI_Comm comm, void *baseptr, MPI_Win *win)`
 34
 35 `int MPI_Win_allocate_shared_c(MPI_Aint size, MPI_Aint disp_unit, MPI_Info info,`
 36 `MPI_Comm comm, void *baseptr, MPI_Win *win)`

37 **Fortran 2008 binding**

38 `MPI_Win_allocate_shared(size, disp_unit, info, comm, baseptr, win, ierror)`
 39 `USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR`
 40 `INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size`
 41 `INTEGER, INTENT(IN) :: disp_unit`
 42 `TYPE(MPI_Info), INTENT(IN) :: info`
 43 `TYPE(MPI_Comm), INTENT(IN) :: comm`
 44 `TYPE(C_PTR), INTENT(OUT) :: baseptr`
 45 `TYPE(MPI_Win), INTENT(OUT) :: win`
 46 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`
 47 `MPI_Win_allocate_shared(size, disp_unit, info, comm, baseptr, win, ierror)`
 48 `!(_c)`

```

USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR           1
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size, disp_unit  2
TYPE(MPI_Info), INTENT(IN) :: info                       3
TYPE(MPI_Comm), INTENT(IN) :: comm                      4
TYPE(C_PTR), INTENT(OUT) :: baseptr                    5
TYPE(MPI_Win), INTENT(OUT) :: win                       6
INTEGER, OPTIONAL, INTENT(OUT) :: ierror               7

```

Fortran binding

```

MPI_WIN_ALLOCATE_SHARED(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, WIN, IERROR)  9
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR                             10
INTEGER DISP_UNIT, INFO, COMM, WIN, IERROR                               11

```

This procedure is collective over the group of `comm`. On each MPI process, it allocates memory of at least `size` bytes that is shared among all MPI processes in `comm`, and returns a pointer to the locally allocated segment in `baseptr` that can be used for load/store accesses on the calling MPI process. The locally allocated memory can be the target of load/store accesses by remote MPI processes; the base pointers for other MPI processes can be queried using the function `MPI_WIN_SHARED_QUERY`. The call also returns a handle to a new window that can be used by all MPI processes in `comm` to perform RMA operations. The size argument may be different at each MPI process and `size = 0` is valid. It is the user's responsibility to ensure that the communicator `comm` represents a group of MPI processes that are in the same *shared memory domain*, i.e., that they can create a *shared memory segment* that can be accessed by all processes in the group. The discussions of rationales for `MPI_ALLOC_MEM` and `MPI_FREE_MEM` in Section 9.2 also apply to `MPI_WIN_ALLOCATE_SHARED`; in particular, see the rationale in Section 9.2 for an explanation of the type used for `baseptr`. The allocated memory is *contiguous across processes in rank order* unless the info key "alloc_shared_noncontig" is specified. Contiguous across processes in rank order means that the first address in the memory segment of MPI process i is consecutive with the last address in the memory segment of MPI process $i - 1$. This may enable the user to calculate remote address offsets with local information only.

If the Fortran compiler provides `TYPE(C_PTR)`, then the following generic interface must be provided in the `mpi` module and should be provided in the (deprecated) `mpif.h` include file through overloading, i.e., with the same routine name as the routine with `INTEGER(KIND=MPI_ADDRESS_KIND) BASEPTR`, but with a different specific procedure name:

```

INTERFACE MPI_WIN_ALLOCATE_SHARED                                     36
  SUBROUTINE MPI_WIN_ALLOCATE_SHARED(SIZE, DISP_UNIT, INFO, COMM, &  37
    BASEPTR, WIN, IERROR)                                           38
    IMPORT :: MPI_ADDRESS_KIND                                       39
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR                   40
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR                 41
  END SUBROUTINE                                                    42
  SUBROUTINE MPI_WIN_ALLOCATE_SHARED_CPTR(SIZE, DISP_UNIT, INFO, COMM, &  43
    BASEPTR, WIN, IERROR)                                           44
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR                   45
    IMPORT :: MPI_ADDRESS_KIND                                       46
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR                   47
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE                           48

```

```
1      TYPE(C_PTR) :: BASEPTR
2      END SUBROUTINE
3  END INTERFACE
4
```

5 The base procedure name of this overloaded function is
6 `MPI_WIN_ALLOCATE_SHARED_CPTR`. The implied specific procedure names are described
7 in Section 19.1.5.

8 The `info` argument can be used to specify hints similar to the `info` argument for
9 `MPI_WIN_CREATE`, `MPI_WIN_ALLOCATE`, and `MPI_ALLOC_MEM`. The additional `info`
10 key `"alloc_shared_noncontig"` allows the library to optimize the layout of the *shared memory*
11 *segments* in memory.

12
13 *Advice to users.* If the `info` key `"alloc_shared_noncontig"` is not set to true, the allocation
14 strategy is to allocate *contiguous memory* across MPI process ranks. This may limit
15 the performance on some architectures because it does not allow the implementation
16 to modify the data layout (e.g., padding to reduce access latency). (*End of advice to*
17 *users.*)

18
19 *Advice to implementors.* If the user sets the `info` key `"alloc_shared_noncontig"` to true,
20 the implementation can allocate the memory requested by each MPI process in a loca-
21 tion that is close to this MPI process. This can be achieved by padding or allocating
22 memory in special memory segments. Both techniques may make the address space
23 across consecutive ranks *noncontiguous*. (*End of advice to implementors.*)

24
25 For *contiguous shared memory* allocations, the default alignment requirements outlined
26 for `MPI_ALLOC_MEM` in Section 9.2 and the `"mpi_minimum_memory_alignment"` `info` key apply
27 to the start of the *contiguous memory* that is returned in `baseptr` to the first MPI process
28 with nonzero `size` argument. For noncontiguous memory allocations, the default alignment
29 requirements and the `"mpi_minimum_memory_alignment"` `info` key apply to all MPI processes
30 with nonzero `size` argument.

31
32 *Advice to users.* If the `info` key `"alloc_shared_noncontig"` is not set to true (or ignored
33 by the MPI implementation), the alignment of the memory returned in `baseptr` to all
34 but the first MPI process with nonzero `size` argument depends on the value of the
35 `size` argument provided by other MPI processes. It is thus the user's responsibility
36 to control the alignment of contiguous memory allocated for these MPI processes by
37 ensuring that each MPI process provides a `size` argument that is an integral multiple
38 of the alignment required for the application. (*End of advice to users.*)

39
40 The consistency of load/store accesses from/to the shared memory as observed by the
41 user program depends on the architecture. For details on how to create a consistent view
42 see the description of `MPI_WIN_SHARED_QUERY`.

MPI_WIN_SHARED_QUERY(win, rank, size, disp_unit, baseptr)			1
IN	win	shared memory window (handle)	2
			3
IN	rank	rank in the group of window win or MPI_PROC_NULL (non-negative integer)	4
			5
OUT	size	size of the window segment (non-negative integer)	6
OUT	disp_unit	local unit size for displacements, in bytes (positive integer)	7
			8
OUT	baseptr	address for load/store access to window segment (choice)	9
			10
			11
			12

C binding

```
int MPI_Win_shared_query(MPI_Win win, int rank, MPI_Aint *size, int *disp_unit,
                        void *baseptr)
```

```
int MPI_Win_shared_query_c(MPI_Win win, int rank, MPI_Aint *size,
                          MPI_Aint *disp_unit, void *baseptr)
```

Fortran 2008 binding

```
MPI_Win_shared_query(win, rank, size, disp_unit, baseptr, ierror)
```

```
USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
```

```
TYPE(MPI_Win), INTENT(IN) :: win
```

```
INTEGER, INTENT(IN) :: rank
```

```
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: size
```

```
INTEGER, INTENT(OUT) :: disp_unit
```

```
TYPE(C_PTR), INTENT(OUT) :: baseptr
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Win_shared_query_c(win, rank, size, disp_unit, baseptr, ierror) !(_c)
```

```
USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
```

```
TYPE(MPI_Win), INTENT(IN) :: win
```

```
INTEGER, INTENT(IN) :: rank
```

```
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: size, disp_unit
```

```
TYPE(C_PTR), INTENT(OUT) :: baseptr
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_SHARED_QUERY(WIN, RANK, SIZE, DISP_UNIT, BASEPTR, IERROR)
```

```
INTEGER WIN, RANK, DISP_UNIT, IERROR
```

```
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR
```

This function queries the MPI process-local address for remote memory segments created with `MPI_WIN_ALLOCATE_SHARED`, `MPI_WIN_ALLOCATE`, and `MPI_WIN_CREATE`. This function can return different MPI process-local addresses for the same physical memory when called by different MPI processes. The returned memory can be used for load/store accesses subject to the constraints defined in Section 12.7. When rank is `MPI_PROC_NULL`, the `baseptr`, `disp_unit`, and `size` returned are the base, displacement unit, and size of the memory segment belonging to the MPI process with the lowest rank in the *shared memory domain* that specified size $\neq 0$. If all MPI processes in the group

1 attached to the window specified `size = 0`, then the call returns `size = 0` and a `baseptr` as
 2 if `MPI_ALLOC_MEM` was called with `size = 0`.

3 Only `MPI_WIN_ALLOCATE_SHARED` is guaranteed to allocate *shared memory*. Im-
 4 plementations are permitted, where possible, to provide *shared memory* for windows cre-
 5 ated with `MPI_WIN_CREATE` and `MPI_WIN_ALLOCATE`. However, availability of *shared*
 6 *memory* is not guaranteed. When the remote memory segment corresponding to a par-
 7 ticular process cannot be accessed directly, this call returns `size = 0` and a `baseptr` as if
 8 `MPI_ALLOC_MEM` was called with `size = 0`.

9
 10 *Rationale.* `MPI_WIN_SHARED_QUERY` may only be called on windows created by
 11 a call to `MPI_WIN_ALLOCATE_SHARED`, `MPI_WIN_ALLOCATE`, or
 12 `MPI_WIN_CREATE`. The potential for multiple memory regions in windows created
 13 through `MPI_WIN_CREATE_DYNAMIC` means that these windows cannot be used as
 14 input for `MPI_WIN_SHARED_QUERY`. (*End of rationale.*)

15
 16 *Advice to users.* For windows allocated using `MPI_WIN_ALLOCATE` or
 17 `MPI_WIN_CREATE`, the group of MPI processes for which the implementation may
 18 provide shared memory can be determined using `MPI_COMM_SPLIT_TYPE` described
 19 in Section 7.4.2. (*End of advice to users.*)

20
 21 The consistency of load/store accesses from/to the *shared memory* as observed by the
 22 user program depends on the architecture. A consistent view can be created in the *unified*
 23 *memory model* (see Section 12.4) by utilizing the window synchronization functions (see
 24 Section 12.5) or explicitly completing outstanding store accesses (e.g., by calling
 25 `MPI_WIN_FLUSH`). MPI does not define the semantics for accessing *shared window memory*
 26 in the *separate memory model*.

27 If the Fortran compiler provides `TYPE(C_PTR)`, then the following generic interface must
 28 be provided in the `mpi` module and should be provided in the (deprecated) `mpif.h` include
 29 file through overloading, i.e., with the same routine name as the routine with
 30 `INTEGER(KIND=MPI_ADDRESS_KIND) BASEPTR`, but with a different specific procedure name:

```
31
32 INTERFACE MPI_WIN_SHARED_QUERY
33     SUBROUTINE MPI_WIN_SHARED_QUERY(WIN, RANK, SIZE, DISP_UNIT, &
34         BASEPTR, IERROR)
35         IMPORT :: MPI_ADDRESS_KIND
36         INTEGER :: WIN, RANK, DISP_UNIT, IERROR
37         INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR
38     END SUBROUTINE
39     SUBROUTINE MPI_WIN_SHARED_QUERY_CPTR(WIN, RANK, SIZE, DISP_UNIT, &
40         BASEPTR, IERROR)
41         USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
42         IMPORT :: MPI_ADDRESS_KIND
43         INTEGER :: WIN, RANK, DISP_UNIT, IERROR
44         INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE
45         TYPE(C_PTR) :: BASEPTR
46     END SUBROUTINE
47 END INTERFACE
48
```


The base procedure name of this overloaded function is `MPI_WIN_SHARED_QUERY_CPTR`. The implied specific procedure names are described in Section 19.1.5.

12.2.4 Window of Dynamically Attached Memory

The previously described window creation procedures require the user to identify the local memory that may be a target of RMA calls at the time the window is created. This has advantages for both the programmer (only this memory can be updated by one-sided operations and provides greater safety) and the MPI implementation (special steps may be taken to make one-sided access to such memory more efficient). However, consider implementing a modifiable linked list using RMA operations; as new items are added to the list, memory must be allocated. In a C or C++ program, this memory is typically allocated using `malloc` or `new` respectively. With the previously described window creation procedures, the programmer must create a window with a predefined amount of memory and then implement routines for allocating memory from within the window's memory. In addition, there is no easy way to handle the situation where the predefined amount of memory turns out to be inadequate. To support this model, the routine `MPI_WIN_CREATE_DYNAMIC` creates a window that makes it possible to expose memory without remote synchronization. It must be used in combination with the local routines `MPI_WIN_ATTACH` and `MPI_WIN_DETACH`.

`MPI_WIN_CREATE_DYNAMIC`(*info*, *comm*, *win*)

IN	<i>info</i>	<i>info</i> argument (handle)
IN	<i>comm</i>	intra-communicator (handle)
OUT	<i>win</i>	window object (handle)

C binding

```
int MPI_Win_create_dynamic(MPI_Info info, MPI_Comm comm, MPI_Win *win)
```

Fortran 2008 binding

```
MPI_Win_create_dynamic(info, comm, win, ierror)
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Win), INTENT(OUT) :: win
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_CREATE_DYNAMIC(INFO, COMM, WIN, IERROR)
  INTEGER INFO, COMM, WIN, IERROR
```

This procedure is collective over the group of `comm`. It returns a window `win` without memory attached. Existing MPI process memory can be attached as described below. This procedure returns a handle to a new window that can be used by MPI processes in the group of `comm` to perform RMA operations on attached memory. Because this window has special properties, it will sometimes be referred to as a **dynamic** window.

The `info` argument can be used to specify hints similar to the `info` argument for `MPI_WIN_CREATE`.

In the case of a window created with `MPI_WIN_CREATE_DYNAMIC`, the `target_disp` for all RMA functions is the address at the target; i.e., the effective `window_base` is `MPI_BOTTOM` and the `disp_unit` is one. For dynamic windows, the `target_disp` argument to RMA communication operations is not restricted to nonnegative values. Users should use `MPI_GET_ADDRESS` at the target process to determine the address of a target memory location and communicate this address to the origin process.

Advice to users. Users are cautioned that displacement arithmetic can overflow in variables of type `MPI_Aint` and result in unexpected values on some platforms. The `MPI_AINT_ADD` and `MPI_AINT_DIFF` functions can be used to safely perform address arithmetic with `MPI_Aint` displacements. (*End of advice to users.*)

Advice to implementors. In environments with heterogeneous data representations, care must be exercised in communicating addresses between MPI processes. For example, it is possible that an address valid at the target MPI process (for example, a 64-bit pointer) cannot be expressed as an address at the origin (for example, the origin uses 32-bit pointers). For this reason, a portable MPI implementation should ensure that the type `MPI_AINT` (see Table 3.3) is able to store addresses from any MPI process. (*End of advice to implementors.*)

Memory at the target cannot be accessed with this window until that memory has been attached using the function `MPI_WIN_ATTACH`. That is, in addition to using `MPI_WIN_CREATE_DYNAMIC` to create an MPI window, the user must use `MPI_WIN_ATTACH` before any local memory may be the target of an MPI RMA operation. Only memory that is currently accessible may be attached.

`MPI_WIN_ATTACH(win, base, size)`

IN	win	window object (handle)
IN	base	initial address of memory to be attached (choice)
IN	size	size of memory to be attached in bytes (non-negative integer)

C binding

```
int MPI_Win_attach(MPI_Win win, void *base, MPI_Aint size)
```

Fortran 2008 binding

```
MPI_Win_attach(win, base, size, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
  INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_ATTACH(WIN, BASE, SIZE, IERROR)
  INTEGER WIN, IERROR
  <type> BASE(*)
  INTEGER(KIND=MPI_ADDRESS_KIND) SIZE
```

Attaches a local memory region beginning at `base` for remote access within the given window. The memory region specified must not contain any part that is already attached to the window `win`, that is, attaching overlapping memory concurrently within the same window is erroneous. The argument `win` must be a window that was created with `MPI_WIN_CREATE_DYNAMIC`. The local memory region attached to the window consists of `size` bytes, starting at address `base`. In C, `base` is the starting address of a memory region. In Fortran, one can pass the first element of a memory region or a whole array, which must be ‘simply contiguous’ (for ‘simply contiguous,’ see Section 19.1.12). Multiple (but nonoverlapping) memory regions may be attached to the same window.

Rationale. Requiring that memory be explicitly attached before it is exposed to one-sided access by other MPI processes can simplify implementations and improve performance. The ability to make memory available for RMA operations without requiring a collective `MPI_WIN_CREATE` call is needed for some one-sided programming models. (*End of rationale.*)

Advice to users. Attaching memory to a window may require the use of scarce resources; thus, attaching large regions of memory is not recommended in portable programs. Attaching memory to a window may fail if sufficient resources are not available; this is similar to the behavior of `MPI_ALLOC_MEM`.

The user is also responsible for ensuring that `MPI_WIN_ATTACH` at the target has returned before an MPI process attempts to target that memory with an MPI RMA operation.

Performing an RMA operation on memory that has not been attached to a window created with `MPI_WIN_CREATE_DYNAMIC` is erroneous. (*End of advice to users.*)

Advice to implementors. A high-quality implementation will attempt to make as much memory available for attaching as possible. Any limitations should be documented by the implementor. (*End of advice to implementors.*)

`MPI_WIN_ATTACH` is a local procedure that is not collective. Memory may be detached with the procedure `MPI_WIN_DETACH`. After memory has been detached, it may not be the target of an MPI RMA operation on that window (unless the memory is re-attached with `MPI_WIN_ATTACH`).

`MPI_WIN_DETACH(win, base)`

IN	<code>win</code>	window object (handle)
IN	<code>base</code>	initial address of memory to be detached (choice)

C binding

```
int MPI_Win_detach(MPI_Win win, const void *base)
```

Fortran 2008 binding

```
MPI_Win_detach(win, base, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```

1 MPI_WIN_DETACH(WIN, BASE, IERROR)
2   INTEGER WIN, IERROR
3   <type> BASE(*)

```

Detaches a previously attached memory region beginning at `base`. The arguments `base` and `win` must match the arguments passed to a previous call to `MPI_WIN_ATTACH`. `MPI_WIN_DETACH` is a local procedure that is not collective.

Advice to users. Detaching memory may permit the implementation to make more efficient use of special memory or provide memory that may be needed by a subsequent `MPI_WIN_ATTACH`. Users are encouraged to detach memory that is no longer needed. Memory should be detached before it is freed by the user. (*End of advice to users.*)

Memory becomes detached when the associated dynamic memory window is freed, see Section 12.2.5.

12.2.5 Window Destruction

```

20 MPI_WIN_FREE(win)

```

```

21   INOUT   win           window object (handle)

```

C binding

```

25 int MPI_Win_free(MPI_Win *win)

```

Fortran 2008 binding

```

27 MPI_Win_free(win, ierror)
28   TYPE(MPI_Win), INTENT(INOUT) :: win
29   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

32 MPI_WIN_FREE(WIN, IERROR)
33   INTEGER WIN, IERROR

```

Frees the window object `win` and returns a null handle (equal to `MPI_WIN_NULL`). This procedure is collective over the group associated with `win`. `MPI_WIN_FREE` can be invoked by an MPI process only after it has completed its involvement in RMA communications on window `win`: e.g., the MPI process has called `MPI_WIN_FENCE`, or called `MPI_WIN_WAIT` to match a previous call to `MPI_WIN_POST`, called `MPI_WIN_COMPLETE` to match a previous call to `MPI_WIN_START`, or called `MPI_WIN_UNLOCK` to match a previous call to `MPI_WIN_LOCK`. The memory associated with windows created by a call to `MPI_WIN_CREATE` may be freed after the call returns. If the window was created with `MPI_WIN_ALLOCATE`, `MPI_WIN_FREE` will free the window memory that was allocated in `MPI_WIN_ALLOCATE`. If the window was created with `MPI_WIN_ALLOCATE_SHARED`, `MPI_WIN_FREE` will free the window memory that was allocated in `MPI_WIN_ALLOCATE_SHARED`.

Table 12.1: C types of attribute value argument to MPI_WIN_GET_ATTR and MPI_WIN_SET_ATTR

Attribute	C Type
MPI_WIN_BASE	void *
MPI_WIN_SIZE	MPI_Aint *
MPI_WIN_DISP_UNIT	int *
MPI_WIN_CREATE_FLAVOR	int *
MPI_WIN_MODEL	int *

Freeing a window that was created with a call to MPI_WIN_CREATE_DYNAMIC detaches all associated memory; i.e., it has the same effect as if all attached memory was detached by calls to MPI_WIN_DETACH.

MPI_WIN_FREE is required to delay its return until all accesses to the local window using passive target synchronization have completed. Therefore, it is synchronizing unless the window was created with the "no_locks" info key set to "true".

12.2.6 Window Attributes

The following attributes are cached with a window when the window is created.

MPI_WIN_BASE	window base address.
MPI_WIN_SIZE	window size, in bytes.
MPI_WIN_DISP_UNIT	displacement unit associated with the window.
MPI_WIN_CREATE_FLAVOR	how the window was created.
MPI_WIN_MODEL	memory model for window.

In C, calls such as MPI_Win_get_attr(win, MPI_WIN_BASE, &base, &flag), MPI_Win_get_attr(win, MPI_WIN_SIZE, &size, &flag), MPI_Win_get_attr(win, MPI_WIN_DISP_UNIT, &disp_unit, &flag), MPI_Win_get_attr(win, MPI_WIN_CREATE_FLAVOR, &create_flavor, &flag), and MPI_Win_get_attr(win, MPI_WIN_MODEL, &memory_model, &flag) will return in base a pointer to the start of the window win and in size, disp_unit, create_kind, and memory_model pointers to the size of the window, the displacement unit of the window, the flavor of the window, and the memory model of the window, respectively. A detailed listing of the type of the pointer in the attribute value argument to MPI_WIN_GET_ATTR and MPI_WIN_SET_ATTR is shown in Table 12.1.

In Fortran, calls such as MPI_WIN_GET_ATTR(win, MPI_WIN_BASE, base, flag, ierror), MPI_WIN_GET_ATTR(win, MPI_WIN_SIZE, size, flag, ierror), MPI_WIN_GET_ATTR(win, MPI_WIN_DISP_UNIT, disp_unit, flag, ierror), MPI_WIN_GET_ATTR(win, MPI_WIN_CREATE_FLAVOR, create_flavor, flag, ierror), and MPI_WIN_GET_ATTR(win, MPI_WIN_MODEL, memory_model, flag, ierror) will return in base, size, disp_unit, create_kind, and memory_model the (integer representation of) the base address of the window, the size of the window, the displacement unit of the window, the flavor of the window, and the memory model of the window, respectively.

The values of create_kind are

MPI_WIN_FLAVOR_CREATE	Window was created with MPI_WIN_CREATE.
MPI_WIN_FLAVOR_ALLOCATE	Window was created with MPI_WIN_ALLOCATE.

1	MPI_WIN_FLAVOR_DYNAMIC	Window was created with
2		MPI_WIN_CREATE_DYNAMIC.
3	MPI_WIN_FLAVOR_SHARED	Window was created with
4		MPI_WIN_ALLOCATE_SHARED.

5 The values of `memory_model` are `MPI_WIN_SEPARATE` and `MPI_WIN_UNIFIED`. The mean-
6 ing of these is described in Section 12.4.

7 In the case of windows created with `MPI_WIN_CREATE_DYNAMIC`, the base address is
8 `MPI_BOTTOM` and the size is 0. In C, pointers are returned, and in Fortran, the values are
9 returned, for the respective attributes. (The window attribute access functions are defined
10 in Section 7.7.3.) The value returned for an attribute on a window is constant over the
11 lifetime of the window.

12 The other “window attribute,” namely the group of MPI processes attached to the
13 window, can be retrieved using the call below.

14
15
16 `MPI_WIN_GET_GROUP(win, group)`

17	IN	win	window object (handle)
18			
19	OUT	group	group of MPI processes that share access to the
20			window (handle)

21 C binding

22 `int MPI_Win_get_group(MPI_Win win, MPI_Group *group)`

23 Fortran 2008 binding

24 `MPI_Win_get_group(win, group, ierror)`
25 `TYPE(MPI_Win), INTENT(IN) :: win`
26 `TYPE(MPI_Group), INTENT(OUT) :: group`
27 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

28 Fortran binding

29 `MPI_WIN_GET_GROUP(WIN, GROUP, IERROR)`
30 `INTEGER WIN, GROUP, IERROR`

31 `MPI_WIN_GET_GROUP` returns in `group` a duplicate of the group of the communicator
32 used to create the window associated with `win`.

33 12.2.7 Window Info

34 Hints specified via `info` (see Section 10) allow a user to provide information to direct op-
35 timization. Providing hints may enable an implementation to deliver increased perfor-
36 mance or use system resources more efficiently. As described in Section 10, an imple-
37 mentation is free to ignore all hints; however, applications must comply with any `info`
38 hints they provide that are used by the MPI implementation (i.e., are returned by a call to
39 `MPI_WIN_GET_INFO`) and that place a restriction on the behavior of the application. Hints
40 are specified on a per window basis, in window creation functions and `MPI_WIN_SET_INFO`,
41 via the opaque `info` object. When an `info` object that specifies a subset of valid hints is passed
42 to `MPI_WIN_SET_INFO` there will be no effect on previously set or default hints that the
43 `info` does not specify.

Advice to implementors. It may happen that a program is coded with hints for one system, and later executes on another system that does not support these hints. In general, unsupported hints should simply be ignored. Needless to say, no hint can be mandatory. However, for each hint used by a specific implementation, a default value must be provided when the user does not specify a value for the hint. (*End of advice to implementors.*)

`MPI_WIN_SET_INFO(win, info)`

INOUT	win	window object (handle)
IN	info	info argument (handle)

C binding

```
int MPI_Win_set_info(MPI_Win win, MPI_Info info)
```

Fortran 2008 binding

```
MPI_Win_set_info(win, info, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  TYPE(MPI_Info), INTENT(IN) :: info
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_SET_INFO(WIN, INFO, IERROR)
  INTEGER WIN, INFO, IERROR
```

`MPI_WIN_SET_INFO` updates the hints of the window associated with `win` using the hints provided in `info`. This operation has no effect on previously set or defaulted hints that are not specified by `info`. It also has no effect on previously set or defaulted hints that are specified by `info`, but are ignored by the MPI implementation in this call to `MPI_WIN_SET_INFO`. The procedure is collective over the group of `win`. The entries in the `info` object may be different on each MPI process, but any `info` entries that an implementation requires to be the same on all MPI processes must appear with the same value in each MPI process's `info` object.

Advice to users. Some `info` items that an implementation can use when it creates a window cannot easily be changed once the window has been created. Thus, an implementation may ignore hints issued in this call that it would have accepted in a creation call. An implementation may also be unable to update certain `info` hints in a call to `MPI_WIN_SET_INFO`. `MPI_WIN_GET_INFO` can be used to determine whether `info` changes were ignored by the implementation. (*End of advice to users.*)

`MPI_WIN_GET_INFO(win, info_used)`

IN	win	window object (handle)
OUT	info_used	new info object (handle)

C binding

```
int MPI_Win_get_info(MPI_Win win, MPI_Info *info_used)
```

Fortran 2008 binding

```

MPI_Win_get_info(win, info_used, ierror)
    TYPE(MPI_Win), INTENT(IN) :: win
    TYPE(MPI_Info), INTENT(OUT) :: info_used
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_WIN_GET_INFO(WIN, INFO_USED, IERROR)
    INTEGER WIN, INFO_USED, IERROR

```

MPI_WIN_GET_INFO returns a new info object containing the hints of the window associated with win. The current setting of all hints related to this window is returned in info_used. An MPI implementation is required to return all hints that are supported by the implementation and have default values specified; any user-supplied hints that were not ignored by the implementation; and any additional hints that were set by the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no key/value pair. The user is responsible for freeing info_used via MPI_INFO_FREE.

12.3 Communication Calls

MPI supports the following RMA communication calls: MPI_PUT and MPI_RPUT transfer data from the caller memory (origin) to the target memory; MPI_GET and MPI_RGET transfer data from the target memory to the caller memory; MPI_ACCUMULATE and MPI_RACCUMULATE perform element-wise atomic updates of locations in the target memory, e.g., by adding to these locations values sent from the caller memory; MPI_GET_ACCUMULATE, MPI_RGET_ACCUMULATE, and MPI_FETCH_AND_OP perform element-wise atomic read-modify-write updates and return each value before the update; and MPI_COMPARE_AND_SWAP performs a remote atomic compare and swap operation. These procedures are *nonblocking*. The operation is completed, at the origin or both the origin and the target, when a subsequent *synchronization* procedure is called by the origin on the involved window object. These synchronization procedures are described in Section 12.5. RMA communication operations can also be completed with calls to flush procedures; see Section 12.5.4 for details. Request-based operations MPI_RPUT, MPI_RGET, MPI_RACCUMULATE, and MPI_RGET_ACCUMULATE can be completed at the origin by using the MPI test or wait procedures described in Section 3.7.3.

The local communication buffer of an RMA operation should not be updated after the operation started and until the operation completes at the origin. The local communication buffer of a get operation should not be accessed after the operation started and until the operation completes at the origin.

Two concurrent accesses are called conflicting if one of the two is a put operation, exactly one of them is an accumulate operation, or one of them is a get operation and the other is a local store access. The outcome of conflicting accesses to the same memory location is undefined; if a location is updated by a put or accumulate operation, then the outcome of loads or other RMA operations is undefined until the updating operation has completed at the target. There is one exception to this rule; namely, the same location can be updated by several concurrent accumulate operations, the outcome being as if these updates occurred in some order. In addition, the outcome of concurrent load/store accesses and RMA updates

to the same memory location is undefined. These restrictions are described in more detail in Section 12.7.

The calls use general datatype arguments to specify communication buffers at the origin and at the target. Thus, a transfer operation may also gather data at the source and scatter it at the destination. However, all arguments specifying both communication buffers are provided by the caller.

For all RMA communication operations, the target process may be identical with the origin process; i.e., an MPI process may use an RMA operation to move data in its memory.

Rationale. The choice of supporting “self-communication” is the same as for message-passing. It simplifies some coding, and is very useful with accumulate operations, to allow atomic updates of local variables. (*End of rationale.*)

MPI_PROC_NULL is a valid target rank in all MPI RMA communication calls. The effect is the same as for MPI_PROC_NULL in MPI point-to-point communication. After any RMA operation with rank MPI_PROC_NULL, it is still necessary to close the RMA epoch with the synchronization method that opened the epoch.

12.3.1 Put

The execution of a put operation is similar to the execution of a send by the origin process and a matching receive by the target process. The obvious difference is that all arguments are provided by one call—the call executed by the origin process.

MPI_PUT(origin_addr, origin_count, origin_datatype, target_rank, target_disp, target_count, target_datatype, win)

IN	origin_addr	initial address of origin buffer (choice)
IN	origin_count	number of entries in origin buffer (non-negative integer)
IN	origin_datatype	datatype of each entry in origin buffer (handle)
IN	target_rank	rank of target (non-negative integer)
IN	target_disp	displacement from start of window to target buffer (non-negative integer)
IN	target_count	number of entries in target buffer (non-negative integer)
IN	target_datatype	datatype of each entry in target buffer (handle)
IN	win	window used for communication (handle)

C binding

```
int MPI_Put(const void *origin_addr, int origin_count,
           MPI_Datatype origin_datatype, int target_rank,
           MPI_Aint target_disp, int target_count,
           MPI_Datatype target_datatype, MPI_Win win)

int MPI_Put_c(const void *origin_addr, MPI_Count origin_count,
             MPI_Datatype origin_datatype, int target_rank,
```

```

1         MPI_Aint target_disp, MPI_Count target_count,
2         MPI_Datatype target_datatype, MPI_Win win)
3

```

Fortran 2008 binding

```

4 MPI_Put(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
5         target_count, target_datatype, win, ierror)
6     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
7     INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
8     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
9     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
10    TYPE(MPI_Win), INTENT(IN) :: win
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12

```

```

13 MPI_Put(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
14         target_count, target_datatype, win, ierror) !(_c)
15    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
16    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
17    TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
18    INTEGER, INTENT(IN) :: target_rank
19    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
20    TYPE(MPI_Win), INTENT(IN) :: win
21    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22

```

Fortran binding

```

23 MPI_PUT(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
24         TARGET_COUNT, TARGET_DATATYPE, WIN, IERROR)
25    <type> ORIGIN_ADDR(*)
26    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
27         TARGET_DATATYPE, WIN, IERROR
28    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
29

```

Transfers `origin_count` successive entries of the type specified by the `origin_datatype`, starting at address `origin_addr` on the origin process, to the target process specified by the `win`, `target_rank` pair. The data are written in the target buffer at address `target_addr = window_base + target_disp × disp_unit`, where `window_base` and `disp_unit` are the base address and window displacement unit specified at window initialization, by the target process.

The target buffer is specified by the arguments `target_count` and `target_datatype`.

The data transfer is the same as that which would occur if the origin process executed a send operation with arguments `origin_addr`, `origin_count`, `origin_datatype`, `target_rank`, `tag`, `comm`, and the target process executed a receive operation with arguments `target_addr`, `target_count`, `target_datatype`, `source`, `tag`, `comm`, where `target_addr` is the target buffer address computed as explained above, the values of `tag` are arbitrary valid matching tag values, and `comm` is a communicator for the group of `win`.

The communication must satisfy the same constraints as for a similar message-passing communication. The `target_datatype` may not specify overlapping entries in the target buffer. The message sent must fit, without truncation, in the target buffer. Furthermore, the target buffer must fit in the target window or in attached memory in a dynamic window.

The `target_datatype` argument is a handle to a datatype object defined at the origin process. However, this object is interpreted at the target process: the outcome is as if the target datatype object was defined at the target process by the same sequence of calls

used to define it at the origin process. The target datatype must contain only relative displacements, not absolute addresses. The same holds for get and accumulate operations.

Advice to users. The `target_datatype` argument is a handle to a datatype object that is defined at the origin process, even though it defines a data layout in the target process memory. This causes no problems in a homogeneous environment, or in a heterogeneous environment if only portable datatypes are used (portable datatypes are defined in Section 2.4).

The performance of a put transfer can be significantly affected, on some systems, by the choice of window location and the shape and location of the origin and target buffer: transfers to a target window in memory allocated by `MPI_ALLOC_MEM` or `MPI_WIN_ALLOCATE` may be much faster on shared memory systems; transfers from contiguous buffers will be faster on most, if not all, systems; the alignment of the communication buffers may also impact performance. (*End of advice to users.*)

Advice to implementors. A high-quality implementation will attempt to prevent remote accesses to memory outside the window that was exposed by the MPI process. This is important both for debugging purposes and for protection with client-server codes that use RMA. That is, a high-quality implementation will check, if possible, window bounds on each RMA call, and raise an error at the origin call if an out-of-bound situation occurs. Note that the condition can be checked at the origin. Of course, the added safety achieved by such checks has to be weighed against the added cost of such checks. (*End of advice to implementors.*)

12.3.2 Get

`MPI_GET(origin_addr, origin_count, origin_datatype, target_rank, target_disp, target_count, target_datatype, win)`

OUT	<code>origin_addr</code>	initial address of origin buffer (choice)
IN	<code>origin_count</code>	number of entries in origin buffer (non-negative integer)
IN	<code>origin_datatype</code>	datatype of each entry in origin buffer (handle)
IN	<code>target_rank</code>	rank of target (non-negative integer)
IN	<code>target_disp</code>	displacement from window start to the beginning of the target buffer (non-negative integer)
IN	<code>target_count</code>	number of entries in target buffer (non-negative integer)
IN	<code>target_datatype</code>	datatype of each entry in target buffer (handle)
IN	<code>win</code>	window used for communication (handle)

C binding

```
int MPI_Get(void *origin_addr, int origin_count, MPI_Datatype origin_datatype,
            int target_rank, MPI_Aint target_disp, int target_count,
            MPI_Datatype target_datatype, MPI_Win win)
```

```

1 int MPI_Get_c(void *origin_addr, MPI_Count origin_count,
2             MPI_Datatype origin_datatype, int target_rank,
3             MPI_Aint target_disp, MPI_Count target_count,
4             MPI_Datatype target_datatype, MPI_Win win)

```

Fortran 2008 binding

```

6 MPI_Get(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
7         target_count, target_datatype, win, ierror)
8     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr
9     INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
10    TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
11    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
12    TYPE(MPI_Win), INTENT(IN) :: win
13    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14
15 MPI_Get(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
16         target_count, target_datatype, win, ierror) !(_c)
17    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr
18    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
19    TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
20    INTEGER, INTENT(IN) :: target_rank
21    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
22    TYPE(MPI_Win), INTENT(IN) :: win
23    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

25 MPI_GET(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
26        TARGET_COUNT, TARGET_DATATYPE, WIN, IERROR)
27    <type> ORIGIN_ADDR(*)
28    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
29        TARGET_DATATYPE, WIN, IERROR
30    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
31

```

Similar to MPI_PUT, except that the direction of data transfer is reversed. Data are copied from the target memory to the origin. The `origin_datatype` may not specify overlapping entries in the origin buffer. The target buffer must be contained within the target window or within attached memory in a dynamic window, and the copied data must fit, without truncation, in the origin buffer.

12.3.3 Examples for Communication Calls

These examples show the use of the MPI_GET procedure. As all MPI RMA communication procedures are nonblocking, the associated operations must be completed by subsequent calls to synchronization procedures. In the following example, completion is accomplished with the routine MPI_WIN_FENCE, introduced in Section 12.5.

Example 12.1. We show how to implement the generic indirect assignment $A = B(\text{map})$, where A , B , and map have the same distribution, and map is a permutation. To simplify, we assume a block distribution with equal size blocks.

```

47 SUBROUTINE MAPVALS(A, B, map, m, comm, p)
48

```

```

USE MPI
INTEGER m, map(m), comm, p
REAL A(m), B(m)

INTEGER otype(p), oindex(m), & ! used to construct origin datatypes
      ttype(p), tindex(m), & ! used to construct target datatypes
      count(p), total(p), &
      disp_int, win, ierr, i, j, k
INTEGER(KIND=MPI_ADDRESS_KIND) lowerbound, size, realextent, disp_aint

! This part does the work that depends on the locations of B.
! Can be reused while this does not change

CALL MPI_TYPE_GET_EXTENT(MPI_REAL, lowerbound, realextent, ierr)
disp_int = realextent
size = m * realextent
CALL MPI_WIN_CREATE(B, size, disp_int, MPI_INFO_NULL, &
      comm, win, ierr)

! This part does the work that depends on the value of map and
! the locations of the arrays.
! Can be reused while these do not change

! Compute number of entries to be received from each process

DO i=1,p
  count(i) = 0
END DO
DO i=1,m
  j = map(i)/m+1
  count(j) = count(j)+1
END DO

total(1) = 0
DO i=2,p
  total(i) = total(i-1) + count(i-1)
END DO

DO i=1,p
  count(i) = 0
END DO

! compute origin and target indices of entries.
! entry i at current process is received from location
! k at process (j-1), where map(i) = (j-1)*m + (k-1),
! j = 1..p and k = 1..m

DO i=1,m
  j = map(i)/m+1
  k = MOD(map(i),m)+1
  count(j) = count(j)+1
  oindex(total(j) + count(j)) = i
  tindex(total(j) + count(j)) = k
END DO

! create origin and target datatypes for each get operation
DO i=1,p

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

1  CALL MPI_TYPE_CREATE_INDEXED_BLOCK(count(i), 1, &
2  oindex(total(i)+1:total(i)+count(i)), &
3  MPI_REAL, otype(i), ierr)
4  CALL MPI_TYPE_COMMIT(otype(i), ierr)
5  CALL MPI_TYPE_CREATE_INDEXED_BLOCK(count(i), 1, &
6  tindex(total(i)+1:total(i)+count(i)), &
7  MPI_REAL, ttype(i), ierr)
8  CALL MPI_TYPE_COMMIT(ttype(i), ierr)
9  END DO
10 ! this part does the assignment itself
11 CALL MPI_WIN_FENCE(0, win, ierr)
12 disp_aint = 0
13 DO i=1,p
14   CALL MPI_GET(A, 1, otype(i), i-1, disp_aint, 1, ttype(i), win, ierr)
15 END DO
16 CALL MPI_WIN_FENCE(0, win, ierr)
17 CALL MPI_WIN_FREE(win, ierr)
18 DO i=1,p
19   CALL MPI_TYPE_FREE(otype(i), ierr)
20   CALL MPI_TYPE_FREE(ttype(i), ierr)
21 END DO
22 RETURN
23 END

```

Example 12.2. A simpler version can be written that does not require that a datatype be built for the target buffer. One then needs a separate get operation for each entry, as illustrated below. This code is much simpler, but usually much less efficient, for large arrays.

```

27 SUBROUTINE MAPVALS(A, B, map, m, comm, p)
28 USE MPI
29 INTEGER m, map(m), comm, p
30 REAL A(m), B(m)
31 INTEGER disp_int, i, j, win, ierr
32 INTEGER(KIND=MPI_ADDRESS_KIND) lowerbound, size, realextent, disp_aint
33 CALL MPI_TYPE_GET_EXTENT(MPI_REAL, lowerbound, realextent, ierr)
34 disp_int = realextent
35 size = m * realextent
36 CALL MPI_WIN_CREATE(B, size, disp_int, MPI_INFO_NULL, &
37 comm, win, ierr)
38 CALL MPI_WIN_FENCE(0, win, ierr)
39 DO i=1,m
40   j = map(i)/m
41   disp_aint = MOD(map(i),m)
42   CALL MPI_GET(A(i), 1, MPI_REAL, j, disp_aint, 1, MPI_REAL, win, ierr)
43 END DO
44 CALL MPI_WIN_FENCE(0, win, ierr)
45 CALL MPI_WIN_FREE(win, ierr)
46 RETURN
47 END

```

12.3.4 Accumulate Functions

It is often useful in a put operation to combine the data moved to the target process with the data that resides at that MPI process, rather than replacing it. This will allow, for example, the accumulation of a sum by having all involved MPI processes add their contributions to the sum variable in the memory of one MPI process. The accumulate functions have slightly different semantics with respect to overlapping data accesses than the put and get functions; see Section 12.7 for details.

Accumulate

```

MPI_ACCUMULATE(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
                target_count, target_datatype, op, win)
IN      origin_addr          initial address of buffer (choice)
IN      origin_count         number of entries in buffer (non-negative integer)
IN      origin_datatype      datatype of each entry (handle)
IN      target_rank          rank of target (non-negative integer)
IN      target_disp          displacement from start of window to beginning of
                             target buffer (non-negative integer)
IN      target_count         number of entries in target buffer (non-negative
                             integer)
IN      target_datatype      datatype of each entry in target buffer (handle)
IN      op                   accumulate operator (handle)
IN      win                  window object (handle)

```

C binding

```

int MPI_Accumulate(const void *origin_addr, int origin_count,
                  MPI_Datatype origin_datatype, int target_rank,
                  MPI_Aint target_disp, int target_count,
                  MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
int MPI_Accumulate_c(const void *origin_addr, MPI_Count origin_count,
                    MPI_Datatype origin_datatype, int target_rank,
                    MPI_Aint target_disp, MPI_Count target_count,
                    MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)

```

Fortran 2008 binding

```

MPI_Accumulate(origin_addr, origin_count, origin_datatype, target_rank,
               target_disp, target_count, target_datatype, op, win, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Win), INTENT(IN) :: win

```

```

1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2 MPI_Accumulate(origin_addr, origin_count, origin_datatype, target_rank,
3               target_disp, target_count, target_datatype, op, win, ierror)
4               !(_c)
5     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
6     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
7     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
8     INTEGER, INTENT(IN) :: target_rank
9     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
10    TYPE(MPI_Op), INTENT(IN) :: op
11    TYPE(MPI_Win), INTENT(IN) :: win
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

14 MPI_ACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK,
15               TARGET_DISP, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, IERROR)
16 <type> ORIGIN_ADDR(*)
17 INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
18               TARGET_DATATYPE, OP, WIN, IERROR
19 INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP

```

Accumulate the contents of the origin buffer (as defined by `origin_addr`, `origin_count`, and `origin_datatype`) to the buffer specified by arguments `target_count` and `target_datatype`, at offset `target_disp`, in the target window specified by `target_rank` and `win`, using the operator `op`. This is like `MPI_PUT` except that data is combined into the target area instead of overwriting it.

Any of the predefined operators for `MPI_REDUCE` can be used. User-defined operators cannot be used. For example, if `op` is `MPI_SUM`, each element of the origin buffer is added to the corresponding element in the target, replacing the former value in the target.

Each datatype argument must be a predefined datatype or a derived datatype, where all basic components are of the same predefined datatype. Both datatype arguments must be constructed from the same predefined datatype. The operator `op` applies to elements of that predefined type. The parameter `target_datatype` must not specify overlapping entries, and the target buffer must fit in the target window.

An additional predefined operator, `MPI_REPLACE`, is defined. It corresponds to the associative function $f(a, b) = b$; i.e., the current value in the target memory is replaced by the value supplied by the origin.

`MPI_REPLACE` can be used only in `MPI_ACCUMULATE`, `MPI_RACCUMULATE`, `MPI_GET_ACCUMULATE`, `MPI_FETCH_AND_OP`, and `MPI_RGET_ACCUMULATE`, but not in collective reduction operations such as `MPI_REDUCE`.

Advice to users. `MPI_PUT` can be considered a special case of `MPI_ACCUMULATE` with the operator `MPI_REPLACE`. Note, however, that `MPI_PUT` and `MPI_ACCUMULATE` have different constraints on concurrent updates. (*End of advice to users.*)

Example 12.3. We want to compute $B(j) = \sum_{\text{map}(i)=j} A(i)$. The arrays `A`, `B`, and `map` are distributed in the same manner. We write the simple version.


```

1  SUBROUTINE SUM(A, B, map, m, comm, p)
2  USE MPI
3  INTEGER m, map(m), comm, p, win, ierr, disp_int, i, j
4  REAL A(m), B(m)
5  INTEGER(KIND=MPI_ADDRESS_KIND) lowerbound, size, realextent, disp_aint
6
7  CALL MPI_TYPE_GET_EXTENT(MPI_REAL, lowerbound, realextent, ierr)
8  size = m * realextent
9  disp_int = realextent
10 CALL MPI_WIN_CREATE(B, size, disp_int, MPI_INFO_NULL, &
11                    comm, win, ierr)
12
13 CALL MPI_WIN_FENCE(0, win, ierr)
14 DO i=1,m
15   j = map(i)/m
16   disp_aint = MOD(map(i),m)
17   CALL MPI_ACCUMULATE(A(i), 1, MPI_REAL, j, disp_aint, 1, MPI_REAL, &
18                     MPI_SUM, win, ierr)
19
20 END DO
21 CALL MPI_WIN_FENCE(0, win, ierr)
22
23 CALL MPI_WIN_FREE(win, ierr)
24 RETURN
25 END

```

This code is identical to the code in Example 12.2, except that a call to `MPI_GET` has been replaced by a call to `MPI_ACCUMULATE`. (Note that, if `map` is one-to-one, the code computes $B = A(\text{map}^{-1})$, which is the reverse assignment to the one computed in that previous example.) In a similar manner, we can replace in Example 12.1, the call to `get` by a call to `accumulate`, thus performing the computation with only one communication between any two MPI processes.

Get Accumulate

It is often useful to have fetch-and-accumulate semantics such that the remote data is returned to the caller before the sent data is accumulated into the remote data. The `get` and `accumulate` steps are executed atomically for each basic element in the datatype (see Section 12.7 for details). The predefined operator `MPI_REPLACE` provides fetch-and-set behavior.

```

MPI_GET_ACCUMULATE(origin_addr, origin_count, origin_datatype, result_addr,
                   result_count, result_datatype, target_rank, target_disp, target_count,
                   target_datatype, op, win)

```

IN	<code>origin_addr</code>	initial address of buffer (choice)
IN	<code>origin_count</code>	number of entries in origin buffer (non-negative integer)
IN	<code>origin_datatype</code>	datatype of each entry in origin buffer (handle)
OUT	<code>result_addr</code>	initial address of result buffer (choice)

1	IN	result_count	number of entries in result buffer (non-negative integer)
2			
3	IN	result_datatype	datatype of each entry in result buffer (handle)
4			
5	IN	target_rank	rank of target (non-negative integer)
6	IN	target_disp	displacement from start of window to beginning of target buffer (non-negative integer)
7			
8	IN	target_count	number of entries in target buffer (non-negative integer)
9			
10			
11	IN	target_datatype	datatype of each entry in target buffer (handle)
12	IN	op	accumulate operator (handle)
13	IN	win	window object (handle)
14			

C binding

```

16 int MPI_Get_accumulate(const void *origin_addr, int origin_count,
17                       MPI_Datatype origin_datatype, void *result_addr,
18                       int result_count, MPI_Datatype result_datatype, int target_rank,
19                       MPI_Aint target_disp, int target_count,
20                       MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
21
22 int MPI_Get_accumulate_c(const void *origin_addr, MPI_Count origin_count,
23                         MPI_Datatype origin_datatype, void *result_addr,
24                         MPI_Count result_count, MPI_Datatype result_datatype,
25                         int target_rank, MPI_Aint target_disp, MPI_Count target_count,
26                         MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
27

```

Fortran 2008 binding

```

28 MPI_Get_accumulate(origin_addr, origin_count, origin_datatype, result_addr,
29                   result_count, result_datatype, target_rank, target_disp,
30                   target_count, target_datatype, op, win, ierror)
31
32   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
33   INTEGER, INTENT(IN) :: origin_count, result_count, target_rank,
34                   target_count
35   TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,
36                   target_datatype
37   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr
38   INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
39   TYPE(MPI_Op), INTENT(IN) :: op
40   TYPE(MPI_Win), INTENT(IN) :: win
41   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Get_accumulate(origin_addr, origin_count, origin_datatype, result_addr,
44                   result_count, result_datatype, target_rank, target_disp,
45                   target_count, target_datatype, op, win, ierror) !(_c)
46   TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
47   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, result_count,
48                   target_count

```

```

TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,
    target_datatype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr
INTEGER, INTENT(IN) :: target_rank
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_GET_ACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_ADDR,
    RESULT_COUNT, RESULT_DATATYPE, TARGET_RANK, TARGET_DISP,
    TARGET_COUNT, TARGET_DATATYPE, OP, WIN, IERROR)
<type> ORIGIN_ADDR(*), RESULT_ADDR(*)
INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_COUNT, RESULT_DATATYPE,
    TARGET_RANK, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, IERROR
INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP

```

Accumulate `origin_count` elements of type `origin_datatype` from the origin buffer (`origin_addr`) to the buffer at offset `target_disp`, in the target window specified by `target_rank` and `win`, using the operator `op` and return in the result buffer `result_addr` the content of the target buffer before the accumulation, specified by `target_disp`, `target_count`, and `target_datatype`. The data transferred from origin to target must fit, without truncation, in the target buffer. Likewise, the data copied from target to origin must fit, without truncation, in the result buffer.

The origin and result buffers (`origin_addr` and `result_addr`) must be disjoint. Each datatype argument must be a predefined datatype or a derived datatype where all basic components are of the same predefined datatype. All datatype arguments must be constructed from the same predefined datatype. The operator `op` applies to elements of that predefined type. `target_datatype` must not specify overlapping entries, and the target buffer must fit in the target window or in attached memory in a dynamic window. The operation is executed atomically for each basic datatype; see Section 12.7 for details.

Any of the predefined operators for `MPI_REDUCE`, as well as `MPI_NO_OP` or `MPI_REPLACE` can be specified as `op`. User-defined functions cannot be used. An additional predefined operator, `MPI_NO_OP`, is defined. It corresponds to the associative function $f(a,b) = a$; i.e., the current value in the target memory is returned in the result buffer at the origin and the target buffer is not updated. If `MPI_NO_OP` is specified as the operator, the `origin_addr`, `origin_count`, and `origin_datatype` arguments are ignored. `MPI_NO_OP` can be used only in `MPI_GET_ACCUMULATE`, `MPI_RGET_ACCUMULATE`, and `MPI_FETCH_AND_OP`. `MPI_NO_OP` cannot be used in `MPI_ACCUMULATE`, `MPI_RACCUMULATE`, or collective reduction operations, such as `MPI_REDUCE` and others.

Advice to users. `MPI_GET` is similar to `MPI_GET_ACCUMULATE`, with the operator `MPI_NO_OP`. Note, however, that `MPI_GET` and `MPI_GET_ACCUMULATE` have different constraints on concurrent updates. (*End of advice to users.*)

Fetch and Op

The generic functionality of `MPI_GET_ACCUMULATE` might limit the performance of fetch-and-increment or fetch-and-add calls that might be supported by special hardware oper-

ations. `MPI_FETCH_AND_OP` thus allows for a fast implementation of a commonly used subset of the functionality of `MPI_GET_ACCUMULATE`.

```
MPI_FETCH_AND_OP(origin_addr, result_addr, datatype, target_rank, target_disp, op, win)
```

IN	<code>origin_addr</code>	initial address of buffer (choice)
OUT	<code>result_addr</code>	initial address of result buffer (choice)
IN	<code>datatype</code>	datatype of the entry in origin, result, and target buffers (handle)
IN	<code>target_rank</code>	rank of target (non-negative integer)
IN	<code>target_disp</code>	displacement from start of window to beginning of target buffer (non-negative integer)
IN	<code>op</code>	accumulate operator (handle)
IN	<code>win</code>	window object (handle)

C binding

```
int MPI_Fetch_and_op(const void *origin_addr, void *result_addr,
                    MPI_Datatype datatype, int target_rank, MPI_Aint target_disp,
                    MPI_Op op, MPI_Win win)
```

Fortran 2008 binding

```
MPI_Fetch_and_op(origin_addr, result_addr, datatype, target_rank, target_disp,
                 op, win, ierror)
    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr
    TYPE(MPI_Datatype), INTENT(IN) :: datatype
    INTEGER, INTENT(IN) :: target_rank
    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
    TYPE(MPI_Op), INTENT(IN) :: op
    TYPE(MPI_Win), INTENT(IN) :: win
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FETCH_AND_OP(ORIGIN_ADDR, RESULT_ADDR, DATATYPE, TARGET_RANK, TARGET_DISP,
                 OP, WIN, IERROR)
    <type> ORIGIN_ADDR(*), RESULT_ADDR(*)
    INTEGER DATATYPE, TARGET_RANK, OP, WIN, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
```

Accumulate one element of type `datatype` from the origin buffer `origin_addr` to the buffer at offset `target_disp`, in the target window specified by `target_rank` and `win`, using the operator `op` and return in the result buffer `result_addr` the content of the target buffer before the accumulation.

The origin and result buffers (`origin_addr` and `result_addr`) must be disjoint. Any of the predefined operators for `MPI_REDUCE`, as well as `MPI_NO_OP` or `MPI_REPLACE`, can be specified as `op`; user-defined functions cannot be used. The `datatype` argument must be a predefined datatype. The operation is executed atomically.

Compare and Swap

Another useful operation is an atomic compare and swap where the value at the origin is compared to the value at the target, which is atomically replaced by a third value only if the values at origin and target are equal.

`MPI_COMPARE_AND_SWAP(origin_addr, compare_addr, result_addr, datatype, target_rank, target_disp, win)`

IN	<code>origin_addr</code>	initial address of buffer (choice)
IN	<code>compare_addr</code>	initial address of compare buffer (choice)
OUT	<code>result_addr</code>	initial address of result buffer (choice)
IN	<code>datatype</code>	datatype of the element in all buffers (handle)
IN	<code>target_rank</code>	rank of target (non-negative integer)
IN	<code>target_disp</code>	displacement from start of window to beginning of target buffer (non-negative integer)
IN	<code>win</code>	window object (handle)

C binding

```
int MPI_Compare_and_swap(const void *origin_addr, const void *compare_addr,
                        void *result_addr, MPI_Datatype datatype, int target_rank,
                        MPI_Aint target_disp, MPI_Win win)
```

Fortran 2008 binding

```
MPI_Compare_and_swap(origin_addr, compare_addr, result_addr, datatype,
                    target_rank, target_disp, win, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr,
                    compare_addr
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: target_rank
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_COMPARE_AND_SWAP(ORIGIN_ADDR, COMPARE_ADDR, RESULT_ADDR, DATATYPE,
                    TARGET_RANK, TARGET_DISP, WIN, IERROR)
<type> ORIGIN_ADDR(*), COMPARE_ADDR(*), RESULT_ADDR(*)
INTEGER DATATYPE, TARGET_RANK, WIN, IERROR
INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
```

This function compares one element of type `datatype` in the compare buffer `compare_addr` with the buffer at offset `target_disp` in the target window specified by `target_rank` and `win` and replaces the value at the target with the value in the origin buffer `origin_addr` if the compare buffer and the target buffer are identical. The original value at the target is returned in the buffer `result_addr`. The parameter `datatype` must belong to one of the following categories of predefined datatypes: C integer, Fortran integer, Logical,

Multi-language types, or Byte as specified in Section 6.9.2. The origin and result buffers (origin_addr and result_addr) must be disjoint.

12.3.5 Request-based RMA Communication Operations

Request-based RMA communication operations allow the user to associate a request handle with the RMA operations and test or wait for the completion of these requests using the functions described in Section 3.7.3. Request-based RMA operations are only valid within a passive target epoch (see Section 12.5).

Upon returning from a completion call in which an RMA operation completes, all fields of the status object, if any, and the results of status query functions (e.g., MPI_GET_COUNT) are undefined with the exception of MPI_ERROR if appropriate (see Section 3.2.5). It is valid to mix different request types (e.g., any combination of RMA requests, collective requests, I/O requests, generalized requests, or point-to-point requests) in functions that enable multiple completions (e.g., MPI_WAITALL). It is erroneous to call MPI_REQUEST_FREE or MPI_CANCEL for a request associated with an RMA operation. RMA requests are not persistent.

The closing of the epoch, or explicit bulk synchronization using MPI_WIN_FLUSH, MPI_WIN_FLUSH_ALL, MPI_WIN_FLUSH_LOCAL, or MPI_WIN_FLUSH_LOCAL_ALL, also indicates completion of request-based RMA operations on the specified window. However, users must still free the request by testing, waiting, or calling MPI_REQUEST_FREE on the request handle to allow the MPI implementation to release any resources associated with these requests.

MPI_RPUT(origin_addr, origin_count, origin_datatype, target_rank, target_disp, target_count, target_datatype, win, request)

IN	origin_addr	initial address of origin buffer (choice)
IN	origin_count	number of entries in origin buffer (non-negative integer)
IN	origin_datatype	datatype of each entry in origin buffer (handle)
IN	target_rank	rank of target (non-negative integer)
IN	target_disp	displacement from start of window to target buffer (non-negative integer)
IN	target_count	number of entries in target buffer (non-negative integer)
IN	target_datatype	datatype of each entry in target buffer (handle)
IN	win	window used for communication (handle)
OUT	request	RMA request (handle)

C binding

```
int MPI_Rput(const void *origin_addr, int origin_count,
            MPI_Datatype origin_datatype, int target_rank,
            MPI_Aint target_disp, int target_count,
            MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)
```

```

int MPI_Rput_c(const void *origin_addr, MPI_Count origin_count,
              MPI_Datatype origin_datatype, int target_rank,
              MPI_Aint target_disp, MPI_Count target_count,
              MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Rput(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
        target_count, target_datatype, win, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
TYPE(MPI_Win), INTENT(IN) :: win
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Rput(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
        target_count, target_datatype, win, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
INTEGER, INTENT(IN) :: target_rank
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
TYPE(MPI_Win), INTENT(IN) :: win
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_RPUT(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
        TARGET_COUNT, TARGET_DATATYPE, WIN, REQUEST, IERROR)
<type> ORIGIN_ADDR(*)
INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
        TARGET_DATATYPE, WIN, REQUEST, IERROR
INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP

```

MPI_RPUT is similar to MPI_PUT (Section 12.3.1), except that it allocates a communication request object and associates it with the request handle (the argument `request`). The completion of the operation at the origin (i.e., after the corresponding test or wait) indicates that the sender is now free to update the locations in the origin buffer. It does not indicate that the data is available at the target window. If remote completion is required, MPI_WIN_FLUSH, MPI_WIN_FLUSH_ALL, MPI_WIN_UNLOCK, or MPI_WIN_UNLOCK_ALL can be used.

```

MPI_RGET(origin_addr, origin_count, origin_datatype, target_rank, target_disp, target_count,
        target_datatype, win, request)

```

OUT	origin_addr	initial address of origin buffer (choice)
IN	origin_count	number of entries in origin buffer (non-negative integer)

1	IN	origin_datatype	datatype of each entry in origin buffer (handle)
2	IN	target_rank	rank of target (non-negative integer)
3			
4	IN	target_disp	displacement from window start to the beginning of the target buffer (non-negative integer)
5			
6	IN	target_count	number of entries in target buffer (non-negative integer)
7			
8	IN	target_datatype	datatype of each entry in target buffer (handle)
9			
10	IN	win	window used for communication (handle)
11	OUT	request	RMA request (handle)
12			

C binding

```

14 int MPI_Rget(void *origin_addr, int origin_count, MPI_Datatype origin_datatype,
15             int target_rank, MPI_Aint target_disp, int target_count,
16             MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)
17
18 int MPI_Rget_c(void *origin_addr, MPI_Count origin_count,
19               MPI_Datatype origin_datatype, int target_rank,
20               MPI_Aint target_disp, MPI_Count target_count,
21               MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)

```

Fortran 2008 binding

```

23 MPI_Rget(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
24          target_count, target_datatype, win, request, ierror)
25     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr
26     INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
27     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
28     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
29     TYPE(MPI_Win), INTENT(IN) :: win
30     TYPE(MPI_Request), INTENT(OUT) :: request
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Rget(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
34          target_count, target_datatype, win, request, ierror) !(_c)
35     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr
36     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
37     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
38     INTEGER, INTENT(IN) :: target_rank
39     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
40     TYPE(MPI_Win), INTENT(IN) :: win
41     TYPE(MPI_Request), INTENT(OUT) :: request
42     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

44 MPI_RGET(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
45          TARGET_COUNT, TARGET_DATATYPE, WIN, REQUEST, IERROR)
46     <type> ORIGIN_ADDR(*)
47     INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
48     TARGET_DATATYPE, WIN, REQUEST, IERROR

```


INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP

MPI_RGET is similar to MPI_GET (Section 12.3.2), except that it allocates a communication request object and associates it with the request handle (the argument `request`) that can be used to wait or test for completion of the operation at the origin, which indicates that the data is available in the origin buffer. If `origin_addr` points to memory attached to a window, then the data becomes available in the private copy of this window.

MPI_RACCUMULATE(`origin_addr`, `origin_count`, `origin_datatype`, `target_rank`, `target_disp`, `target_count`, `target_datatype`, `op`, `win`, `request`)

IN	<code>origin_addr</code>	initial address of buffer (choice)
IN	<code>origin_count</code>	number of entries in buffer (non-negative integer)
IN	<code>origin_datatype</code>	datatype of each entry in origin buffer (handle)
IN	<code>target_rank</code>	rank of target (non-negative integer)
IN	<code>target_disp</code>	displacement from start of window to beginning of target buffer (non-negative integer)
IN	<code>target_count</code>	number of entries in target buffer (non-negative integer)
IN	<code>target_datatype</code>	datatype of each entry in target buffer (handle)
IN	<code>op</code>	accumulate operator (handle)
IN	<code>win</code>	window object (handle)
OUT	<code>request</code>	RMA request (handle)

C binding

```
int MPI_Raccumulate(const void *origin_addr, int origin_count,
                   MPI_Datatype origin_datatype, int target_rank,
                   MPI_Aint target_disp, int target_count,
                   MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
                   MPI_Request *request)
```

```
int MPI_Raccumulate_c(const void *origin_addr, MPI_Count origin_count,
                     MPI_Datatype origin_datatype, int target_rank,
                     MPI_Aint target_disp, MPI_Count target_count,
                     MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
                     MPI_Request *request)
```

Fortran 2008 binding

```
MPI_Raccumulate(origin_addr, origin_count, origin_datatype, target_rank,
                target_disp, target_count, target_datatype, op, win, request,
                ierror)
```

```
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
TYPE(MPI_Op), INTENT(IN) :: op
```

```

1     TYPE(MPI_Win), INTENT(IN) :: win
2     TYPE(MPI_Request), INTENT(OUT) :: request
3     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
4
5     MPI_Raccumulate(origin_addr, origin_count, origin_datatype, target_rank,
6                   target_disp, target_count, target_datatype, op, win, request,
7                   ierror) !(_c)
8     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
9     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
10    TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
11    INTEGER, INTENT(IN) :: target_rank
12    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
13    TYPE(MPI_Op), INTENT(IN) :: op
14    TYPE(MPI_Win), INTENT(IN) :: win
15    TYPE(MPI_Request), INTENT(OUT) :: request
16    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

17    MPI_RACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK,
18                  TARGET_DISP, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, REQUEST,
19                  IERROR)
20
21    <type> ORIGIN_ADDR(*)
22    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
23          TARGET_DATATYPE, OP, WIN, REQUEST, IERROR
24    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP

```

MPI_RACCUMULATE is similar to MPI_ACCUMULATE (Section 12.3.4), except that it allocates a communication request object and associates it with the request handle (the argument `request`) that can be used to wait or test for completion. The completion of the operation at the origin (i.e., after the corresponding test or wait) indicates that the origin buffer is free to be updated. It does not indicate that the operation has completed at the target window.

```

33    MPI_RGET_ACCUMULATE(origin_addr, origin_count, origin_datatype, result_addr,
34                      result_count, result_datatype, target_rank, target_disp, target_count,
35                      target_datatype, op, win, request)
36
37    IN      origin_addr      initial address of buffer (choice)
38    IN      origin_count     number of entries in origin buffer (non-negative
39                          integer)
40    IN      origin_datatype  datatype of each entry in origin buffer (handle)
41    OUT     result_addr      initial address of result buffer (choice)
42    IN      result_count     number of entries in result buffer (non-negative
43                          integer)
44    IN      result_datatype  datatype of entries in result buffer (handle)
45    IN      target_rank      rank of target (non-negative integer)

```

48

IN	target_disp	displacement from start of window to beginning of target buffer (non-negative integer)	1 2
IN	target_count	number of entries in target buffer (non-negative integer)	3 4 5
IN	target_datatype	datatype of each entry in target buffer (handle)	6
IN	op	accumulate operator (handle)	7
IN	win	window object (handle)	8 9
OUT	request	RMA request (handle)	10 11

C binding

```

int MPI_Rget_accumulate(const void *origin_addr, int origin_count,
    MPI_Datatype origin_datatype, void *result_addr,
    int result_count, MPI_Datatype result_datatype, int target_rank,
    MPI_Aint target_disp, int target_count,
    MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
    MPI_Request *request)

```

```

int MPI_Rget_accumulate_c(const void *origin_addr, MPI_Count origin_count,
    MPI_Datatype origin_datatype, void *result_addr,
    MPI_Count result_count, MPI_Datatype result_datatype,
    int target_rank, MPI_Aint target_disp, MPI_Count target_count,
    MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
    MPI_Request *request)

```

Fortran 2008 binding

```

MPI_Rget_accumulate(origin_addr, origin_count, origin_datatype, result_addr,
    result_count, result_datatype, target_rank, target_disp,
    target_count, target_datatype, op, win, request, ierror)

```

```

TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
INTEGER, INTENT(IN) :: origin_count, result_count, target_rank,
    target_count

```

```

TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,
    target_datatype

```

```

TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr

```

```

INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp

```

```

TYPE(MPI_Op), INTENT(IN) :: op

```

```

TYPE(MPI_Win), INTENT(IN) :: win

```

```

TYPE(MPI_Request), INTENT(OUT) :: request

```

```

INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Rget_accumulate(origin_addr, origin_count, origin_datatype, result_addr,
    result_count, result_datatype, target_rank, target_disp,
    target_count, target_datatype, op, win, request, ierror) !(_c)

```

```

TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr

```

```

INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, result_count,
    target_count

```

```

TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,
    target_datatype

```

```

1     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr
2     INTEGER, INTENT(IN) :: target_rank
3     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
4     TYPE(MPI_Op), INTENT(IN) :: op
5     TYPE(MPI_Win), INTENT(IN) :: win
6     TYPE(MPI_Request), INTENT(OUT) :: request
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

9     MPI_RGET_ACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_ADDR,
10                        RESULT_COUNT, RESULT_DATATYPE, TARGET_RANK, TARGET_DISP,
11                        TARGET_COUNT, TARGET_DATATYPE, OP, WIN, REQUEST, IERROR)
12     <type> ORIGIN_ADDR(*), RESULT_ADDR(*)
13     INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_COUNT, RESULT_DATATYPE,
14            TARGET_RANK, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, REQUEST,
15            IERROR
16     INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP

```

MPI_RGET_ACCUMULATE is similar to MPI_GET_ACCUMULATE (Section 12.3.4), except that it allocates a communication request object and associates it with the request handle (the argument `request`) that can be used to wait or test for completion. The completion of the operation at the origin (i.e., after the corresponding test or wait) indicates that the data is available in the result buffer and the origin buffer is free to be updated. It does not indicate that the operation has been completed at the target window.

12.4 Memory Model

The memory semantics of RMA are best understood by using the concept of *public* and *private* window copies. We assume that systems have a public memory region that is addressable by all MPI processes (e.g., the shared memory in shared memory machines or the exposed main memory in distributed memory machines). In addition, most machines have fast private buffers (e.g., transparent caches or explicit communication buffers) local to each MPI process where copies of data elements from the main memory can be stored for faster access. Such buffers are either coherent, i.e., all updates to main memory are reflected in all private copies consistently, or noncoherent, i.e., conflicting accesses to main memory need to be synchronized and updated in all private copies explicitly. Coherent systems allow direct updates to remote memory without any participation of the remote side. Noncoherent systems, however, need to call RMA functions in order to reflect updates to the public window in their private memory. Thus, in coherent memory, the public and the private window are identical while they remain logically separate in the noncoherent case. MPI thus differentiates between two **memory models** called **RMA unified**, if public and private window are logically identical, and **RMA separate**, otherwise.

In the RMA separate model, there is only one instance of each variable in MPI process memory, but a distinct *public* copy of the variable for each window that contains it. A load accesses the instance in MPI process memory (this includes MPI sends). A local store accesses and updates the instance in MPI process memory (this includes MPI receives), but the update may affect other public copies of the same locations. A get on a window accesses the public copy of that window. A put or accumulate on a window accesses and updates




Figure 12.1: Schematic description of the public/private window operations in the MPI_WIN_SEPARATE memory model for two overlapping windows

the public copy of that window, but the update may affect the private copy of the same locations in MPI process memory, and public copies of other overlapping windows. This is illustrated in Figure 12.1.

In the RMA unified model, public and private copies are identical and updates via put or accumulate operations are eventually observed by load accesses without additional RMA procedure calls. A store access to a window is eventually visible to remote get or accumulate operations without additional RMA procedure calls. These stronger semantics of the RMA unified model allow the user to omit some synchronization calls and potentially improve performance.

Advice to users. If accesses in the RMA unified model are not synchronized (with locks or flushes, see Section 12.5.3), load/store accesses might observe changes to the memory while they are in progress. The order in which data is written is not specified unless further synchronization is used. This might lead to inconsistent views on memory and programs that assume that a transfer is complete by only checking parts of the message are erroneous. (*End of advice to users.*)

The memory model for a particular RMA window can be determined by accessing the attribute MPI_WIN_MODEL. If the memory model is the unified model, the value of this attribute is MPI_WIN_UNIFIED; otherwise, the value is MPI_WIN_SEPARATE.

12.5 Synchronization Calls

RMA communications fall in two categories:

active target communication, where data is moved from the memory of one MPI process to the memory of another, and both are explicitly involved in the communication. This communication pattern is similar to message passing, except that all the data transfer

1 arguments are provided by the origin process, and the target process only participates
2 in the synchronization.

3
4 **passive target communication**, where data is moved from the memory of one MPI pro-
5 cess to the memory of another, and only the origin process is explicitly involved in the
6 transfer. Thus, two origin processes may communicate by accessing the same location
7 in a target window. The MPI process that owns the target window may be distinct
8 from the two communicating MPI processes, in which case it does not participate
9 explicitly in the communication. This communication paradigm is closest to a shared
10 memory model, where shared data can be accessed by all MPI processes, irrespective
11 of location.

12 RMA communication calls with argument `win` must occur at an origin process only
13 within an **access epoch** for `win`. Such an epoch is opened with an RMA synchronization
14 call on `win`; it proceeds with zero or more RMA communication calls (e.g., `MPI_PUT`,
15 `MPI_GET` or `MPI_ACCUMULATE`) on `win`; it is closed with another synchronization call on
16 `win`. This allows users to amortize one synchronization with multiple data transfers and
17 provide implementors more flexibility in the implementation of RMA operations.

18 Distinct access epochs for `win` at the same MPI process must be disjoint. On the other
19 hand, epochs pertaining to different `win` arguments may overlap. Load/store accesses or
20 other MPI calls may also occur during an epoch.

21 In active target communication, a target window can be accessed by RMA operations
22 only within an **exposure epoch**. Such an epoch is opened and closed by RMA synchroniza-
23 tion calls executed by the target process. Distinct exposure epochs at an MPI process on
24 the same window must be disjoint, but such an exposure epoch may overlap with exposure
25 epochs on other windows or with access epochs for the same or other window arguments.
26 There is a one-to-one matching between access epochs at origin processes and exposure
27 epochs on target processes: RMA operations issued by an origin process for a target win-
28 dow will access that target window during the same exposure epoch if and only if they were
29 issued during the same access epoch.

30 In passive target communication the target process does not execute RMA synchroni-
31 zation calls, and there is no concept of an exposure epoch.

32 MPI provides three synchronization mechanisms:

- 33
34 1. The `MPI_WIN_FENCE` collective synchronization call supports a simple synchroniza-
35 tion pattern that is often used in parallel computations: namely a loosely-synchronous
36 model, where global computation phases alternate with global communication phases.
37 This mechanism is most useful for loosely synchronous algorithms where the graph
38 of communicating MPI processes changes very frequently, or where each MPI process
39 communicates with many others.

40 This call is used for active target communication. An access epoch at an origin
41 process or an exposure epoch at a target process is opened and closed by calls to
42 `MPI_WIN_FENCE`. An origin process can access windows at all target processes in the
43 group of `win` during such an access epoch, and the local window can be accessed by
44 all MPI processes in the group of `win` during such an exposure epoch.

- 45
46 2. The four functions `MPI_WIN_START`, `MPI_WIN_COMPLETE`, `MPI_WIN_POST`, and
47 `MPI_WIN_WAIT` can be used to restrict synchronization to the minimum: only pairs
48

of communicating MPI processes synchronize, and they do so only when a synchronization is needed to order RMA accesses to a window correctly with respect to local accesses to that same window. This mechanism may be more efficient when each MPI process communicates with few (logical) neighbors, and the communication graph is fixed or changes infrequently.

These calls are used for active target communication. An access epoch is opened at the origin process with a call to `MPI_WIN_START` and is closed by a call to `MPI_WIN_COMPLETE`. The start call has a group argument that specifies the group of target processes for that epoch. An exposure epoch is opened at the target process by a call to `MPI_WIN_POST` and is closed by a call to `MPI_WIN_WAIT`. The post call has a group argument that specifies the set of origin processes for that epoch.

3. Finally, *shared lock* access is provided by the functions `MPI_WIN_LOCK`, `MPI_WIN_LOCK_ALL`, `MPI_WIN_UNLOCK`, and `MPI_WIN_UNLOCK_ALL`. `MPI_WIN_LOCK` and `MPI_WIN_UNLOCK` also provide *exclusive lock* capability. Lock synchronization is useful for MPI applications that emulate a shared memory model via MPI calls; e.g., in a “bulletin board” model, where MPI processes can, at random times, access or update different parts of the bulletin board.

These four calls provide passive target communication. An access epoch is opened by a call to `MPI_WIN_LOCK` or `MPI_WIN_LOCK_ALL` and closed by a call to `MPI_WIN_UNLOCK` or `MPI_WIN_UNLOCK_ALL`, respectively.

Figure 12.2 illustrates the general synchronization pattern for active target communication. The synchronization between `post` and `start` ensures that the put operation of the origin process does not start until the target process exposes the window (with the `post` call); the target process will expose the window only after preceding local accesses to the window have completed. The synchronization between `complete` and `wait` ensures that the put operation of the origin process completes at the origin and the target before the window is unexposed (with the `wait` call). The target process will execute subsequent local accesses to the target window only after the `wait` returned.

Figure 12.2 shows operations occurring in the natural temporal order implied by the synchronizations: the `post` occurs before the matching `start`, and `complete` occurs before the matching `wait`. However, such **strong synchronization** is more than needed for correct ordering of window accesses. The semantics of MPI calls allow **weak synchronization**, as illustrated in Figure 12.3. The access to the target window is delayed until the window is exposed, after the `post`. However the `start` may return before the exposure epoch opens at the target. Similarly, the `put` and `complete` calls may also return before the exposure epoch opens at the target, if put data is buffered by the implementation. The synchronization calls correctly order window accesses, but do not necessarily synchronize other operations. This weaker synchronization semantic allows for more efficient implementations.

Figure 12.4 illustrates the general synchronization pattern for passive target communication. The first origin process communicates data to the second origin process, through the memory of the target process; the target process is not explicitly involved in the communication. The `lock` and `unlock` calls ensure that the two RMA accesses do not occur concurrently. However, they do *not* ensure that the `put` by origin 1 will precede the `get` by origin 2.

Rationale. RMA does not define fine-grained mutexes in memory (only logical coarse-grained window locks). MPI provides the primitives (compare and swap, accumulate,

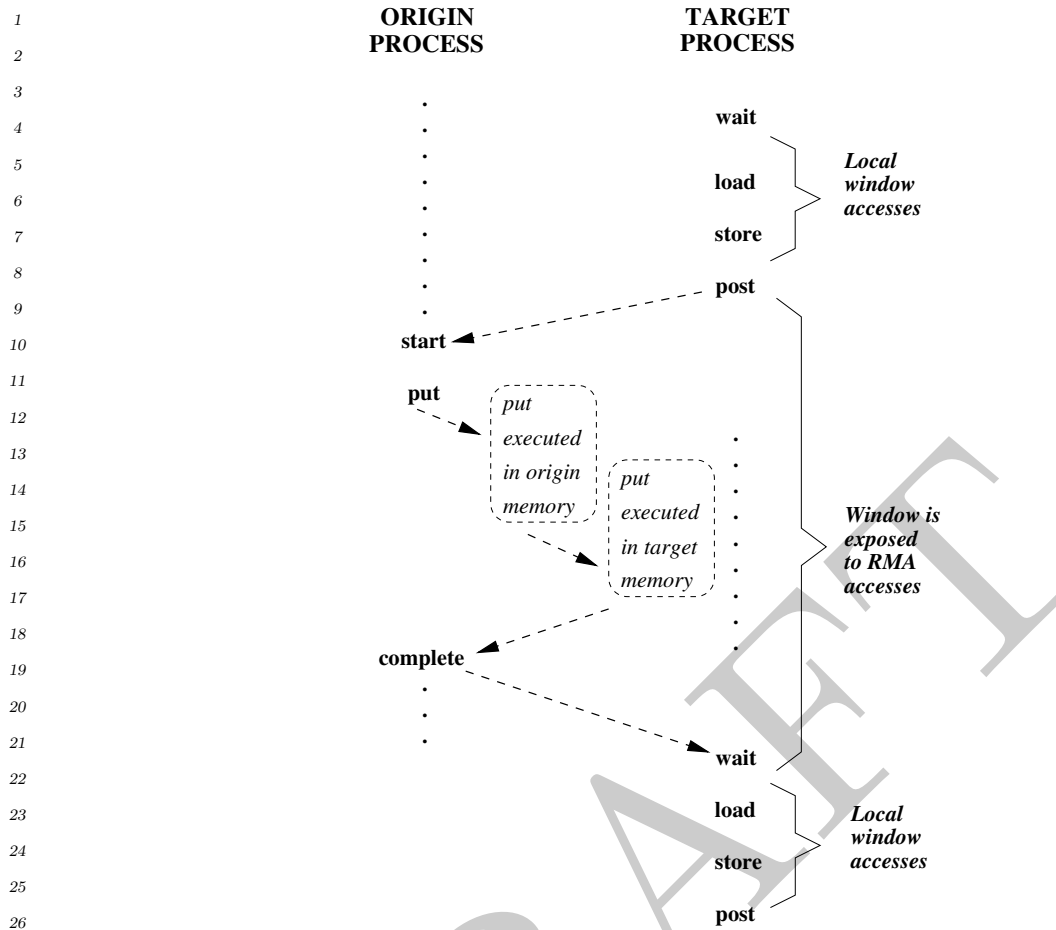


Figure 12.2: Active target communication. Dashed arrows represent synchronizations (ordering of events).

send/receive, etc.) needed to implement high-level synchronization operations. (*End of rationale.*)

12.5.1 Fence

`MPI_WIN_FENCE(assert, win)`

IN	assert	program assertion (integer)
IN	win	window object (handle)

C binding

```
int MPI_Win_fence(int assert, MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_fence(assert, win, ierror)
  INTEGER, INTENT(IN) :: assert
  TYPE(MPI_Win), INTENT(IN) :: win
```

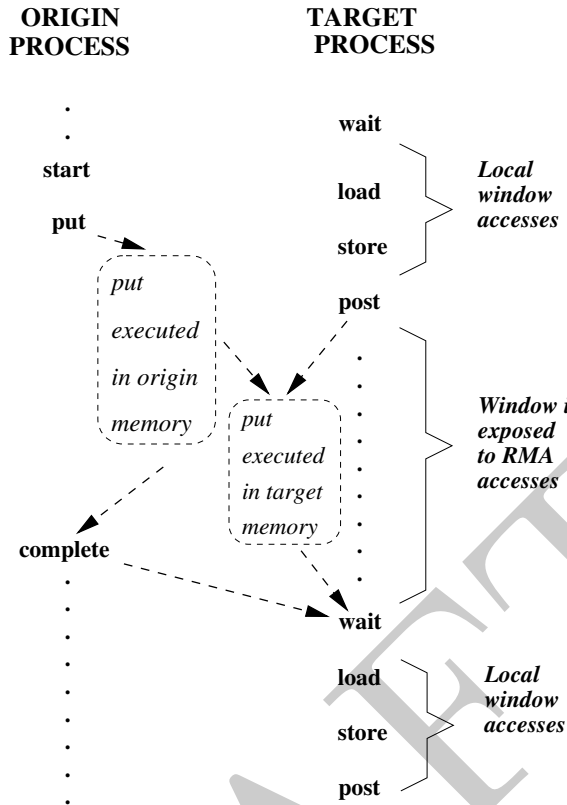



Figure 12.3: Active target communication, with weak synchronization. Dashed arrows represent synchronizations (ordering of events).

INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

MPI_WIN_FENCE(ASSERT, WIN, IERROR)
 INTEGER ASSERT, WIN, IERROR

MPI_WIN_FENCE synchronizes RMA communication operations on win. The procedure is collective over the group of win. All RMA operations on win originating at a given origin process and started before the fence call will complete at that MPI process before the fence call returns. They will be completed at their target before the fence call returns at the target. Store accesses to shared-memory of win will become visible before the fence call returns at the target. RMA operations on win started by an origin process after the fence call returns will access their target window only after MPI_WIN_FENCE has been called by the target process.

The call closes an RMA access epoch if it was preceded by another fence call and the local MPI process initiated any RMA communication operations on win between these two calls. The call closes an RMA exposure epoch if it was preceded by another fence call and the local window was the target of RMA accesses between these two calls. The call opens an RMA access epoch if it is followed by another fence call and by RMA communication calls issued between these two fence calls. The call opens an exposure epoch if it is followed by another fence call and the local window is the target of RMA accesses between these two fence calls. Thus, the fence call is equivalent to calls to a subset of post, start, complete,

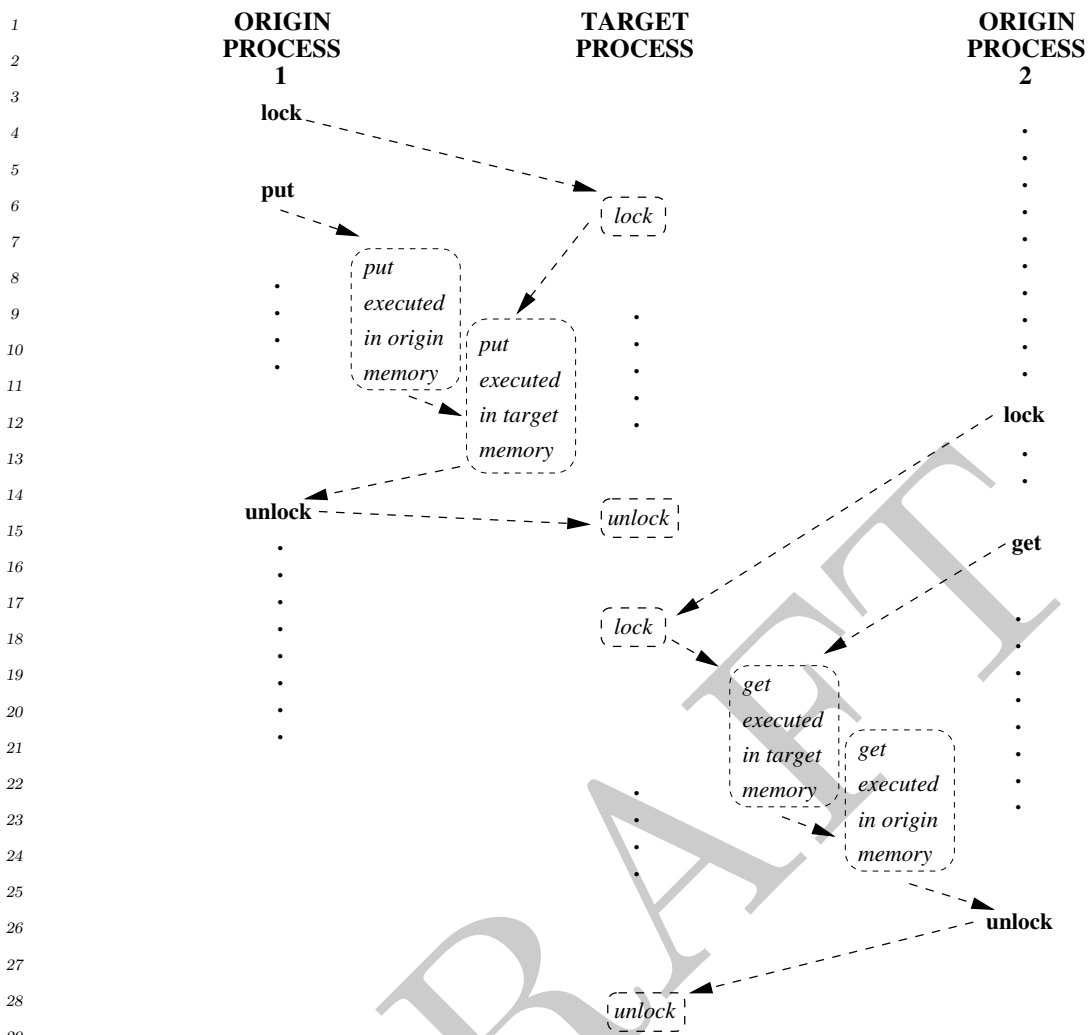


Figure 12.4: Passive target communication. Dashed arrows represent synchronizations (ordering of events).

wait.

A call to `MPI_WIN_FENCE` is usually synchronizing. However, a call to `MPI_WIN_FENCE` that is known not to close any epoch (in particular, a call with the `MPI_MODE_NOPRECEDE` assert set) is not necessarily synchronizing.

The `assert` argument is used to provide assertions on the context of the call that may be used for various optimizations. This is described in Section 12.5.5. A value of `assert = 0` is always valid.

Advice to users. Calls to `MPI_WIN_FENCE` should both precede and follow calls to RMA communication procedures that are synchronized with fence calls. (*End of advice to users.*)

12.5.2 General Active Target Synchronization

`MPI_WIN_START(group, assert, win)`

IN	group	group of target processes (handle)
IN	assert	program assertion (integer)
IN	win	window object (handle)

C binding

```
int MPI_Win_start(MPI_Group group, int assert, MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_start(group, assert, win, ierror)
```

```
TYPE(MPI_Group), INTENT(IN) :: group
INTEGER, INTENT(IN) :: assert
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_START(GROUP, ASSERT, WIN, IERROR)
```

```
INTEGER GROUP, ASSERT, WIN, IERROR
```

Opens an RMA access epoch for `win`. RMA calls issued on `win` during this epoch must access only windows at MPI processes in `group`. Each MPI process in `group` must issue a matching call to `MPI_WIN_POST`. RMA accesses to each target window will be delayed, if necessary, until the target process executed the matching call to `MPI_WIN_POST`. `MPI_WIN_START` is allowed to delay its return until the corresponding calls to `MPI_WIN_POST` have occurred, but is not required to.

The `assert` argument is used to provide assertions on the context of the call that may be used for various optimizations. This is described in Section 12.5.5. A value of `assert = 0` is always valid.

`MPI_WIN_COMPLETE(win)`

IN	win	window object (handle)
----	-----	------------------------

C binding

```
int MPI_Win_complete(MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_complete(win, ierror)
```

```
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_COMPLETE(WIN, IERROR)
```

```
INTEGER WIN, IERROR
```

Closes an RMA access epoch on `win` opened by a call to `MPI_WIN_START`. All RMA communication operations initiated on `win` during this epoch will have completed at the

origin when the call returns. All updates to shared memory in win through load/store accesses executed during this epoch will be visible at the target when the call returns.

MPI_WIN_COMPLETE enforces completion of preceding RMA operations and visibility of load/store accesses at the origin, but not at the target. A put or accumulate operation may not have completed at the target when it has completed at the origin.

Consider the sequence of calls in the example below.

Example 12.4. Use of MPI_WIN_START and MPI_WIN_COMPLETE.

```
MPI_Win_start(group, flag, win);
MPI_Put(..., win);
MPI_Win_complete(win);
```

The call to MPI_WIN_COMPLETE does not return until the put operation has completed at the origin; and the target window will be accessed by the put operation only after the call to MPI_WIN_START has matched a call to MPI_WIN_POST by the target process.

Advice to implementors. The semantics described above still leave much choice to implementors. The return from the call to MPI_WIN_START can block until the matching call to MPI_WIN_POST occurs at all target processes. One can also have implementations where the call to MPI_WIN_START returns immediately, but the call to MPI_WIN_COMPLETE delays its return until the call to MPI_WIN_POST occurred; or implementations where all three calls can complete before any target process has called MPI_WIN_POST—the data put must be buffered, in this last case, so as to allow the put to complete at the origin ahead of its completion at the target. However, once the call to MPI_WIN_POST is issued, the sequence above must complete, without further dependencies. (*End of advice to implementors.*)

Advice to users. In order to ensure a portable deadlock free program, users must assume that MPI_WIN_START may delay its return until the corresponding call to MPI_WIN_POST has occurred. (*End of advice to users.*)

MPI_WIN_POST(group, assert, win)

IN	group	group of origin processes (handle)
IN	assert	program assertion (integer)
IN	win	window object (handle)

C binding

```
int MPI_Win_post(MPI_Group group, int assert, MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_post(group, assert, win, ierror)
  TYPE(MPI_Group), INTENT(IN) :: group
  INTEGER, INTENT(IN) :: assert
  TYPE(MPI_Win), INTENT(IN) :: win
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_POST(GROUP, ASSERT, WIN, IERROR)
```

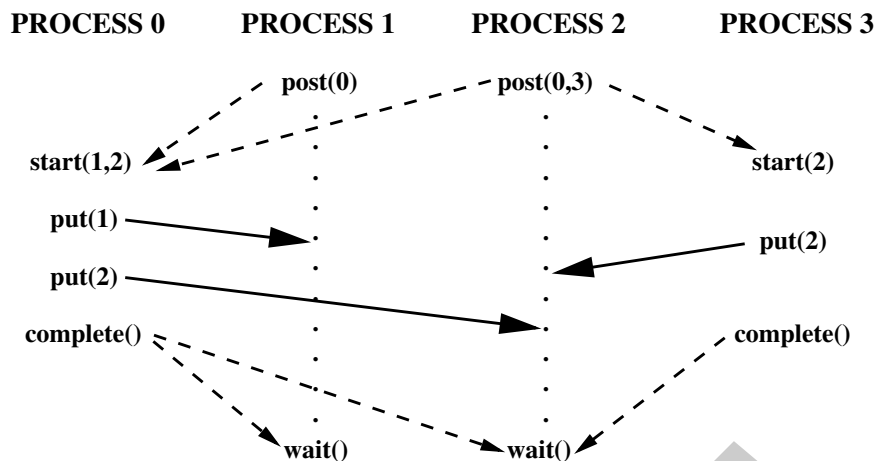


Figure 12.5: Active target communication. Dashed arrows represent synchronizations and solid arrows represent data transfer.

INTEGER GROUP, ASSERT, WIN, IERROR

Opens an RMA exposure epoch for the local window associated with win. Only MPI processes in group may access the window with RMA calls on win during this epoch. Each MPI process in group must issue a matching call to MPI_WIN_START. MPI_WIN_POST is a local procedure.

MPI_WIN_WAIT(win)

IN win window object (handle)

C binding

int MPI_Win_wait(MPI_Win win)

Fortran 2008 binding

MPI_Win_wait(win, ierror)

TYPE(MPI_Win), INTENT(IN) :: win

INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

MPI_WIN_WAIT(WIN, IERROR)

INTEGER WIN, IERROR

Closes an RMA exposure epoch opened by a call to MPI_WIN_POST on win. This call matches calls to MPI_WIN_COMPLETE on win issued by each of the origin processes that were granted access to the window during this epoch. The call to MPI_WIN_WAIT will return only after all matching calls to MPI_WIN_COMPLETE have occurred. This guarantees that all these origin processes have completed their RMA operations and shared-memory load/store accesses have become visible on the local window. When the call returns, all these RMA accesses will have completed at the target window.

Figure 12.5 illustrates the use of these four functions. Process 0 puts data in the windows of processes 1 and 2 and process 3 puts data in the window of process 2. Each start call lists the ranks of the MPI processes whose windows will be accessed; each post call

1 lists the ranks of the MPI processes that access the local window. The figure illustrates a
 2 possible timing for the events, assuming strong synchronization; in a weak synchronization,
 3 the start, put or complete calls may occur ahead of the matching post calls.
 4

5
 6 **MPI_WIN_TEST**(win, flag)

7 IN win window object (handle)
 8 OUT flag success flag (logical)
 9

10 **C binding**

11 `int MPI_Win_test(MPI_Win win, int *flag)`

12 **Fortran 2008 binding**

13 `MPI_Win_test(win, flag, ierror)`
 14 TYPE(MPI_Win), INTENT(IN) :: win
 15 LOGICAL, INTENT(OUT) :: flag
 16 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
 17

18 **Fortran binding**

19 `MPI_WIN_TEST(WIN, FLAG, IERROR)`
 20 INTEGER WIN, IERROR
 21 LOGICAL FLAG
 22

23 `MPI_WIN_TEST` is a local procedure. Repeated calls to `MPI_WIN_TEST` with the
 24 same `win` argument will eventually return `flag = true` once all accesses to the local window
 25 by the group to which it was exposed by the corresponding call to `MPI_WIN_POST` have
 26 been completed as indicated by matching `MPI_WIN_COMPLETE` calls, and `flag = false`
 27 otherwise. In the former case `MPI_WIN_WAIT` would have returned immediately. The
 28 effect of return of `MPI_WIN_TEST` with `flag = true` is the same as the effect of a return of
 29 `MPI_WIN_WAIT`. If `flag = false` is returned, then the call has no visible effect.

30 `MPI_WIN_TEST` should be called only where `MPI_WIN_WAIT` can be called. Once the
 31 call has returned `flag = true`, it must not be called again, until the window is posted again.

32 Assume that window `win` is associated with a “hidden” communicator `wincomm`, used
 33 for communication by the MPI processes in the group of `win`. The rules for matching of post
 34 and start calls and for matching complete and wait calls can be derived from the rules for
 35 matching sends and receives, by considering the following (partial) model implementation.
 36

37 **`MPI_WIN_POST(group,0,win)`** initiates a nonblocking send with tag `tag0` to each MPI
 38 process in `group`, using `wincomm`.

39 **`MPI_WIN_START(group,0,win)`** initiates a nonblocking receive with tag `tag0` from each
 40 process in `group`, using `wincomm`. An RMA access to a target process is delayed until
 41 the receive from that MPI process is completed.
 42

43 **`MPI_WIN_COMPLETE(win)`** initiates a nonblocking send with tag `tag1` to each MPI pro-
 44 cess in the group of the preceding start call.
 45

46 **`MPI_WIN_WAIT(win)`** initiates a nonblocking receive with tag `tag1` from each MPI process
 47 in the group of the preceding post call. Wait for the completion of all receives.
 48

No races can occur in a correct program: each of the sends matches a unique receive, and vice versa.

Rationale. The design for general active target synchronization requires the user to provide complete information on the communication pattern, at each end of a communication link: each origin specifies a list of targets, and each target specifies a list of origins. This provides maximum flexibility (hence, efficiency) for the implementor: each synchronization can be initiated by either side, since each “knows” the identity of the other. This also provides maximum protection from possible races. On the other hand, the design requires more information than RMA needs: in general, it is sufficient for the origin to know the rank of the target, but not vice versa. Users that want more “anonymous” communication will be required to use the fence or lock mechanisms. (*End of rationale.*)

Advice to users. Assume a communication pattern that is represented by a directed graph $G = \langle V, E \rangle$, where $V = \{0, \dots, n - 1\}$ and $ij \in E$ if origin process i accesses the window at target process j . Then each MPI process i issues a call to `MPI_WIN_POST(ingroupi, ...)`, followed by a call to `MPI_WIN_START(outgroupi, ...)`, where $outgroup_i = \{j : ij \in E\}$ and $ingroup_i = \{j : ji \in E\}$. A call is a noop, and can be skipped, if the `group` argument is empty. After the communications calls, each MPI process that issued a start will issue a complete. Finally, each MPI process that issued a post will issue a wait.

Note that each MPI process may call with a `group` argument that has different members. (*End of advice to users.*)

12.5.3 Lock

Locks are used to protect accesses to the locked target window effected by RMA calls issued between the lock and unlock calls, and to protect load/store accesses to a locked local or shared memory window executed between the lock and unlock calls. Accesses that are protected by an **exclusive lock** (acquired using `MPI_LOCK_EXCLUSIVE`) will not be concurrent at the window site with other accesses to the same window that are lock protected. Accesses that are protected by a **shared lock** (acquired using `MPI_LOCK_SHARED`) will not be concurrent at the window site with accesses protected by an *exclusive lock* to the same window.

`MPI_WIN_LOCK(lock_type, rank, assert, win)`

IN	<code>lock_type</code>	either <code>MPI_LOCK_EXCLUSIVE</code> or <code>MPI_LOCK_SHARED</code> (state)
IN	<code>rank</code>	rank of locked window (non-negative integer)
IN	<code>assert</code>	program assertion (integer)
IN	<code>win</code>	window object (handle)

C binding

`int MPI_Win_lock(int lock_type, int rank, int assert, MPI_Win win)`

Fortran 2008 binding

```

1 MPI_Win_lock(lock_type, rank, assert, win, ierror)
2   INTEGER, INTENT(IN) :: lock_type, rank, assert
3   TYPE(MPI_Win), INTENT(IN) :: win
4   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

7 MPI_WIN_LOCK(LOCK_TYPE, RANK, ASSERT, WIN, IERROR)
8   INTEGER LOCK_TYPE, RANK, ASSERT, WIN, IERROR

```

10 Opens an RMA access epoch. The window at the MPI process with a rank of rank in
 11 the group of win can be accessed by RMA operations on win during that epoch. Multiple
 12 RMA access epochs (with calls to MPI_WIN_LOCK) can occur simultaneously; however,
 13 each access epoch must target a different MPI process.

```

15 MPI_WIN_LOCK_ALL(assert, win)

```

```

17   IN      assert                program assertion (integer)
18   IN      win                   window object (handle)

```

C binding

```

21 int MPI_Win_lock_all(int assert, MPI_Win win)

```

Fortran 2008 binding

```

24 MPI_Win_lock_all(assert, win, ierror)
25   INTEGER, INTENT(IN) :: assert
26   TYPE(MPI_Win), INTENT(IN) :: win
27   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

29 MPI_WIN_LOCK_ALL(ASSERT, WIN, IERROR)
30   INTEGER ASSERT, WIN, IERROR

```

32 Opens an RMA access epoch to all MPI processes in win, with a lock type of
 33 MPI_LOCK_SHARED. During the epoch, the calling MPI process can access the window
 34 memory on all MPI processes in win by using RMA operations. A window locked with
 35 MPI_WIN_LOCK_ALL must be unlocked with MPI_WIN_UNLOCK_ALL. This routine is not
 36 collective—the ALL refers to a lock on all members of the group of the window.

37
 38 *Advice to users.* There may be additional overheads associated with using
 39 MPI_WIN_LOCK and MPI_WIN_LOCK_ALL concurrently on the same window. These
 40 overheads could be avoided by specifying the assertion MPI_MODE_NOCHECK when
 41 possible (see Section 12.5.5). (*End of advice to users.*)

MPI_WIN_UNLOCK(rank, win)			1
IN	rank	rank of window (non-negative integer)	2
IN	win	window object (handle)	3
			4
			5

C binding

```
int MPI_Win_unlock(int rank, MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_unlock(rank, win, ierror)
  INTEGER, INTENT(IN) :: rank
  TYPE(MPI_Win), INTENT(IN) :: win
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_UNLOCK(RANK, WIN, IERROR)
  INTEGER RANK, WIN, IERROR
```

Closes an RMA access epoch opened by a call to `MPI_WIN_LOCK` on window `win`. RMA operations issued during this period will have completed both at the origin and at the target when the call returns.

```
MPI_WIN_UNLOCK_ALL(win)
```

IN	win	window object (handle)
----	-----	------------------------

C binding

```
int MPI_Win_unlock_all(MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_unlock_all(win, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_UNLOCK_ALL(WIN, IERROR)
  INTEGER WIN, IERROR
```

Closes a shared RMA access epoch opened by a call to `MPI_WIN_LOCK_ALL` on window `win`. RMA operations issued during this epoch will have completed both at the origin and at the target when the call returns.

It is erroneous to have a window locked and exposed (in an exposure epoch) concurrently. For example, an MPI process may not call `MPI_WIN_LOCK` to lock a target window if the target process has called `MPI_WIN_POST` and has not yet called `MPI_WIN_WAIT`; it is erroneous to call `MPI_WIN_POST` while the local window is locked.

Rationale. An alternative is to require MPI to enforce mutual exclusion between exposure epochs and locking periods. But this would entail additional overheads when locks or active target synchronization do not interact in support of those rare interactions between the two mechanisms. The programming style that we encourage

1 here is that a set of windows is used with only one synchronization mechanism at
 2 a time, with shifts from one mechanism to another being rare and involving global
 3 synchronization. (*End of rationale.*)
 4

5 *Advice to users.* Users need to use explicit synchronization code in order to enforce
 6 mutual exclusion between locking periods and exposure epochs on a window. (*End of*
 7 *advice to users.*)
 8

9 Implementors may restrict the use of RMA communication that is synchronized by
 10 lock calls to windows in memory allocated by MPI_ALLOC_MEM (Section 9.2),
 11 MPI_WIN_ALLOCATE (Section 12.2.2), MPI_WIN_ALLOCATE_SHARED (Section 12.2.3),
 12 or attached with MPI_WIN_ATTACH (Section 12.2.4). Locks can be used portably only in
 13 such memory.

14
 15 *Rationale.* The implementation of passive target communication between processes
 16 in different *shared memory domains* may require an asynchronous software agent.
 17 Such an agent can be implemented more easily, and can achieve better performance,
 18 if restricted to specially allocated memory. It can be avoided altogether if *shared*
 19 *memory* is used. It seems natural to impose restrictions that allow the use of shared
 20 memory for RMA communication in shared memory machines.

21 (*End of rationale.*)
 22

23 Consider the sequence of calls in the example below.

24 **Example 12.5.** Use of MPI_WIN_LOCK and MPI_WIN_UNLOCK.

```
25 MPI_Win_lock(MPI_LOCK_EXCLUSIVE, rank, assert, win);
26 MPI_Put(..., rank, ..., win);
27 MPI_Win_unlock(rank, win);
28
```

29 The call to MPI_WIN_UNLOCK will not return until the put transfer has completed at the
 30 origin and at the target.
 31

32 *Advice to implementors.* The semantics described above still leave much freedom
 33 to implementors. Return from the call to MPI_WIN_LOCK may be delayed until an
 34 exclusive lock on the window is acquired; or, the first two calls may return immedi-
 35 ately, while return from MPI_WIN_UNLOCK is delayed until a lock is acquired—the
 36 update of the target window is then postponed until the call to MPI_WIN_UNLOCK
 37 occurs. However, if the call to MPI_WIN_LOCK is used to lock a window accessible
 38 via load/store accesses (i.e., a local window or a window at an MPI process for which
 39 a pointer to shared memory can be obtained via MPI_WIN_SHARED_QUERY), then
 40 the call must not return before the lock is acquired, since the lock may protect load-
 41 /store accesses to the window issued after the lock call returns. (*End of advice to*
 42 *implementors.*)
 43

44 *Advice to users.* In order to ensure a portable deadlock free program, a user must
 45 assume that MPI_WIN_LOCK may delay its return until the desired lock on the window
 46 has been acquired. (*End of advice to users.*)
 47
 48

12.5.4 Flush and Sync

All flush and sync functions can be called only within passive target epochs.

MPI_WIN_FLUSH(rank, win)

IN	rank	rank of target MPI process in the group of the window win (non-negative integer)
IN	win	window object (handle)

C binding

```
int MPI_Win_flush(int rank, MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_flush(rank, win, ierror)
  INTEGER, INTENT(IN) :: rank
  TYPE(MPI_Win), INTENT(IN) :: win
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_FLUSH(RANK, WIN, IERROR)
  INTEGER RANK, WIN, IERROR
```

All outstanding RMA operations on win initiated by the MPI process calling this procedure to the target with rank in the group of the specified window will have completed when MPI_WIN_FLUSH returns. The operations are completed both at the origin and at the target.

MPI_WIN_FLUSH_ALL(win)

IN	win	window object (handle)
----	-----	------------------------

C binding

```
int MPI_Win_flush_all(MPI_Win win)
```

Fortran 2008 binding

```
MPI_Win_flush_all(win, ierror)
  TYPE(MPI_Win), INTENT(IN) :: win
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_WIN_FLUSH_ALL(WIN, IERROR)
  INTEGER WIN, IERROR
```

All RMA operations initiated by the MPI process calling this procedure to any target on the specified window prior to this call will have completed both at the origin and at the target when MPI_WIN_FLUSH_ALL returns.

1 MPI_WIN_FLUSH_LOCAL(rank, win)

2 IN rank rank of target MPI process in the group of the
3 window win (non-negative integer)
4
5 IN win window object (handle)

6
7 **C binding**

8 int MPI_Win_flush_local(int rank, MPI_Win win)

9
10 **Fortran 2008 binding**

11 MPI_Win_flush_local(rank, win, ierror)
12 INTEGER, INTENT(IN) :: rank
13 TYPE(MPI_Win), INTENT(IN) :: win
14 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

15 **Fortran binding**

16 MPI_WIN_FLUSH_LOCAL(RANK, WIN, IERROR)
17 INTEGER RANK, WIN, IERROR

18 All outstanding RMA operations initiated on win by the MPI process calling this procedure to the target with rank in the group of the specified window will have completed at the origin when MPI_WIN_FLUSH_LOCAL returns. For example, after this procedure returns, the user may reuse any buffers provided to put, get, or accumulate operations.

23
24 MPI_WIN_FLUSH_LOCAL_ALL(win)

25 IN win window object (handle)

26
27 **C binding**

28 int MPI_Win_flush_local_all(MPI_Win win)

29
30 **Fortran 2008 binding**

31 MPI_Win_flush_local_all(win, ierror)
32 TYPE(MPI_Win), INTENT(IN) :: win
33 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

34 **Fortran binding**

35 MPI_WIN_FLUSH_LOCAL_ALL(WIN, IERROR)
36 INTEGER WIN, IERROR

37
38 All RMA operations initiated by the MPI process calling this procedure to any target on the specified window prior to this call will have completed at the origin when MPI_WIN_FLUSH_LOCAL_ALL returns.

41
42 MPI_WIN_SYNC(win)

43 IN win window object (handle)

44
45 **C binding**

46 int MPI_Win_sync(MPI_Win win)

Fortran 2008 binding

```

MPI_Win_sync(win, ierror)
    TYPE(MPI_Win), INTENT(IN) :: win
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_WIN_SYNC(WIN, IERROR)
    INTEGER WIN, IERROR

```

For windows in the separate memory model, a call to `MPI_WIN_SYNC` synchronizes the private and public window copies of `win` at the calling MPI process, as described in Section 12.7.

In the unified memory model, `MPI_WIN_SYNC` may be used to order load and store accesses to shared memory and to ensure visibility of store updates in shared memory for other threads and MPI processes.

A call to `MPI_WIN_SYNC` does not open or close an epoch and does not complete any pending RMA operations. A call to `MPI_WIN_SYNC` does not guarantee *progress* of any pending MPI operation.

12.5.5 Assertions

The `assert` argument in the calls `MPI_WIN_POST`, `MPI_WIN_START`, `MPI_WIN_FENCE`, `MPI_WIN_LOCK`, and `MPI_WIN_LOCK_ALL` is used to provide assertions on the context of the call that may be used to optimize performance. The `assert` argument does not change program semantics if it provides correct information on the program—it is erroneous to provide incorrect information. Users may always provide `assert = 0` to indicate a general case where no guarantees are made.

Advice to users. Many implementations may not take advantage of the information in `assert`; some of the information is relevant only for noncoherent shared memory machines. Users should consult their implementation’s manual to find which information is useful on each system. On the other hand, applications that provide correct assertions whenever applicable are portable and will take advantage of assertion specific optimizations whenever available. (*End of advice to users.*)

Advice to implementors. Implementations can always ignore the `assert` argument. Implementors should document which `assert` values are significant on their implementation. (*End of advice to implementors.*)

`assert` is the bit vector OR of zero or more of the following integer constants: `MPI_MODE_NOCHECK`, `MPI_MODE_NOSTORE`, `MPI_MODE_NOPUT`, `MPI_MODE_NOPRECEDE`, and `MPI_MODE_NOSUCCEED`. The significant options are listed below for each synchronization procedure.

Advice to users. C/C++ users can use bit vector OR (`|`) to combine these constants; Fortran 90 users can use the bit vector IOR intrinsic. Alternatively, Fortran users can portably use integer addition to OR the constants (each constant should appear at most once in the addition). (*End of advice to users.*)

MPI_WIN_START:

MPI_MODE_NOCHECK: the matching calls to MPI_WIN_POST have already completed on all target processes when the call to MPI_WIN_START is made. This option can be specified in a start call if and only if it is specified in each matching post call. This is similar to the optimization of “ready-send” that may save a handshake when the handshake is implicit in the code. However, ready-send is matched by a regular receive, whereas both start and post must specify the MPI_MODE_NOCHECK option.

MPI_WIN_POST:

MPI_MODE_NOCHECK: the matching calls to MPI_WIN_START have not yet occurred on any origin processes when the call to MPI_WIN_POST is made. This option can be specified by a post call if and only if it is specified by each matching start call.

MPI_MODE_NOSTORE: the local window was not updated by stores (or get or receive operations) since the last synchronization. This may avoid the need for cache synchronization during the post call.

MPI_MODE_NOPUT: the local window will not be updated by put or accumulate operations after the post call, until the ensuing (wait) synchronization. This may avoid the need for cache synchronization during the wait call.

MPI_WIN_FENCE:

MPI_MODE_NOSTORE: the local window was not updated by stores (or get or receive operations) since the last synchronization.

MPI_MODE_NOPUT: the local window will not be updated by put or accumulate operations after the fence call, until the ensuing (fence) synchronization.

MPI_MODE_NOPRECEDE: the fence does not complete any sequence of RMA operations initiated by the calling MPI process. If this assertion is given by any MPI process in the group of the window, then it must be given by all MPI processes in the group.

MPI_MODE_NOSUCCEED: the fence does not start any sequence of RMA operations initiated by the calling MPI process. If the assertion is given by any MPI process in the group of the window, then it must be given by all MPI processes in the group.

MPI_WIN_LOCK, MPI_WIN_LOCK_ALL:

MPI_MODE_NOCHECK: no other MPI process holds, or will attempt to acquire, a conflicting lock, while the calling MPI process holds the window lock. This is useful when mutual exclusion is achieved by other means, but the coherence operations that may be attached to the lock and unlock calls are still required.

Advice to users. The MPI_MODE_NOSTORE and MPI_MODE_NOPRECEDE options provide information on what happened *before* the call; the MPI_MODE_NOPUT and MPI_MODE_NOSUCCEED options provide information on what will happen *after* the call. (*End of advice to users.*)

Table 12.2: Error classes in one-sided communication routines

MPI_ERR_WIN	invalid win argument	
MPI_ERR_BASE	invalid base argument	
MPI_ERR_SIZE	invalid size argument	
MPI_ERR_DISP	invalid disp argument	
MPI_ERR_LOCKTYPE	invalid locktype argument	
MPI_ERR_ASSERT	invalid assert argument	
MPI_ERR_RMA_CONFLICT	conflicting accesses to window	
MPI_ERR_RMA_SYNC	invalid synchronization of RMA calls	
MPI_ERR_RMA_RANGE	target memory is not part of the window (in the case of a window created with MPI_WIN_CREATE_DYNAMIC, target memory is not attached)	
MPI_ERR_RMA_ATTACH	memory cannot be attached (e.g., because of resource exhaustion)	
MPI_ERR_RMA_SHARED	memory cannot be shared (e.g., some MPI process in the group of the specified communicator cannot expose <i>shared memory</i>)	
MPI_ERR_RMA_FLAVOR	passed window has the wrong flavor for the called function	

12.5.6 Miscellaneous Clarifications

Once an RMA procedure call returns, it is safe to free any opaque objects passed as arguments to that procedure. For example, the `datatype` argument of a `MPI_PUT` call can be freed as soon as the call returns, even though the communication may not be complete.

As in message-passing, datatypes must be committed before they can be used in RMA communication.

12.6 Error Handling

12.6.1 Error Handlers

Errors occurring during calls to routines that create MPI windows (e.g., `MPI_WIN_CREATE`) cause an error to be raised on the communicator provided to that procedure call. All other RMA calls have an input window argument on which errors will be raised if they occur.

The error handler `MPI_ERRORS_ARE_FATAL` is associated with the window during its creation. Users may change this default by explicitly associating a new error handler with the window (see Section 9.3).

12.6.2 Error Classes

The error classes for one-sided communication are defined in Table 12.2. RMA routines may (and almost certainly will) use other MPI error classes, such as `MPI_ERR_OP` or `MPI_ERR_RANK`.

12.7 Semantics and Correctness

The following rules specify the latest point in the execution of the application an operation must complete at the origin or the target. The update initiated by a call to `MPI_GET` in the origin process memory is visible when the get operation is complete at the origin (or earlier); the update initiated by a call to `MPI_PUT` or an accumulate procedure in the public copy of the target window is visible when the put or accumulate operation has completed at the target (or earlier). The rules also specify the latest point at which an update of one window copy becomes visible in another overlapping copy.

1. An RMA operation is completed at the origin by the ensuing call to `MPI_WIN_COMPLETE`, `MPI_WIN_FENCE`, `MPI_WIN_FLUSH`, `MPI_WIN_FLUSH_ALL`, `MPI_WIN_FLUSH_LOCAL`, `MPI_WIN_FLUSH_LOCAL_ALL`, `MPI_WIN_UNLOCK`, or `MPI_WIN_UNLOCK_ALL` that synchronizes this access at the origin.
2. If an RMA operation is completed at the origin by a call to `MPI_WIN_FENCE` then the operation is completed at the target by the matching call to `MPI_WIN_FENCE` by the target process.
3. If an RMA operation is completed at the origin by a call to `MPI_WIN_COMPLETE` then the operation is completed at the target by the matching call to `MPI_WIN_WAIT` by the target process.
4. If an RMA operation is completed at the origin by a call to `MPI_WIN_UNLOCK` or `MPI_WIN_FLUSH` (with `rank=target`), `MPI_WIN_UNLOCK_ALL`, or `MPI_WIN_FLUSH_ALL`, then the operation is completed at the target by that same call.
5. An update of a location in a private window copy in MPI process memory becomes visible in the public window copy at the latest when an ensuing call to `MPI_WIN_POST`, `MPI_WIN_FENCE`, `MPI_WIN_UNLOCK`, `MPI_WIN_UNLOCK_ALL`, or `MPI_WIN_SYNC` is executed on that window by the window owner. In the RMA unified memory model, an update of a location in a private window in MPI process memory becomes visible without additional RMA calls.
6. An update by a put or accumulate operation to a public window copy becomes visible in the private copy in MPI process memory at the latest when an ensuing call to `MPI_WIN_WAIT`, `MPI_WIN_FENCE`, `MPI_WIN_LOCK`, `MPI_WIN_LOCK_ALL`, or `MPI_WIN_SYNC` is executed on that window by the window owner. In the RMA unified memory model, an update by a put or accumulate operation to a public window copy eventually becomes visible in the private copy in MPI process memory without additional RMA calls.

The `MPI_WIN_FENCE` or `MPI_WIN_WAIT` call that completes the transfer from public copy to private copy (Rule 6) is the same call that completes the put or accumulate operation in the window copy (Rule 2, Rule 3). If a put or accumulate access was synchronized with a lock, then the update of the public window copy is complete as soon as the updating origin process executed `MPI_WIN_UNLOCK` or `MPI_WIN_UNLOCK_ALL`. In the RMA separate memory model, the update of a private copy in the target process memory may be delayed

until the target process executes a synchronization call on that window (Rule 6). Thus, updates to target process memory can always be delayed in the RMA separate memory model until the target process executes a suitable synchronization call, while they must complete in the RMA unified model without additional synchronization calls. If fence or post-start-complete-wait synchronization is used, updates to a public window copy can be delayed in both memory models until the window owner executes a synchronization call. When passive target synchronization is used, it is necessary to update the public window copy even if the window owner does not execute any related synchronization call.

The rules above also define, by implication, when an update to a public window copy becomes visible in another overlapping public window copy. Consider, for example, two overlapping windows, `win1` and `win2`. A call to `MPI_WIN_FENCE` on `win1` by the window owner makes visible in the target process memory previous updates to window `win1` by origin processes. A subsequent call to `MPI_WIN_FENCE` on `win2` makes these updates visible in the public copy of `win2`.

The behavior of some MPI RMA operations may be *undefined* in certain situations. For example, the result of several origin processes performing concurrent put operations to the same target location is undefined. In addition, the result of a single origin process performing multiple put operations to the same target location within the same access epoch is also undefined. The result at the target may have all of the data from one of the put operations (the “last” one, in some sense), some bytes from each of the operations, or something else. In MPI-2, such operations were *erroneous*. That meant that an MPI implementation was permitted to raise an error. Thus, user programs or tools that used MPI RMA could not portably permit such operations, even if the application code could function correctly with such an undefined result. Starting with MPI-3, these operations are not erroneous, but do not have a defined behavior.

Rationale. As discussed in [8], requiring operations such as overlapping puts to be erroneous makes it difficult to use MPI RMA to implement programming models—such as Unified Parallel C (UPC) or SHMEM—that permit these operations. Further, while MPI-2 defined these operations as erroneous, the MPI Forum is unaware of any implementation that enforces this rule, as it would require significant overhead. Thus, relaxing this condition does not impact existing implementations or applications. (*End of rationale.*)

Advice to implementors. Overlapping accesses are undefined. However, to assist users in debugging code, implementations may wish to provide a mode in which such operations are detected and reported to the user. Note, however, that starting with MPI-3, such operations must not raise an error. (*End of advice to implementors.*)

A program with a well-defined outcome in the `MPI_WIN_SEPARATE` memory model must obey the following rules.

- S1. A location in a window must not be accessed with load/store accesses once an update to that location has started, until the update becomes visible in the private window copy in target process memory.
- S2. A location in a window must not be accessed as a target of an RMA operation once an update to that location has started, until the update becomes visible in the public window copy. There is one exception to this rule, in the case where the same variable

1 is updated by two concurrent accumulates with the same predefined datatype, on
2 the same window. Additional restrictions on the operation apply, see the info key
3 `accumulate_ops` in Section 12.2.1.

- 4
- 5 S3. A put or accumulate must not access a target window once a store or a put or ac-
6 cumulate update to another (overlapping) window has started on a location in the
7 target window, until the update becomes visible in the public copy of the window.
8 Conversely, a store to MPI process memory to a location in a window must not start
9 once a put or accumulate update to that target window has started, until the put
10 or accumulate update becomes visible in target process memory. In both cases, the
11 restriction applies to operations even if they access disjoint locations in the window.

12 *Rationale.* The last constraint on correct RMA accesses may seem unduly restric-
13 tive, as it forbids concurrent accesses to nonoverlapping locations in a window. The
14 reason for this constraint is that, on some architectures, explicit coherence restoring
15 operations may be needed at synchronization points. A different operation may be
16 needed for locations that were updated by stores and for locations that were remotely
17 updated by put or accumulate operations. Without this constraint, the MPI library
18 would have to track precisely which locations in a window were updated by a put or
19 accumulate operation. The additional overhead of maintaining such information is
20 considered prohibitive. (*End of rationale.*)

21

22 Note that `MPI_WIN_SYNC` may be used within a passive target epoch to synchronize
23 the private and public window copies (that is, updates to one are made visible to the other).

24 In the `MPI_WIN_UNIFIED` memory model, the rules are simpler because the public and
25 private windows are the same. However, there are restrictions to avoid concurrent access
26 to the same memory locations by different MPI processes. The rules that a program with a
27 well-defined outcome must obey in this case are:

- 28
- 29 U1. A location in a window must not be accessed with load/store accesses once an update
30 to that location has started, until the update is complete, subject to the special case
31 laid out in Rule 2.
- 32
- 33 U2. Accessing a location in the window that is also the target of a remote update is valid
34 (not erroneous) but the precise result will depend on the behavior of the implemen-
35 tation. Updates from an origin process will appear in the memory of the target, but
36 there are no atomicity or ordering guarantees if more than one byte is updated. Up-
37 dates are stable in the sense that once data appears in the memory of the target, the
38 data remains until replaced by another update. This permits polling on a location for
39 a change from zero to nonzero or for a particular value, but not polling and compar-
40 ing the relative magnitude of values. Users are cautioned that polling on one memory
41 location and then accessing a different memory location has defined behavior only if
42 the other rules given here and in this chapter are followed.

43 *Advice to users.* Some compiler optimizations can result in code that maintains
44 the sequential semantics of the program, but violates this rule by introducing
45 temporary values into locations in memory. Most compilers only apply such
46 transformations under very high levels of optimization and users should be aware
47 that such aggressive optimization may produce unexpected results. (*End of*
48 *advice to users.*)

- U3. Updating a location in the window with a store access that is also the target of a remote read (but not update) is valid (not erroneous) but the precise result will depend on the behavior of the implementation. Store updates will appear in memory, but there are no atomicity or ordering guarantees if more than one byte is updated. Updates are stable in the sense that once data appears in memory, the data remains until replaced by another update. This permits updates to memory with store accesses without requiring an RMA epoch. Users are cautioned that remote accesses to a window that is updated by the local MPI process has defined behavior only if the other rules given here and elsewhere in this chapter are followed.
- U4. A location in a window must not be accessed as a target of an RMA operation once an update to that location has started and until the update completes at the target. There is one exception to this rule: in the case where the same location is updated by two concurrent accumulates with the same predefined datatype on the same window. Additional restrictions on the operation apply; see the info key `accumulate_ops` in Section 12.2.1.
- U5. A put or accumulate must not access a target window once a store, put, or accumulate update to another (overlapping) target window has started on the same location in the target window and until the update completes at the target window. Conversely, a store access to a location in a window must not be executed once a put or accumulate update to the same location in that target window has started and until the put or accumulate update completes at the target.

Advice to users. In the unified memory model, in the case where the window is in *shared memory*, `MPI_WIN_SYNC` can be used to order store accesses and make store updates to the window visible to other MPI processes and threads. Use of this routine is necessary to ensure portable behavior when point-to-point, collective, or *shared memory* synchronization is used in place of an RMA synchronization routine. `MPI_WIN_SYNC` should be called by both the reader and the writer of a shared memory variable between any non-RMA synchronization and access to that variable, as shown in Example 12.23. The calls to `MPI_WIN_SYNC` can be replaced by language level memory synchronization operations, if available. (*End of advice to users.*)

A program that violates these rules has undefined behavior.

Advice to users. A user can write correct programs by following the following rules:

fence: During each period between fence calls, each window is either updated by put or accumulate operation, or updated by stores, but not both. Locations updated by put or accumulate operations should not be accessed during the same period (with the exception of concurrent updates to the same location by accumulate operations). Locations accessed by get operations should not be updated during the same period.

post-start-complete-wait: A window should not be updated with store accesses while posted if it is being updated by put or accumulate operations. Locations updated by put or accumulate operations should not be accessed while the window is posted (with the exception of concurrent updates to the same location by accumulate operations). Locations accessed by get operations should not be updated while the window is posted.

With the post-start synchronization, the target process can tell the origin process that its window is now ready for RMA access; with the complete-wait synchronization, the origin process can tell the target process that it has finished its RMA accesses to the window.

lock: Updates to the window are protected by *exclusive locks* if they may conflict. Nonconflicting accesses (such as read-only accesses or accumulate accesses) are protected by *shared locks*, both for load/store accesses and for RMA accesses.

changing window or synchronization mode: One can change synchronization mode, or change the window used to access a location that belongs to two overlapping windows, when the MPI process memory and the window copy are guaranteed to have the same values. This is true for an MPI process after it has returned from MPI_WIN_FENCE, if RMA accesses to the window are synchronized with fences; after it has returned from MPI_WIN_WAIT, if the accesses are synchronized with post-start-complete-wait; it is true at the origin and target after the origin returned from a call to MPI_WIN_UNLOCK or MPI_WIN_UNLOCK_ALL if the accesses are synchronized with locks.

In addition, an origin process should not access the local buffer of a get operation until the operation is complete, and should not update the local buffer of a put or accumulate operation until that operation is complete.

The RMA synchronization operations define when updates are guaranteed to become visible in public and private windows. Updates may become visible earlier, but such behavior is implementation dependent. (*End of advice to users.*)

The following examples illustrate these semantics.

Example 12.6. The following example demonstrates updating a memory location inside a window for the separate memory model, according to Rule 5. The MPI_WIN_LOCK and MPI_WIN_UNLOCK calls around the store to X in process B are necessary to ensure consistency between the public and private copies of the window.

Process A	Process B
	window location X
	MPI_Win_lock (EXCLUSIVE, B)
	store X /* local update to private copy of B */
	MPI_Win_unlock (B)
	/* now visible in public window copy */
MPI_Barrier	MPI_Barrier
MPI_Win_lock (EXCLUSIVE, B)	
MPI_Get (X) /* ok, read from public window */	
MPI_Win_unlock (B)	

Example 12.7. In the RMA unified model, although the public and private copies of the windows are synchronized, caution must be used when combining load/store accesses with multi-process synchronization. Although the following example appears correct, the compiler or hardware may delay the store to X after the barrier, possibly resulting in the

MPI_GET returning an incorrect value of X.

Process A	Process B
	window location X
	store X /* update to private & public copy of B */
MPI_Barrier	MPI_Barrier
MPI_Win_lock_all	
MPI_Get(X) /* ok, read from window */	
MPI_Win_flush_local(B)	
/* read value in X */	
MPI_Win_unlock_all	

MPI_BARRIER provides process synchronization, but not memory synchronization. The example could potentially be made safe through the use of compiler- and hardware-specific notations to ensure the store to X occurs before process B enters the MPI_BARRIER. The use of one-sided synchronization calls, as shown in Example 12.6, also ensures the correct result.

Example 12.8. The following example demonstrates the reading of a memory location updated by an origin process (Rule 6) in the RMA separate memory model. Although the call to MPI_WIN_UNLOCK on process A and the MPI_BARRIER ensure that the public copy on process B reflects the updated value of X, the call to MPI_WIN_LOCK by process B is necessary to synchronize the private copy with the public copy.

Process A	Process B
	window location X
MPI_Win_lock(EXCLUSIVE, B)	
MPI_Put(X) /* update to public window */	
MPI_Win_unlock(B)	
MPI_Barrier	MPI_Barrier
	MPI_Win_lock(EXCLUSIVE, B)
	/* now visible in private copy of B */
	load X
	MPI_Win_unlock(B)

Note that in this example, the barrier is not critical to the semantic correctness. The use of *exclusive locks* guarantees no other MPI process will modify the public copy after MPI_WIN_LOCK synchronizes the private and public copies. A polling implementation looking for changes in X on process B would be semantically correct. The barrier is required to ensure that process A completes the put operation at the target before process B executes the load of X.

Example 12.9. Similar to Example 12.7, the following example is unsafe even in the unified model, because the load of X cannot be guaranteed to occur after the MPI_BARRIER. While Process B does not need to explicitly synchronize the public and private copies through

MPI_WIN_LOCK as the MPI_PUT will update both the public and private copies of the window, the scheduling of the load could result in old values of X being returned. Compiler and hardware specific notations could ensure the load occurs after the data is updated, or explicit one-sided synchronization calls can be used to ensure the proper result.

Process A	Process B
	window location X
MPI_Win_lock_all	
MPI_Put(X) /* update to window */	
MPI_Win_flush(B)	
MPI_Barrier	MPI_Barrier
	load X /* may return an obsolete value */
MPI_Win_unlock_all	

Example 12.10. The following example further clarifies Rule 5. MPI_WIN_LOCK and MPI_WIN_LOCK_ALL do *not* update the public copy of a window with changes to the private copy. Therefore, there is no guarantee that process A in the following sequence will see the value of X as updated by the store by process B before the lock.

Process A	Process B
	window location X
	store X /* update to private copy of B */
	MPI_Win_lock(SHARED, B)
MPI_Barrier	MPI_Barrier
MPI_Win_lock(SHARED, B)	
MPI_Get(X) /* X may be the X before the store */	
MPI_Win_unlock(B)	MPI_Win_unlock(B)
	/* update on X now visible in public window */

The addition of a call to MPI_WIN_SYNC before the call to MPI_BARRIER by process B would guarantee process A would see the updated value of X, as the public copy of the window would be explicitly synchronized with the private copy.

Example 12.11. Similar to the previous example, Rule 5 can have unexpected implications for general active target synchronization with the RMA separate memory model. It is *not* guaranteed that process B reads the value of X as per the local update by process A, because neither the call to MPI_WIN_WAIT nor the call to MPI_WIN_COMPLETE by process A ensure visibility in the public window copy.

Process A	Process B
window location X	
window location Y	
store Y	
MPI_Win_post(A, B) /* Y visible in public window */	
MPI_Win_start(A)	MPI_Win_start(A)

```

store X /* update to private window */
1
2
MPI_Win_complete          MPI_Win_complete
3
MPI_Win_wait
4
/* update on X may not yet be visible in the public window copy */
5
6
MPI_Barrier              MPI_Barrier
7
8
                             MPI_Win_lock(EXCLUSIVE, A)
9
                             MPI_Get(X) /* may return an obsolete value */
10
                             MPI_Get(Y)
11
                             MPI_Win_unlock(A)
12

```

To allow process B to read the value of X stored by A, the local store must be replaced by a local put operation that updates the public window copy. Note that by this replacement X may become visible in the private copy of process A only after the MPI_WIN_WAIT call in process A. The update to Y made before the MPI_WIN_POST call is visible in the public window after the MPI_WIN_POST call and therefore process B will read the proper value of Y. The get of Y could be moved to the epoch opened by MPI_WIN_START, and process B would still get the value stored by process A.

Example 12.12. The following example demonstrates the interaction of general active target synchronization with load accesses in the RMA separate memory model. Rules 5 and 6 do *not* guarantee that the private copy of X at process B has been updated before the load access is executed.

Process A	Process B
	window location X
MPI_Win_lock (EXCLUSIVE, B)	
MPI_Put (X) /* update to public window */	
MPI_Win_unlock (B)	
MPI_Barrier	MPI_Barrier
	MPI_Win_post (B)
	MPI_Win_start (B)
	load X /* access to private window */
	/* may return an obsolete value */
	MPI_Win_complete
	MPI_Win_wait

To ensure that the value put by process A is read, the load access must be replaced with a get operation, or must be placed after the call to MPI_WIN_WAIT.

12.7.1 Atomicity

The outcome of concurrent accumulate operations to the same location with the same predefined datatype is as if the accumulate operations were done at that location in some serial order. Additional restrictions on the operation apply; see the info key `accumulate_ops` in Section 12.2.1. Concurrent accumulate operations with different origin and target pairs are not ordered. Thus, there is no guarantee of atomicity beyond element-wise atomicity. The effect of this lack of atomicity is limited: The previous correctness conditions imply that a location updated by a an accumulate operation cannot be accessed by a load access or an RMA operation other than another accumulate operation until the accumulate operation has completed (at the target). Different interleavings can lead to different results only to the extent that computer arithmetics are not truly associative or commutative. The outcome of accumulate operations with overlapping types of different sizes or target displacements is undefined.

12.7.2 Ordering

Accumulate operations enable element-wise atomic read and write to window memory locations. MPI specifies ordering between accumulate operations from an origin process to the same (or overlapping) memory locations at a target process on a per-datatype granularity. The default ordering is strict ordering, which guarantees that overlapping updates from the same origin to a remote location are committed in program order and that reads (e.g., with `MPI_GET_ACCUMULATE`) and writes (e.g., with `MPI_ACCUMULATE`) are executed and committed in program order. Ordering only applies to operations originating at the same origin that access overlapping target memory regions. MPI does not provide any guarantees for accesses or updates from different origin processes to overlapping target memory regions.

The default strict ordering may incur a significant performance penalty. MPI specifies the info key `"accumulate_ordering"` to allow relaxation of the ordering semantics when specified to any window creation function. The values for this key are as follows. If set to `"none"`, then no ordering will be guaranteed for accumulate operations. This was the behavior for RMA in MPI-2 but has *not* been the default since MPI-3. The key can be set to a comma-separated list of required access orderings at the target. Allowed values in the comma-separated list are `"rar"`, `"war"`, `"raw"`, and `"waw"` for read-after-read, write-after-read, read-after-write, and write-after-write ordering, respectively. These indicate whether operations of the specified type complete in the order they were issued. For example, `"raw"` means that any writes must complete at the target before subsequent reads. These ordering requirements apply only to operations issued by the same origin process and targeting the same target process. The default value for `"accumulate_ordering"` is `"rar,raw,war,waw"`, which implies that writes complete at the target in the order in which they were issued, reads complete at the target before any writes that are issued after the reads, and writes complete at the target before any reads that are issued after the writes. Any subset of these four orderings can be specified. For example, if only read-after-read and write-after-write ordering is required, then the value of the `"accumulate_ordering"` key could be set to `"rar,waw"`. The order of values is not significant.

Note that the above ordering semantics apply only to accumulate operations, not to put and get operations. Put and get operations within an epoch are unordered.




Figure 12.6: Symmetric communication

12.7.3 Progress

One-sided communication has the same progress requirements as point-to-point communication: once a communication is enabled it is guaranteed to complete. RMA calls must have local semantics, except when required for synchronization with other RMA calls.

There is some fuzziness in the definition of the time when an RMA communication becomes enabled. This fuzziness provides to the implementor more flexibility than with point-to-point communication. Access to a target window becomes enabled once the corresponding synchronization (such as `MPI_WIN_FENCE` or `MPI_WIN_POST`) has executed. On the origin process, an RMA communication operation may become enabled as soon as the corresponding `put`, `get` or `accumulate` call has occurred, or as late as when the ensuing synchronization call is issued. Once the operation is enabled both at the origin and at the target, the operation must complete.

Consider the code fragment in Example 12.4. Some of the calls may have to delay their return until the target window has been posted. However, if the target window is posted, then the code fragment must complete. The data transfer may start as soon as the `put` call occurs, but may be delayed until the ensuing `complete` call occurs.

Consider the code fragment in Example 12.5. Some of the calls may delay their return until the lock is acquired if another MPI process holds a conflicting lock. However, if no conflicting lock is held, then the code fragment must complete.

Consider the code illustrated in Figure 12.6. Each MPI process updates the window of the other MPI process using a `put` operation, then accesses its own window. The `post` calls are local. Once the `post` calls occur, RMA access to the windows is enabled, so that each MPI process should complete the sequence of `start-put-complete`. Once these are done, the `wait` calls should complete at both MPI processes. Thus, this communication should not deadlock, irrespective of the amount of data transferred.

Assume, in the last example, that the order of the `post` and `start` calls is reversed at each MPI process. Then, the code may deadlock, as each MPI process may not return from the `start` call, waiting for the matching `post` to occur. Similarly, the program will deadlock if the order of the `complete` and `wait` calls is reversed at each MPI process.

The following two examples illustrate the fact that the synchronization between `complete` and `wait` is not symmetric: the `wait` call returns only once the `complete` occurs, but

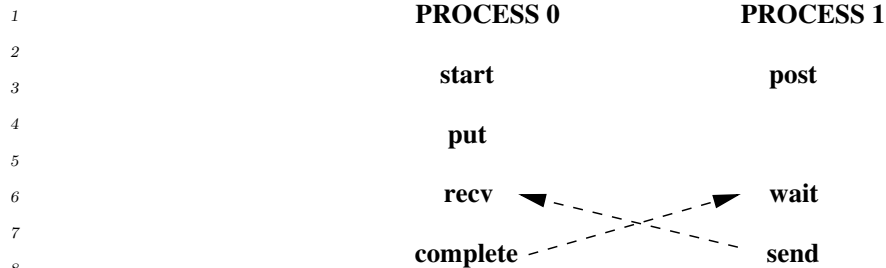


Figure 12.7: Deadlock situation

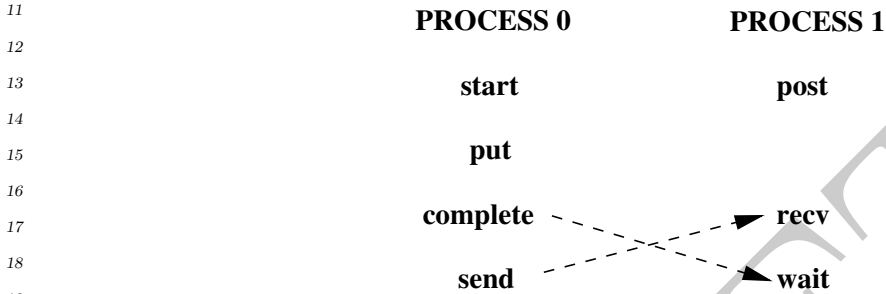


Figure 12.8: No deadlock

23 not vice versa. Consider the code illustrated in Figure 12.7. This code will deadlock: the
 24 wait of process 1 completes only once process 0 calls complete, and the receive of process
 25 0 completes once process 1 calls send. Consider, on the other hand, the code illustrated
 26 in Figure 12.8. This code will not deadlock. Once process 1 calls post, then the sequence
 27 start-put-complete on process 0 can proceed. Process 0 will reach the send call, allowing
 28 the receive call of process 1 to return.

30 *Rationale.* MPI implementations must guarantee that an MPI process makes *progress*
 31 on all enabled communications it participates in, while blocked on an MPI call. This
 32 is true for send-receive communication and applies to RMA communication as well.
 33 Thus, in the example in Figure 12.8, the put and complete calls of process 0 should
 34 complete while process 1 is waiting for the receive operation to complete. This may
 35 require the involvement of process 1, e.g., to transfer the data.

36 A similar issue is whether such progress must occur while an MPI process is busy
 37 computing, or blocked in a non-MPI call. Suppose that in the last example the send-
 38 receive pair is replaced by a write-to-socket/read-from-socket pair. Then MPI does
 39 not specify whether deadlock is avoided. Suppose that the blocking receive of process
 40 1 is replaced by a very long compute loop. Then, according to one interpretation
 41 of the MPI standard, process 0 must return from the complete call after a bounded
 42 delay, even if process 1 does not reach any MPI call in this period of time. According
 43 to another interpretation, the complete call may block until process 1 reaches the
 44 wait call, or reaches another MPI call. The qualitative behavior is the same, under
 45 both interpretations, unless an MPI process is caught in an infinite compute loop, in
 46 which case the difference may not matter. However, the quantitative expectations
 47 are different. Different MPI implementations reflect these different interpretations.
 48 While this ambiguity is unfortunate, the MPI Forum decided not to define which

interpretation of the standard is the correct one, since the issue is contentious. See also Section 2.9 on *progress*. (*End of rationale*.)

The use of shared memory loads and/or stores for synchronizing purposes between MPI processes does not guarantee progress, and therefore a *deadlock* may occur if an MPI implementation does not provide *strong progress*, as shown in Example 12.13.

Example 12.13. Possible *deadlock* due to the use of a shared memory variable for synchronization.

`comm_sm` shall be a shared memory communicator (e.g., returned from a call to `MPI_COMM_SPLIT_TYPE` with `split_type=MPI_COMM_TYPE_SHARED`) with at least two MPI processes. `win_sm` is a shared memory window with the `AckInRank0` as window portion in MPI process with rank 0. The ranks in `comm_sm` and `win_sm` should be the same. According to Section 12.7 rules U2 and U3, a volatile store to `AckInRank0` will be visible in the other MPI process without further RMA calls.

```
int volatile_load(int *addr) {return *(volatile int *)addr;}
void volatile_store(int *addr, int val) {*((volatile int *)addr) = val;}

Process with rank 0                               Process with rank 1
-----
MPI_Win_shared_query(win,                          MPI_Win_shared_query(win,
/*rank=*/ 0, ..., AckInRank0);                    /*rank=*/ 0, ..., AckInRank0);

volatile_store(AckInRank0, 0);
MPI_Win_fence(win_sm)
MPI_Buffer_attach(myHugeBuffer, ...);
MPI_Bsend(myHugeMessage, ...,
/*rank=*/ 1, ..., comm_sm);
sleep(10); // to guarantee that
// the while-loop starts
// after rank 1 is
// blocked in MPI_Recv
MPI_Recv(&myHugeMessage, ...
/*rank=*/ 0, ..., comm_sm, ...);
volatile_store(AckInRank0, 222);

while(volatile_load(AckInRank0)!=222)
/*empty polling loop*/;
MPI_Buffer_detach(&pTemp, &size);
// deadlock
```

While the call to `MPI_Recv` in the MPI process with rank 1 delays its return (until an unspecified MPI procedure call in the MPI process with rank 0 happens to send the buffered data), the subsequent statement cannot change the value of the shared window buffer `AckInRank0`. As long as this value is not changed, the while loop in the MPI process with rank 0 will continue and therefore the next MPI procedure call (`MPI_Buffer_detach`) cannot happen, which is then a *deadlock*.

Note that both communication patterns (A) `BSEND-RECV-DETACH` and (B) the shared memory store/load for synchronization purpose, can be in different software layers and each layer would work properly, but the combination of (A) and (B) can cause the *deadlock*.

12.7.4 Registers and Compiler Optimizations

Advice to users. All the material in this section is an advice to users. (*End of advice to users.*)

A coherence problem exists between variables kept in registers and the memory values of these variables. An RMA call may access a variable in memory (or cache), while the up-to-date value of this variable is in register. A get will not return the latest variable value, and a put may be overwritten when the register is stored back in memory. Note that these issues are unrelated to the RMA memory model; that is, these issues apply even if the memory model is MPI_WIN_UNIFIED.

The problem is illustrated in Example 12.14.

Example 12.14.

In this example, variable `buff` is allocated in the register `reg_A` and therefore `ccc` will have the old value of `buff` and not the new value 777.

Source of Process 1	Source of Process 2	Executed in Process 2
<code>bbbb = 777</code>	<code>buff = 999</code>	<code>reg_A:=999</code>
<code>call MPI_WIN_FENCE</code>	<code>call MPI_WIN_FENCE</code>	
<code>call MPI_PUT(bbbb</code> <code>into buff of process 2)</code>		<code>stop appl.thread</code> <code>buff:=777 in PUT handler</code> <code>continue appl.thread</code>
<code>call MPI_WIN_FENCE</code>	<code>call MPI_WIN_FENCE</code>	
	<code>ccc = buff</code>	<code>ccc:=reg_A</code>

This problem, which also afflicts in some cases send/receive communication, is discussed more at length in Section 19.1.16.

Programs written in C avoid this problem, because of the semantics of C. Many Fortran compilers will avoid this problem, without disabling compiler optimizations. However, in order to avoid register coherence problems in a completely portable manner, users should restrict their use of RMA windows to variables stored in modules or COMMON blocks. To prevent problems with the argument copying and register optimization done by Fortran compilers, please note the hints in Sections 19.1.10–19.1.20. Sections 19.1.17 to 19.1.17 discuss several solutions for the problem in this example.

12.8 Examples

Example 12.15. The following example shows a generic loosely synchronous, iterative code, using MPI_FENCE for synchronization. The window at each MPI process consists of array A, which contains the origin and target buffers of the put operations.

```

...
while (!converged(A)) {
  update(A);
  MPI_Win_fence(MPI_MODE_NOPRECEDE, win);
  for(i=0; i < toneighbors; i++)
    MPI_Put(&frombuf[i], 1, fromtype[i], toneighbor[i],
           todisp[i], 1, totype[i], win);
  MPI_Win_fence((MPI_MODE_NOSTORE | MPI_MODE_NOSUCCEED), win);

```

```
}

```

The same code could be written with get rather than put. Note that, during the communication phase, each window is concurrently read (as origin buffer of puts) and written (as target buffer of puts). This is OK, provided that there is no overlap between the target buffer of a put and another communication buffer.

Example 12.16. Same generic example, with more computation/communication overlap. We assume that the update phase is broken into two subphases: the first, where the “boundary,” which is involved in communication, is updated, and the second, where the “core,” which neither uses nor provides communicated data, is updated.

```
...
while (!converged(A)) {
    update_boundary(A);
    MPI_Win_fence((MPI_MODE_NOPUT | MPI_MODE_NOPRECEDE), win);
    for(i=0; i < fromneighbors; i++)
        MPI_Get(&tobuf[i], 1, totype[i], fromneighbor[i],
                fromdisp[i], 1, fromtype[i], win);
    update_core(A);
    MPI_Win_fence(MPI_MODE_NOSUCCEED, win);
}

```

The get communication can be concurrent with the core update, since they do not access the same locations, and the local update of the origin buffer by the get operation can be concurrent with the local update of the core by the `update_core` call. In order to get similar overlap with put communication we would need to use separate windows for the core and for the boundary. This is required because we do not allow local stores to be concurrent with puts on the same, or on overlapping, windows.

Example 12.17. Same code as in Example 12.15, rewritten using post-start-complete-wait.

```
...
while (!converged(A)) {
    update(A);
    MPI_Win_post(fromgroup, 0, win);
    MPI_Win_start(togroup, 0, win);
    for(i=0; i < toneighbors; i++)
        MPI_Put(&frombuf[i], 1, fromtype[i], toneighbor[i],
                todisp[i], 1, totype[i], win);
    MPI_Win_complete(win);
    MPI_Win_wait(win);
}

```

Example 12.18. Same example, with post-start-complete-wait, as in Example 12.16.

```
...
while (!converged(A)) {
    update_boundary(A);

```

```

1  MPI_Win_post(togroup, MPI_MODE_NOPUT, win);
2  MPI_Win_start(fromgroup, 0, win);
3  for(i=0; i < fromneighbors; i++)
4      MPI_Get(&tobuf[i], 1, totype[i], fromneighbor[i],
5             fromdisp[i], 1, fromtype[i], win);
6  update_core(A);
7  MPI_Win_complete(win);
8  MPI_Win_wait(win);
9  }

```

Example 12.19. A checkerboard, or double buffer communication pattern, that allows more computation/communication overlap. Array A_0 is updated using values of array A_1 , and vice versa. We assume that communication is symmetric: if process A gets data from process B, then process B gets data from process A. Window $wini$ consists of array A_i .

```

16  ...
17  if (!converged(A0,A1))
18      MPI_Win_post(neighbors, (MPI_MODE_NOCHECK | MPI_MODE_NOPUT), win0);
19  MPI_Barrier(comm0);
20  /* the barrier is needed because the start call inside the
21  loop uses the nocheck option */
22  while (!converged(A0, A1)) {
23      /* communication on A0 and computation on A1 */
24      update2(A1, A0); /* local update of A1 that depends on A0 (and A1) */
25      MPI_Win_start(neighbors, MPI_MODE_NOCHECK, win0);
26      for(i=0; i < fromneighbors; i++)
27          MPI_Get(&tobuf0[i], 1, totype0[i], neighbor[i],
28                 fromdisp0[i], 1, fromtype0[i], win0);
29      update1(A1); /* local update of A1 that is
30                  concurrent with communication that updates A0 */
31      MPI_Win_post(neighbors, (MPI_MODE_NOCHECK | MPI_MODE_NOPUT), win1);
32      MPI_Win_complete(win0);
33      MPI_Win_wait(win0);
34
35      /* communication on A1 and computation on A0 */
36      update2(A0, A1); /* local update of A0 that depends on A1 (and A0) */
37      MPI_Win_start(neighbors, MPI_MODE_NOCHECK, win1);
38      for(i=0; i < fromneighbors; i++)
39          MPI_Get(&tobuf1[i], 1, totype1[i], neighbor[i],
40                 fromdisp1[i], 1, fromtype1[i], win1);
41      update1(A0); /* local update of A0 that depends on A0 only,
42                  concurrent with communication that updates A1 */
43      if (!converged(A0,A1))
44          MPI_Win_post(neighbors, (MPI_MODE_NOCHECK | MPI_MODE_NOPUT), win0);
45      MPI_Win_complete(win1);
46      MPI_Win_wait(win1);
47  }

```

An MPI process posts the local window associated with w_0 before it completes RMA accesses to the remote windows associated with w_1 . When the call to `MPI_WIN_WAIT` on w_1 returns, then all neighbors of the calling MPI process have posted the windows associated with w_0 . Conversely, when the call to `MPI_WIN_WAIT` on w_0 returns, then

all neighbors of the calling MPI process have posted the windows associated with win1. Therefore, the MPI_MODE_NOCHECK option can be used with the calls to MPI_WIN_START. Put operations can be used, instead of get operations, if the area of array A0 (resp. A1) used by update(A1, A0) (resp. update(A0, A1)) is disjoint from the area modified by the RMA operation. On some systems, a put operation may be more efficient than a get operation, as it requires information exchange only in one direction.

In the next several examples, for conciseness, the expression

```
z = MPI_Get_accumulate(...)
```

means to perform a get-accumulate operation with the result buffer (given by result_addr in the description of MPI_GET_ACCUMULATE) on the left side of the assignment, in this case, z. This format is also used with MPI_COMPARE_AND_SWAP and MPI_COMM_SIZE. Process B... refers to any process other than A.

Example 12.20. The following example implements a naive, nonscalable counting semaphore. The example demonstrates the use of MPI_WIN_SYNC to manipulate the public copy of X, as well as MPI_WIN_FLUSH to complete operations without closing the access epoch opened with MPI_WIN_LOCK_ALL. To avoid the rules regarding synchronization of the public and private copies of windows, MPI_ACCUMULATE and MPI_GET_ACCUMULATE are used to write to or read from the local public copy.

Process A	Process B...
MPI_Win_lock_all	MPI_Win_lock_all
window location X	
X=MPI_Comm_size()	
MPI_Win_sync	
MPI_Barrier	MPI_Barrier
MPI_Accumulate(X, MPI_SUM, -1)	MPI_Accumulate(X, MPI_SUM, -1)
stack variable z	stack variable z
do	do
z = MPI_Get_accumulate(X,	z = MPI_Get_accumulate(X,
MPI_NO_OP, 0)	MPI_NO_OP, 0)
MPI_Win_flush(A)	MPI_Win_flush(A)
while(z!=0)	while(z!=0)
MPI_Win_unlock_all	MPI_Win_unlock_all

Example 12.21. Implementing a critical region between two MPI processes (Peterson's algorithm). Despite their appearance in the following example, MPI_WIN_LOCK_ALL and MPI_WIN_UNLOCK_ALL are not collective calls, but it is frequently useful to open shared access epochs to all MPI processes from all other MPI processes in a window. Once the access epochs are opened, accumulate operations as well as flush and sync synchronization can be used to read from or write to the public copy of the window.

Process A	Process B
window location X	window location Y
window location T	

```

1  MPI_Win_lock_all
2  X=1
3  MPI_Win_sync
4  MPI_Barrier
5  MPI_Accumulate(T, MPI_REPLACE, 1)
6  stack variables t,y
7  t=1
8  y=MPI_Get_accumulate(Y,
9  MPI_NO_OP, 0)
10 while(y==1 && t==1) do
11   y=MPI_Get_accumulate(Y,
12   MPI_NO_OP, 0)
13   t=MPI_Get_accumulate(T,
14   MPI_NO_OP, 0)
15   MPI_Win_flush_all
16 done
17 // critical region
18 MPI_Accumulate(X, MPI_REPLACE, 0)
19 MPI_Win_unlock_all

```

```

1  MPI_Win_lock_all
2  Y=1
3  MPI_Win_sync
4  MPI_Barrier
5  MPI_Accumulate(T, MPI_REPLACE, 0)
6  stack variable t,x
7  t=0
8  x=MPI_Get_accumulate(X,
9  MPI_NO_OP, 0)
10 while(x==1 && t==0) do
11   x=MPI_Get_accumulate(X,
12   MPI_NO_OP, 0)
13   t=MPI_Get_accumulate(T,
14   MPI_NO_OP, 0)
15   MPI_Win_flush(A)
16 done
17 // critical region
18 MPI_Accumulate(Y, MPI_REPLACE, 0)
19 MPI_Win_unlock_all

```

Example 12.22. Implementing a critical region between multiple MPI processes with compare and swap. The call to `MPI_WIN_SYNC` is necessary on Process A after local initialization of A to guarantee the public copy has been updated with the initialization value found in the private copy. It would also be valid to call `MPI_ACCUMULATE` with `MPI_REPLACE` to directly initialize the public copy. A call to `MPI_WIN_FLUSH` would be necessary to assure A in the public copy of Process A had been updated before the barrier.

Process A	Process B...
<pre> 27 MPI_Win_lock_all 28 atomic location A 29 A=0 30 MPI_Win_sync 31 MPI_Barrier 32 stack variable r = 1 33 while(r != 0) do 34 r = MPI_Compare_and_swap(A, 0, 1) 35 MPI_Win_flush(A) 36 done 37 // critical region 38 r = MPI_Compare_and_swap(A, 1, 0) 39 MPI_Win_unlock_all </pre>	<pre> 27 MPI_Win_lock_all 30 MPI_Barrier 31 stack variable r = 1 32 while(r != 0) do 33 r = MPI_Compare_and_swap(A, 0, 1) 34 MPI_Win_flush(A) 35 done 36 // critical region 37 r = MPI_Compare_and_swap(A, 1, 0) 38 MPI_Win_unlock_all </pre>

Example 12.23. The following example demonstrates the proper synchronization in the unified memory model when a data transfer is implemented with load and store accesses in the case of windows in *shared memory* (instead of using `MPI_PUT` or `MPI_GET`) and the synchronization between MPI processes is performed using point-to-point communication. The synchronization between MPI processes must be supplemented with a memory synchronization through calls to `MPI_WIN_SYNC`, which act locally as a processor-memory barrier. In Fortran, if `MPI_ASYNC_PROTECTS_NONBLOCKING` is `.FALSE.` or the variable X is not declared as `ASYNCHRONOUS`, reordering of the accesses to the variable X must be prevented

with `MPI_F_SYNC_REG` operations. (No equivalent function is needed in C.) The variable `X` is contained within a *shared memory window* and `X` corresponds to the same memory location at both processes. The first call to `MPI_WIN_SYNC` performed by process A ensures completion of the load/store accesses issued by process A. The first call to `MPI_WIN_SYNC` performed by process B ensures that process A's updates to `X` are visible to process B. Similarly, the second call to `MPI_WIN_SYNC` on each process ensures correct ordering of the point-to-point communication and thus that the load/store operations on process B have completed before any subsequent load/store accesses to the variable `X` in process A.

Process A	Process B
<code>MPI_WIN_LOCK_ALL (</code>	<code>MPI_WIN_LOCK_ALL (</code>
<code>MPI_MODE_NOCHECK , win)</code>	<code>MPI_MODE_NOCHECK , win)</code>
<code>DO ...</code>	<code>DO ...</code>
<code>X=...</code>	
<code>MPI_F_SYNC_REG(X)</code>	
<code>MPI_WIN_SYNC(win)</code>	
<code>MPI_SEND</code>	
 	<code>MPI_RECV</code>
	<code>MPI_WIN_SYNC(win)</code>
	<code>MPI_F_SYNC_REG(X)</code>
	<code>print X</code>
	<code>MPI_F_SYNC_REG(X)</code>
	<code>MPI_WIN_SYNC(win)</code>
	<code>MPI_SEND</code>
<code>MPI_RECV</code>	
<code>MPI_WIN_SYNC(win)</code>	
<code>MPI_F_SYNC_REG(X)</code>	
<code>END DO</code>	<code>END DO</code>
<code>MPI_WIN_UNLOCK_ALL(win)</code>	<code>MPI_WIN_UNLOCK_ALL(win)</code>

Example 12.24. The following example shows how request-based operations can be used to overlap communication with computation. Each MPI process fetches, processes, and writes the result for `NSTEPS` chunks of data. Instead of a single buffer, `M` local buffers are used to allow up to `M` communication operations to overlap with computation.

```

int          i, j;
MPI_Win     win;
MPI_Request  put_req[M] = { MPI_REQUEST_NULL };
MPI_Request  get_req;
double      *baseptr;
double      data[M][N];

MPI_Win_allocate(NSTEPS*N*sizeof(double), sizeof(double), MPI_INFO_NULL,
                MPI_COMM_WORLD, &baseptr, &win);

MPI_Win_lock_all(0, win);

```

```

1  for (i = 0; i < NSTEPS; i++) {
2    if (i < M)
3      j = i;
4    else
5      MPI_Waitany(M, put_req, &j, MPI_STATUS_IGNORE);
6
7    MPI_Rget(data[j], N, MPI_DOUBLE, target, i*N, N, MPI_DOUBLE, win,
8             &get_req);
9    MPI_Wait(&get_req, MPI_STATUS_IGNORE);
10   compute(i, data[j], ...);
11   MPI_Rput(data[j], N, MPI_DOUBLE, target, i*N, N, MPI_DOUBLE, win,
12            &put_req[j]);
13 }
14 MPI_Waitall(M, put_req, MPI_STATUSES_IGNORE);
15 MPI_Win_unlock_all(win);

```

Example 12.25. The following example constructs a distributed shared linked list using dynamic windows. Initially process 0 creates the head of the list, attaches it to the window, and broadcasts the pointer to all MPI processes. All MPI processes then concurrently append N new elements to the list. When an MPI process attempts to attach its element to the tail of the list it may discover that its tail pointer is stale and it must chase ahead to the new tail before the element can be attached. This example requires some modification to work in an environment where the layout of the structures is different on different MPI processes.

```

25  ...
26  #define NUM_ELEMS 10
27
28  #define LLIST_ELEM_NEXT_RANK ( offsetof(llist_elem_t, next) + \
29                               offsetof(llist_ptr_t, rank) )
30  #define LLIST_ELEM_NEXT_DISP ( offsetof(llist_elem_t, next) + \
31                                 offsetof(llist_ptr_t, disp) )
32
33  /* Linked list pointer */
34  typedef struct {
35    MPI_Aint disp;
36    int rank;
37  } llist_ptr_t;
38
39  /* Linked list element */
40  typedef struct {
41    llist_ptr_t next;
42    int value;
43  } llist_elem_t;
44
45  const llist_ptr_t nil = { (MPI_Aint) MPI_BOTTOM, -1 };
46
47  /* List of locally allocated list elements. */
48  static llist_elem_t **my_elems = NULL;
49  static int my_elems_size = 0;
50  static int my_elems_count = 0;
51
52  /* Allocate a new shared linked list element */

```

```

1 MPI_Aint alloc_elem(int value, MPI_Win win) {
2     MPI_Aint disp;
3     llist_elem_t *elem_ptr;
4
5     /* Allocate the new element and register it with the window */
6     MPI_Alloc_mem(sizeof(llist_elem_t), MPI_INFO_NULL, &elem_ptr);
7     elem_ptr->value = value;
8     elem_ptr->next = nil;
9     MPI_Win_attach(win, elem_ptr, sizeof(llist_elem_t));
10
11     /* Add the element to the list of local elements so we can free
12        it later. */
13     if (my_elems_size == my_elems_count) {
14         my_elems_size += 100;
15         my_elems = realloc(my_elems, my_elems_size*sizeof(void*));
16     }
17     my_elems[my_elems_count] = elem_ptr;
18     my_elems_count++;
19
20     MPI_Get_address(elem_ptr, &disp);
21     return disp;
22 }
23
24 int main(int argc, char *argv[]) {
25     int          procid, nproc, i;
26     MPI_Win      llist_win;
27     llist_ptr_t  head_ptr, tail_ptr;
28
29     MPI_Init(&argc, &argv);
30
31     MPI_Comm_rank(MPI_COMM_WORLD, &procid);
32     MPI_Comm_size(MPI_COMM_WORLD, &nproc);
33
34     MPI_Win_create_dynamic(MPI_INFO_NULL, MPI_COMM_WORLD, &llist_win);
35
36     /* Process 0 creates the head node */
37     if (procid == 0)
38         head_ptr.disp = alloc_elem(-1, llist_win);
39
40     /* Broadcast the head pointer to everyone */
41     head_ptr.rank = 0;
42     MPI_Bcast(&head_ptr.disp, 1, MPI_AINT, 0, MPI_COMM_WORLD);
43     tail_ptr = head_ptr;
44
45     /* Lock the window for shared access to all targets */
46     MPI_Win_lock_all(0, llist_win);
47
48     /* All processes concurrently append NUM_ELEMS elements to the list */
49     for (i = 0; i < NUM_ELEMS; i++) {
50         llist_ptr_t new_elem_ptr;
51         int success;
52
53         /* Create a new list element and attach it to the window */
54         new_elem_ptr.rank = procid;
55         new_elem_ptr.disp = alloc_elem(procid, llist_win);
56
57         /* Append the new node to the list. This might take multiple

```

```

1      attempts if others have already appended and our tail pointer
2      is stale. */
3      do {
4          llist_ptr_t next_tail_ptr = nil;
5
6          MPI_Compare_and_swap((void*) &new_elem_ptr.rank, (void*) &nil.rank,
7                               (void*)&next_tail_ptr.rank, MPI_INT, tail_ptr.rank,
8                               MPI_Aint_add(tail_ptr.disp, LLIST_ELEM_NEXT_RANK),
9                               llist_win);
10
11         MPI_Win_flush(tail_ptr.rank, llist_win);
12         success = (next_tail_ptr.rank == nil.rank);
13
14         if (success) {
15             MPI_Accumulate(&new_elem_ptr.disp, 1, MPI_AINT, tail_ptr.rank,
16                           MPI_Aint_add(tail_ptr.disp, LLIST_ELEM_NEXT_DISP), 1,
17                           MPI_AINT, MPI_REPLACE, llist_win);
18
19             MPI_Win_flush(tail_ptr.rank, llist_win);
20             tail_ptr = new_elem_ptr;
21
22         } else {
23             /* Tail pointer is stale, fetch the displacement. May take
24              * multiple tries if it is being updated. */
25             do {
26                 MPI_Get_accumulate(NULL, 0, MPI_AINT, &next_tail_ptr.disp,
27                                     1, MPI_AINT, tail_ptr.rank,
28                                     MPI_Aint_add(tail_ptr.disp, LLIST_ELEM_NEXT_DISP),
29                                     1, MPI_AINT, MPI_NO_OP, llist_win);
30
31                 MPI_Win_flush(tail_ptr.rank, llist_win);
32             } while (next_tail_ptr.disp == nil.disp);
33             tail_ptr = next_tail_ptr;
34         }
35     } while (!success);
36 }
37
38 MPI_Win_unlock_all(llist_win);
39 MPI_Barrier(MPI_COMM_WORLD);
40
41 /* Free all the elements in the list */
42 for ( ; my_elems_count > 0; my_elems_count--) {
43     MPI_Win_detach(llist_win, my_elems[my_elems_count-1]);
44     MPI_Free_mem(my_elems[my_elems_count-1]);
45 }
46 MPI_Win_free(&llist_win);
47 ...
48

```

Chapter 13

External Interfaces

13.1 Introduction

This chapter contains calls used to create **generalized requests**, which allow users to create new nonblocking operations with an interface similar to what is present in MPI. These calls can be used to layer new functionality on top of MPI. Section 13.3 deals with setting the information found in `status`. This functionality is needed for generalized requests.

13.2 Generalized Requests

The goal of generalized requests is to allow users to define new nonblocking operations. Such an outstanding nonblocking operation is represented by a (generalized) request. A fundamental property of nonblocking operations is that *progress* toward the completion of this operation occurs asynchronously, i.e., concurrently with normal program execution. Typically, this requires execution of code concurrently with the execution of the user code, e.g., in a separate thread or in a signal handler. Operating systems provide a variety of mechanisms in support of concurrent execution. MPI does not attempt to standardize or to replace these mechanisms: it is assumed programmers who wish to define new asynchronous operations will use the mechanisms provided by the underlying operating system. Thus, the calls in this section only provide a means for defining the effect of MPI calls such as `MPI_WAIT` or `MPI_CANCEL` when they apply to generalized requests, and for signaling to MPI the completion of a generalized operation.

Rationale. It is tempting to also define an MPI standard mechanism for achieving concurrent execution of user-defined nonblocking operations. However, it is difficult to define such a mechanism without consideration of the specific mechanisms used in the operating system. The Forum feels that concurrency mechanisms are a proper part of the underlying operating system and should not be standardized by MPI; the MPI standard should only deal with the interaction of such mechanisms with MPI. (*End of rationale.*)

For a regular request, the operation associated with the request is performed by the MPI implementation, and the operation completes without intervention by the application. For a generalized request, the operation associated with the request is performed by the application; therefore, the application must notify MPI through a call to `MPI_GREQUEST_COMPLETE` when the operation completes. MPI maintains the “completion” status of generalized requests. Any other request state has to be maintained by the user.

A new generalized request is started with

```

1 MPI_GREQUEST_START(query_fn, free_fn, cancel_fn, extra_state, request)
2     IN      query_fn      callback function invoked when request status is
3                          queried (function)
4
5     IN      free_fn       callback function invoked when request is freed
6                          (function)
7
8     IN      cancel_fn     callback function invoked when request is cancelled
9                          (function)
10
11    IN      extra_state    extra state
12
13    OUT     request        generalized request (handle)

```

C binding

```

14 int MPI_Grequest_start(MPI_Grequest_query_function *query_fn,
15                       MPI_Grequest_free_function *free_fn,
16                       MPI_Grequest_cancel_function *cancel_fn, void *extra_state,
17                       MPI_Request *request)

```

Fortran 2008 binding

```

18 MPI_Grequest_start(query_fn, free_fn, cancel_fn, extra_state, request, ierror)
19
20 PROCEDURE(MPI_Grequest_query_function) :: query_fn
21 PROCEDURE(MPI_Grequest_free_function) :: free_fn
22 PROCEDURE(MPI_Grequest_cancel_function) :: cancel_fn
23 INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
24 TYPE(MPI_Request), INTENT(OUT) :: request
25 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

26 MPI_GREQUEST_START(QUERY_FN, FREE_FN, CANCEL_FN, EXTRA_STATE, REQUEST, IERROR)
27
28 EXTERNAL QUERY_FN, FREE_FN, CANCEL_FN
29 INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
30 INTEGER REQUEST, IERROR

```

Advice to users. Note that a generalized request is of the same type as regular requests, in C and Fortran. (*End of advice to users.*)

The call starts a generalized request and returns a handle to it in `request`.

The syntax and meaning of the callback functions are listed below. All callback functions are passed the `extra_state` argument that was associated with the request by the starting call `MPI_GREQUEST_START`; `extra_state` can be used to maintain user-defined state for the request.

In C, the query procedure is

```

41 typedef int MPI_Grequest_query_function(void *extra_state, MPI_Status *status);

```

in Fortran with the `mpi_f08` module

```

42 ABSTRACT INTERFACE
43
44 SUBROUTINE MPI_Grequest_query_function(extra_state, status, ierror)
45     INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
46     TYPE(MPI_Status) :: status

```

```
INTEGER :: ierror
```

```
in Fortran with the mpi module and (deprecated) mpif.h include file
SUBROUTINE GREQUEST_QUERY_FUNCTION(EXTRA_STATE, STATUS, IERROR)
  INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
  INTEGER STATUS(MPI_STATUS_SIZE), IERROR
```

The `query_fn` function computes the status that should be returned for the generalized request. The status also includes information about successful/unsuccesful cancellation of the request (result to be returned by `MPI_TEST_CANCELLED`).

The `query_fn` callback is invoked by the `MPI_{WAIT|TEST}{ANY|SOME|ALL}` call that completed the generalized request associated with this callback. The callback function is also invoked by calls to `MPI_REQUEST_GET_STATUS`, if the request is complete when the call occurs. In both cases, the callback is passed a reference to the corresponding status variable passed by the user to the MPI call; the status set by the callback function is returned by the MPI call. If the user provided `MPI_STATUS_IGNORE` or `MPI_STATUSES_IGNORE` to the MPI procedure that causes `query_fn` to be called, then MPI will pass a valid status object to `query_fn`, and this status will be ignored upon return of the callback function. Note that `query_fn` is invoked only after `MPI_GREQUEST_COMPLETE` is called on the request; it may be invoked several times for the same generalized request, e.g., if the user calls `MPI_REQUEST_GET_STATUS` several times for this request. Note also that a call to `MPI_{WAIT|TEST}{SOME|ALL}` may cause multiple invocations of `query_fn` callback functions, one for each generalized request that is completed by the MPI call. The order of these invocations is not specified by MPI.

In C, the free procedures is

```
typedef int MPI_Grequest_free_function(void *extra_state);
```

```
in Fortran with the mpi_f08 module
ABSTRACT INTERFACE
  SUBROUTINE MPI_Grequest_free_function(extra_state, ierror)
    INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
    INTEGER :: ierror
```

```
in Fortran with the mpi module and (deprecated) mpif.h include file
SUBROUTINE GREQUEST_FREE_FUNCTION(EXTRA_STATE, IERROR)
  INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
  INTEGER IERROR
```

The `free_fn` function is invoked to clean up user-allocated resources when the generalized request is freed.

The `free_fn` callback is invoked by the `MPI_{WAIT|TEST}{ANY|SOME|ALL}` call that completed the generalized request associated with this callback. `free_fn` is invoked after the call to `query_fn` for the same request. However, if the MPI call completed multiple generalized requests, the order in which `free_fn` callback functions are invoked is not specified by MPI.

The `free_fn` callback is also invoked for generalized requests that are freed by a call to `MPI_REQUEST_FREE` (no call to `MPI_{WAIT|TEST}{ANY|SOME|ALL}` will occur for such a request). In this case, the callback function will be called either in the MPI call `MPI_REQUEST_FREE(request)`, or in the MPI call `MPI_GREQUEST_COMPLETE(request)`, whichever happens last, i.e., in this case the actual freeing code is executed as soon as both calls `MPI_REQUEST_FREE` and `MPI_GREQUEST_COMPLETE` have occurred. The request

1 is not deallocated until after `free_fn` completes. Note that `free_fn` will be invoked only once
 2 per request by a correct program.

3
 4 *Advice to users.* Calling `MPI_REQUEST_FREE(request)` will cause the `request` handle
 5 to be set to `MPI_REQUEST_NULL`. This handle to the generalized request is no longer
 6 valid. However, user copies of this handle are valid until after `free_fn` completes since
 7 MPI does not deallocate the object until then. Since `free_fn` is not called until after
 8 `MPI_GREQUEST_COMPLETE`, the user copy of the handle can be used to make this
 9 call. Users should note that MPI will deallocate the object after `free_fn` executes. At
 10 this point, user copies of the `request` handle no longer point to a valid request. MPI will
 11 not set user copies to `MPI_REQUEST_NULL` in this case, so it is up to the user to avoid
 12 accessing this stale handle. This is a special case in which MPI defers deallocating the
 13 object until a later time that is known by the user. (*End of advice to users.*)

14
 15 In C, the cancel procedure is

```
16 typedef int MPI_Grequest_cancel_function(void *extra_state, int complete);
```

17 in Fortran with the `mpi_f08` module

```
18 ABSTRACT INTERFACE
```

```
19   SUBROUTINE MPI_Grequest_cancel_function(extra_state, complete, ierror)
```

```
20     INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
```

```
21     LOGICAL :: complete
```

```
22     INTEGER :: ierror
```

23
 24 in Fortran with the `mpi` module and (deprecated) `mpif.h` include file

```
25 SUBROUTINE GREQUEST_CANCEL_FUNCTION(EXTRA_STATE, COMPLETE, IERROR)
```

```
26   INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
```

```
27   LOGICAL COMPLETE
```

```
28   INTEGER IERROR
```

29 The `cancel_fn` function is invoked to start the cancelation of a generalized request. It
 30 is called by `MPI_CANCEL(request)`. MPI passes `complete = true` to the callback function
 31 if `MPI_GREQUEST_COMPLETE` was already called on the request, and `complete = false`
 32 otherwise.

33 All callback functions return an error code. The code is passed back and dealt with as
 34 appropriate for the error code by the MPI procedure that invoked the callback function. For
 35 example, if error codes are returned then the error code returned by the callback function
 36 will be returned by the MPI procedure that invoked the callback function. In the case of
 37 an `MPI_{WAIT|TEST}{ANY}` call that invokes both `query_fn` and `free_fn`, the MPI call will
 38 return the error code returned by the last callback, namely `free_fn`. If one or more of the
 39 requests in a call to `MPI_{WAIT|TEST}{SOME|ALL}` failed, then the MPI call will return
 40 `MPI_ERR_IN_STATUS`. In such a case, if the MPI call was passed an array of statuses, then
 41 MPI will return in each of the statuses that correspond to a completed generalized request
 42 the error code returned by the corresponding invocation of its `free_fn` callback function.
 43 However, if the MPI procedure was passed `MPI_STATUSES_IGNORE`, then the individual
 44 error codes returned by each callback functions will be lost.

45
 46 *Advice to users.* `query_fn` must *not* set the error field of `status` since `query_fn` may be
 47 called by `MPI_WAIT` or `MPI_TEST`, in which case the error field of `status` should not
 48 change. The MPI library knows the “context” in which `query_fn` is invoked and can

decide correctly when to put the returned error code in the error field of `status`. (*End of advice to users.*)

`MPI_GREQUEST_COMPLETE(request)`

INOUT request generalized request (handle)

C binding

`int MPI_Grequest_complete(MPI_Request request)`

Fortran 2008 binding

`MPI_Grequest_complete(request, ierror)`
 TYPE(MPI_Request), INTENT(IN) :: request
 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding

`MPI_GREQUEST_COMPLETE(REQUEST, IERROR)`
 INTEGER REQUEST, IERROR

The call informs MPI that the operations represented by the generalized request `request` are complete (see definitions in Section 2.4). A call to `MPI_WAIT(request, status)` will return and a call to `MPI_TEST(request, flag, status)` will return `flag = true` only after a call to `MPI_GREQUEST_COMPLETE` has declared that these operations are complete.

MPI imposes no restrictions on the code executed by the callback functions. However, new nonblocking operations should be defined so that the general semantic rules about MPI calls such as `MPI_TEST`, `MPI_REQUEST_FREE`, or `MPI_CANCEL` still hold. For example, these calls are supposed to be local and nonblocking. Therefore, the callback functions `query_fn`, `free_fn`, or `cancel_fn` should invoke blocking MPI communication calls only if the context is such that these calls are guaranteed to return in finite time. Once `MPI_CANCEL` is invoked, the cancelled operation should complete in finite time, irrespective of the state of other MPI processes (the operation has acquired “local” semantics). It should either succeed, or fail without side-effects. The user should guarantee these same properties for newly defined operations.

Advice to implementors. A call to `MPI_GREQUEST_COMPLETE` may unblock a blocked user process/thread. The MPI library should ensure that the blocked user computation will resume. (*End of advice to implementors.*)

13.2.1 Examples

Example 13.1. This example shows the code for a user-defined reduce operation on an `int` using a binary tree: each nonroot node receives two messages, sums them, and sends them up. We assume that no status is returned and that the operation cannot be cancelled.

```
typedef struct {
  MPI_Comm comm;
  int tag;
  int root;
  int valin;
  int *valout;
```

```

1  MPI_Request request;
2  } ARGS;
3
4
5  int myreduce(MPI_Comm comm, int tag, int root,
6              int valin, int *valout, MPI_Request *request)
7  {
8      ARGS *args;
9      pthread_t thread;
10
11     /* start request */
12     MPI_Grequest_start(query_fn, free_fn, cancel_fn, NULL, request);
13
14     args = (ARGS*)malloc(sizeof(ARGS));
15     args->comm = comm;
16     args->tag = tag;
17     args->root = root;
18     args->valin = valin;
19     args->valout = valout;
20     args->request = *request;
21
22     /* spawn thread to handle request */
23     /* The availability of the pthread_create call is system dependent */
24     pthread_create(&thread, NULL, reduce_thread, args);
25
26     return MPI_SUCCESS;
27 }
28
29 /* thread code */
30 void* reduce_thread(void *ptr)
31 {
32     int lchild, rchild, parent, lval, rval, val;
33     MPI_Request req[2];
34     ARGS *args;
35
36     args = (ARGS*)ptr;
37
38     /* compute left and right child and parent in tree; set
39        to MPI_PROC_NULL if does not exist */
40     /* code not shown */
41     ...
42
43     MPI_Irecv(&lval, 1, MPI_INT, lchild, args->tag, args->comm, &req[0]);
44     MPI_Irecv(&rval, 1, MPI_INT, rchild, args->tag, args->comm, &req[1]);
45     MPI_Waitall(2, req, MPI_STATUSES_IGNORE);
46     val = lval + args->valin + rval;
47     MPI_Send(&val, 1, MPI_INT, parent, args->tag, args->comm);
48     if (parent == MPI_PROC_NULL) *(args->valout) = val;
49     MPI_Grequest_complete((args->request));
50     free(ptr);
51     return(NULL);
52 }

```

```

1  int query_fn(void *extra_state, MPI_Status *status)
2  {
3      /* always send just one int */
4      MPI_Status_set_elements(status, MPI_INT, 1);
5      /* can never cancel so always true */
6      MPI_Status_set_cancelled(status, 0);
7      /* choose not to return a value for this */
8      status->MPI_SOURCE = MPI_UNDEFINED;
9      /* tag has no meaning for this generalized request */
10     status->MPI_TAG = MPI_UNDEFINED;
11     /* this generalized request never fails */
12     return MPI_SUCCESS;
13 }
14
15 int free_fn(void *extra_state)
16 {
17     /* this generalized request does not need to do any freeing */
18     /* as a result it never fails here */
19     return MPI_SUCCESS;
20 }
21
22 int cancel_fn(void *extra_state, int complete)
23 {
24     /* This generalized request does not support cancelling.
25      * Abort if not already done.
26      * If done then treat as if cancel failed.*/
27     if (!complete) {
28         fprintf(stderr,
29             "Cannot cancel generalized request - aborting program\n");
30         MPI_Abort(MPI_COMM_WORLD, 99);
31     }
32     return MPI_SUCCESS;
33 }

```

13.3 Associating Information with Status

MPI supports several different types of requests besides those for point-to-point operations. These range from MPI calls for I/O to generalized requests. It is desirable to allow these calls to use the same request mechanism, which allows one to wait or test on different types of requests. However, `MPI_{TEST|WAIT}{ANY|SOME|ALL}` returns a status with information about the request. With the generalization of requests, one needs to define what information will be returned in the status object.

Each MPI call fills in the appropriate fields in the status object. Any unused fields will have undefined values. A call to `MPI_{TEST|WAIT}{ANY|SOME|ALL}` can modify any of the fields in the status object. Specifically, it can modify fields that are undefined. The fields with meaningful values for a given request are defined in the sections with the new request.

Generalized requests raise additional considerations. Here, the user provides the func-

tions to deal with the request. Unlike other MPI calls, the user needs to provide the information to be returned in the status. The status argument is provided directly to the callback function where the status needs to be set. Users can directly set the values in 3 of the 5 status values. The count and cancel fields are opaque. To overcome this, these calls are provided:

```
MPI_STATUS_SET_ELEMENTS(status, datatype, count)
```

INOUT	status	status with which to associate count (status)
IN	datatype	datatype associated with count (handle)
IN	count	number of elements to associate with status (integer)

C binding

```
int MPI_Status_set_elements(MPI_Status *status, MPI_Datatype datatype,
                           int count)
```

```
int MPI_Status_set_elements_c(MPI_Status *status, MPI_Datatype datatype,
                              MPI_Count count)
```

Fortran 2008 binding

```
MPI_Status_set_elements(status, datatype, count, ierror)
```

```
TYPE(MPI_Status), INTENT(INOUT) :: status
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
INTEGER, INTENT(IN) :: count
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_Status_set_elements(status, datatype, count, ierror) !(_c)
```

```
TYPE(MPI_Status), INTENT(INOUT) :: status
```

```
TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_STATUS_SET_ELEMENTS(STATUS, DATATYPE, COUNT, IERROR)
```

```
INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR
```

This procedure modifies the opaque part of status so calls to MPI_GET_ELEMENTS will return count. Calls to MPI_GET_COUNT will return a compatible value.

Rationale. The number of elements is set instead of the count because the former can deal with a non-integer number of datatypes. (*End of rationale.*)

A subsequent call to MPI_GET_COUNT or MPI_GET_ELEMENTS must use a datatype argument that has the same type signature as the datatype argument that was used in the call to MPI_STATUS_SET_ELEMENTS.

Rationale. The requirement of matching type signatures for these calls is similar to the restriction that holds when count is set by a receive operation: in that case, calls to MPI_GET_COUNT and MPI_GET_ELEMENTS must use a datatype with the same signature as the datatype used in the receive call. (*End of rationale.*)

MPI_STATUS_SET_CANCELLED(status, flag)			1
INOUT	status	status with which to associate cancel flag (status)	2
			3
IN	flag	if true, indicates request was cancelled (logical)	4
			5

C binding

```
int MPI_Status_set_cancelled(MPI_Status *status, int flag)
```

Fortran 2008 binding

```
MPI_Status_set_cancelled(status, flag, ierror)
```

```
TYPE(MPI_Status), INTENT(INOUT) :: status
```

```
LOGICAL, INTENT(IN) :: flag
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_STATUS_SET_CANCELLED(STATUS, FLAG, IERROR)
```

```
INTEGER STATUS(MPI_STATUS_SIZE), IERROR
```

```
LOGICAL FLAG
```

If flag is set to true then a subsequent call to `MPI_TEST_CANCELLED` will also return flag = true, otherwise it will return false.

Advice to users. Users are advised not to reuse the status fields for values other than those for which they were intended. Doing so may lead to unexpected results when using the status object. For example, calling `MPI_GET_ELEMENTS` may cause an error if the value is out of range or it may be impossible to detect such an error. The `extra_state` argument provided with a generalized request can be used to return information that does not logically belong in status. Furthermore, modifying the values in a status set internally by MPI, e.g., `MPI_RECV`, may lead to unpredictable results and is strongly discouraged. (*End of advice to users.*)

While the `MPI_SOURCE`, `MPI_TAG`, and `MPI_ERROR` status values are directly accessible by the user, for convenience in some contexts, users can also modify them via the procedure calls described below. Procedures for querying these fields from a status object are defined in Section 3.2.5.

MPI_STATUS_SET_SOURCE(status, source)			36
INOUT	status	status with which to associate source rank (status)	37
			38
IN	source	rank to set in the <code>MPI_SOURCE</code> field (integer)	39
			40

C binding

```
int MPI_Status_set_source(MPI_Status *status, int source)
```

Fortran 2008 binding

```
MPI_Status_set_source(status, source, ierror)
```

```
TYPE(MPI_Status), INTENT(INOUT) :: status
```

```
INTEGER, INTENT(IN) :: source
```

```
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```

MPI_STATUS_SET_SOURCE(STATUS, SOURCE, IERROR)
    INTEGER STATUS(MPI_STATUS_SIZE), SOURCE, IERROR

```

Set the MPI_SOURCE field in the status object to the provided source argument.

```

MPI_STATUS_SET_TAG(status, tag)

```

INOUT	status	status with which to associate tag (status)
IN	tag	tag to set in the MPI_TAG field (integer)

C binding

```

int MPI_Status_set_tag(MPI_Status *status, int tag)

```

Fortran 2008 binding

```

MPI_Status_set_tag(status, tag, ierror)
    TYPE(MPI_Status), INTENT(INOUT) :: status
    INTEGER, INTENT(IN) :: tag
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_STATUS_SET_TAG(STATUS, TAG, IERROR)
    INTEGER STATUS(MPI_STATUS_SIZE), TAG, IERROR

```

Set the MPI_TAG field in the status object to the provided tag argument.

```

MPI_STATUS_SET_ERROR(status, err)

```

INOUT	status	status with which to associate error (status)
IN	err	error to set in the MPI_ERROR field (integer)

C binding

```

int MPI_Status_set_error(MPI_Status *status, int err)

```

Fortran 2008 binding

```

MPI_Status_set_error(status, err, ierror)
    TYPE(MPI_Status), INTENT(INOUT) :: status
    INTEGER, INTENT(IN) :: err
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_STATUS_SET_ERROR(STATUS, ERR, IERROR)
    INTEGER STATUS(MPI_STATUS_SIZE), ERR, IERROR

```

Set the MPI_ERROR field in the status object to the provided err error code.

Rationale. These functions exist for convenience when using MPI from languages other than C and Fortran, where having a function in the MPI library with a known API reduces the need for utility code written in C. (*End of rationale.*)

Chapter 14

I/O

14.1 Introduction

POSIX provides a model of a widely portable file system, but the portability and optimization needed for parallel I/O cannot be achieved with the POSIX interface.

The significant optimizations required for efficiency (e.g., grouping [55], collective buffering [9, 17, 56, 60, 67], and disk-directed I/O [50]) can only be implemented if the parallel I/O system provides a high-level interface supporting partitioning of file data among processes and a collective interface supporting complete transfers of global data structures between process memories and files. In addition, further efficiencies can be gained via support for asynchronous I/O, strided accesses, and control over physical file layout on storage devices (disks). The I/O environment described in this chapter provides these facilities.

Instead of defining I/O access modes to express the common patterns for accessing a shared file (broadcast, reduction, scatter, gather), we chose another approach in which data partitioning is expressed using derived datatypes. Compared to a limited set of predefined access patterns, this approach has the advantage of added flexibility and expressiveness.

14.1.1 Definitions

file: An MPI file is an ordered collection of typed data items. MPI supports random or sequential access to any integral set of these items. A file is opened collectively by a group of processes. All collective I/O calls on a file are collective over this group.

displacement: A file *displacement* is an absolute byte position relative to the beginning of a file. The displacement defines the location where a *view* begins. Note that a “file displacement” is distinct from a “typemap displacement.”

etype: An *etype* (*elementary* datatype) is the unit of data access and positioning. It can be any MPI predefined or derived datatype. Derived etypes can be constructed using any of the MPI datatype constructor routines, provided all resulting typemap displacements are nonnegative and monotonically nondecreasing. Data access is performed in etype units, reading or writing whole data items of type etype. Offsets are expressed as a count of etypes; file pointers point to the beginning of etypes. Depending on context, the term “etype” is used to describe one of three aspects of an elementary datatype: a particular MPI type, a data item of that type, or the extent of that type.

filetype: A *filetype* is the basis for partitioning a file among processes and defines a template for accessing the file. A filetype is either a single etype or a derived MPI datatype constructed from multiple instances of the same etype. In addition, the extent of any hole in the filetype must be a multiple of the etype’s extent. The displacements in the

1 typemap of the filetype are not required to be distinct, but they must be nonnegative
2 and monotonically nondecreasing.

3
4 **view:** A *view* defines the current set of data visible and accessible from an open file as
5 an ordered set of etypes. Each process has its own view of the file, defined by three
6 quantities: a displacement, an etype, and a filetype. The pattern described by a
7 filetype is repeated, beginning at the displacement, to define the view. The pattern
8 of repetition is defined to be the same pattern that `MPI_TYPE_CONTIGUOUS` would
9 produce if it were passed the filetype and an arbitrarily large count. Figure 14.1 shows
10 how the tiling works; note that the filetype in this example must have explicit lower
11 and upper bounds set in order for the initial and final holes to be repeated in the
12 view. Views can be changed by the user during program execution. The default view
13 is a linear byte stream (displacement is zero, etype and filetype equal to `MPI_BYTE`).

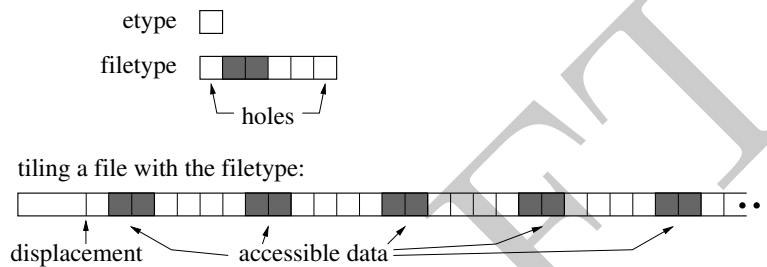


Figure 14.1: Etypes and filetypes

14
15
16
17
18
19
20
21
22
23
24
25 A group of processes can use complementary views to achieve a global data distribution
26 such as a scatter/gather pattern (see Figure 14.2).



Figure 14.2: Partitioning a file among parallel processes

27
28
29
30
31
32
33
34
35
36
37
38
39 **offset:** An *offset* is a position in the file relative to the current view, expressed as a count of
40 etypes. Holes in the view’s filetype are skipped when calculating this position. Offset 0
41 is the location of the first etype visible in the view (after skipping the displacement and
42 any initial holes in the view). For example, an offset of 2 for process 1 in Figure 14.2 is
43 the position of the eighth etype in the file after the displacement. An “explicit offset”
44 is an offset that is used as an argument in explicit data access routines.

45
46
47
48
file size and end of file: The *size* of an MPI file is measured in bytes from the beginning
of the file. A newly created file has a size of zero bytes. Using the size as an absolute
displacement gives the position of the byte immediately following the last byte in the

file. For any given view, the *end of file* is the offset of the first etype accessible in the current view starting after the last byte in the file.

file pointer: A *file pointer* is an implicit offset maintained by MPI. “Individual file pointers” are file pointers that are local to each process that opened the file. A “shared file pointer” is a file pointer that is shared by the group of processes that opened the file.

file handle: A *file handle* is an opaque object created by `MPI_FILE_OPEN` and freed by `MPI_FILE_CLOSE`. All operations on an open file reference the file through the file handle.

14.2 File Manipulation

14.2.1 Opening a File

`MPI_FILE_OPEN(comm, filename, amode, info, fh)`

IN	comm	communicator (handle)
IN	filename	name of file to open (string)
IN	amode	file access mode (integer)
IN	info	info object (handle)
OUT	fh	new file handle (handle)

C binding

```
int MPI_File_open(MPI_Comm comm, const char *filename, int amode,
                 MPI_Info info, MPI_File *fh)
```

Fortran 2008 binding

```
MPI_File_open(comm, filename, amode, info, fh, ierror)
  TYPE(MPI_Comm), INTENT(IN) :: comm
  CHARACTER(LEN=*), INTENT(IN) :: filename
  INTEGER, INTENT(IN) :: amode
  TYPE(MPI_Info), INTENT(IN) :: info
  TYPE(MPI_File), INTENT(OUT) :: fh
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_OPEN(COMM, FILENAME, AMODE, INFO, FH, IERROR)
  INTEGER COMM, AMODE, INFO, FH, IERROR
  CHARACTER*(*) FILENAME
```

`MPI_FILE_OPEN` opens the file identified by the file name `filename` on all processes in the `comm` communicator group. `MPI_FILE_OPEN` is a collective routine: all processes must provide the same value for `amode`, and all processes must provide `filenames` that reference the same file. (Values for `info` may vary.) `comm` must be an intra-communicator; it is erroneous to pass an inter-communicator to `MPI_FILE_OPEN`. Errors in `MPI_FILE_OPEN` are raised using the default file error handler (see Section 14.7). When using the World Model (Section 11.1), a process can open a file independently of other processes by using

1 the MPI_COMM_SELF communicator. Applications using the Sessions Model (Section 11.3)
 2 can achieve the same result using communicators created from the "mpi://SELF" process
 3 set. The file handle returned, fh, can be subsequently used to access the file until the file is
 4 closed using MPI_FILE_CLOSE. Before calling MPI_FINALIZE, the user is required to close
 5 (via MPI_FILE_CLOSE) all files that were opened with MPI_FILE_OPEN. Note that the
 6 communicator comm is unaffected by MPI_FILE_OPEN and continues to be usable in all
 7 MPI routines (e.g., MPI_SEND). Furthermore, the use of comm will not interfere with I/O
 8 behavior.

9 The format for specifying the file name in the filename argument is implementation
 10 dependent and must be documented by the implementation.

11
 12 *Advice to implementors.* An implementation may require that filename include a
 13 string or strings specifying additional information about the file. Examples include
 14 the type of filesystem (e.g., a prefix of ufs:), a remote hostname (e.g., a prefix of
 15 machine.univ.edu:), or a file password (e.g., a suffix of /PASSWORD=SECRET). (*End of*
 16 *advice to implementors.*)

17
 18 *Advice to users.* On some implementations of MPI, the file namespace may not be
 19 identical from all processes of all applications. For example, "/tmp/foo" may denote
 20 different files on different processes, or a single file may have many names, dependent
 21 on process location. The user is responsible for ensuring that a single file is referenced
 22 by the filename argument, as it may be impossible for an implementation to detect
 23 this type of namespace error. (*End of advice to users.*)

24
 25 Initially, all processes view the file as a linear byte stream, and each process views data
 26 in its own native representation (no data representation conversion is performed). (POSIX
 27 files are linear byte streams in the native representation.) The file view can be changed via
 28 the MPI_FILE_SET_VIEW routine.

29 The following access modes are supported (specified in amode, a bit vector OR of the
 30 following integer constants):

31 MPI_MODE_RDONLY	read only
32 MPI_MODE_RDWR	reading and writing
33 MPI_MODE_WRONLY	write only
34 MPI_MODE_CREATE	create the file if it does not exist
35 MPI_MODE_EXCL	error if creating file that already exists
36 MPI_MODE_DELETE_ON_CLOSE	delete file on close
37 MPI_MODE_UNIQUE_OPEN	file will not be concurrently opened elsewhere
38 MPI_MODE_SEQUENTIAL	file will only be accessed sequentially
39 MPI_MODE_APPEND	set initial position of all file pointers to end of file

40
 41 *Advice to users.* C users can use bit vector OR (|) to combine these constants; Fortran
 42 90 users can use the bit vector IOR intrinsic. Fortran 77 users can use (nonportably)
 43 bit vector IOR on systems that support it. Alternatively, Fortran users can portably
 44 use integer addition to OR the constants (each constant should appear at most once
 45 in the addition.). (*End of advice to users.*)

46
 47 *Advice to implementors.* The values of these constants must be defined such that
 48 the bitwise OR and the sum of any distinct set of these constants is equivalent. (*End*
of advice to implementors.)

The modes `MPI_MODE_RDONLY`, `MPI_MODE_RDWR`, `MPI_MODE_WRONLY`, `MPI_MODE_CREATE`, and `MPI_MODE_EXCL` have identical semantics to their POSIX counterparts [45]. Exactly one of `MPI_MODE_RDONLY`, `MPI_MODE_RDWR`, or `MPI_MODE_WRONLY`, must be specified. It is erroneous to specify `MPI_MODE_CREATE` or `MPI_MODE_EXCL` in conjunction with `MPI_MODE_RDONLY`; it is erroneous to specify `MPI_MODE_SEQUENTIAL` together with `MPI_MODE_RDWR`.

The `MPI_MODE_DELETE_ON_CLOSE` mode causes the file to be deleted (equivalent to performing an `MPI_FILE_DELETE`) when the file is closed.

The `MPI_MODE_UNIQUE_OPEN` mode allows an implementation to optimize access by eliminating the overhead of file locking. It is erroneous to open a file in this mode unless the file will not be concurrently opened elsewhere.

Advice to users. For `MPI_MODE_UNIQUE_OPEN`, *not opened elsewhere* includes both inside and outside the MPI environment. In particular, one needs to be aware of potential external events that may open files (e.g., automated backup facilities). When `MPI_MODE_UNIQUE_OPEN` is specified, the user is responsible for ensuring that no such external events take place. (*End of advice to users.*)

The `MPI_MODE_SEQUENTIAL` mode allows an implementation to optimize access to some sequential devices (tapes and network streams). It is erroneous to attempt nonsequential access to a file that has been opened in this mode.

Specifying `MPI_MODE_APPEND` only guarantees that all shared and individual file pointers are positioned at the initial end of file when `MPI_FILE_OPEN` returns. Subsequent positioning of file pointers is application dependent. In particular, the implementation does not ensure that all writes are appended.

Errors related to the access mode are raised in the class `MPI_ERR_AMODE`.

The `info` argument is used to provide information regarding file access patterns and file system specifics (see Section 14.2.8). The constant `MPI_INFO_NULL` can be used when no `info` needs to be specified.

Advice to users. Some file attributes are inherently implementation dependent (e.g., file permissions). These attributes must be set using either the `info` argument or facilities outside the scope of MPI. (*End of advice to users.*)

Files are opened by default using nonatomic mode file consistency semantics (see Section 14.6.1). The more stringent atomic mode consistency semantics, required for atomicity of conflicting accesses, can be set using `MPI_FILE_SET_ATOMICITY`.

14.2.2 Closing a File

`MPI_FILE_CLOSE(fh)`

INOUT `fh` file handle (handle)

C binding

`int MPI_File_close(MPI_File *fh)`

Fortran 2008 binding

`MPI_File_close(fh, ierror)`

```

1      TYPE(MPI_File), INTENT(INOUT) :: fh
2      INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

4      MPI_FILE_CLOSE(FH, IERROR)
5      INTEGER FH, IERROR

```

6
7 MPI_FILE_CLOSE first synchronizes file state (equivalent to performing an
8 MPI_FILE_SYNC), then closes the file associated with fh. The file is deleted if it was opened
9 with access mode MPI_MODE_DELETE_ON_CLOSE (equivalent to performing an
10 MPI_FILE_DELETE). MPI_FILE_CLOSE is a collective routine.

11
12 *Advice to users.* If the file is deleted on close, and there are other processes currently
13 accessing the file, the status of the file and the behavior of future accesses by these
14 processes are implementation dependent. (*End of advice to users.*)

15
16 The user is responsible for ensuring that all outstanding nonblocking requests and
17 split collective operations associated with fh made by a process have completed before that
18 process calls MPI_FILE_CLOSE.

19 The MPI_FILE_CLOSE routine deallocates the file handle object and sets fh to
20 MPI_FILE_NULL.

14.2.3 Deleting a File

```

25      MPI_FILE_DELETE(filename, info)
26      IN          filename          name of file to delete (string)
27      IN          info              info object (handle)

```

C binding

```

31      int MPI_File_delete(const char *filename, MPI_Info info)

```

Fortran 2008 binding

```

33      MPI_File_delete(filename, info, ierror)
34      CHARACTER(LEN=*) INTENT(IN) :: filename
35      TYPE(MPI_Info), INTENT(IN) :: info
36      INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

39      MPI_FILE_DELETE(FILENAME, INFO, IERROR)
40      CHARACTER*(*) FILENAME
41      INTEGER INFO, IERROR

```

42 MPI_FILE_DELETE deletes the file identified by the file name filename. If the file does
43 not exist, MPI_FILE_DELETE raises an error in the class MPI_ERR_NO_SUCH_FILE.

44 The info argument can be used to provide information regarding file system specifics
45 (see Section 14.2.8). The constant MPI_INFO_NULL refers to the null info, and can be used
46 when no info needs to be specified.

47 If a process currently has the file open, the behavior of any access to the file (as well
48 as the behavior of any outstanding accesses) is implementation dependent. In addition,

whether an open file is deleted or not is also implementation dependent. If the file is not deleted, an error in the class `MPI_ERR_FILE_IN_USE` or `MPI_ERR_ACCESS` will be raised. Errors are raised using the default file error handler (see Section 14.7).

14.2.4 Resizing a File

`MPI_FILE_SET_SIZE(fh, size)`

INOUT	fh	file handle (handle)
IN	size	size to truncate or expand file (integer)

C binding

```
int MPI_File_set_size(MPI_File fh, MPI_Offset size)
```

Fortran 2008 binding

```
MPI_File_set_size(fh, size, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_SET_SIZE(FH, SIZE, IERROR)
  INTEGER FH, IERROR
  INTEGER(KIND=MPI_OFFSET_KIND) SIZE
```

`MPI_FILE_SET_SIZE` resizes the file associated with the file handle `fh`. `size` is measured in bytes from the beginning of the file. `MPI_FILE_SET_SIZE` is collective; all processes in the group must pass identical values for `size`.

If `size` is smaller than the current file size, the file is truncated at the position defined by `size`. The implementation is free to deallocate file blocks located beyond this position.

If `size` is larger than the current file size, the file size becomes `size`. Regions of the file that have been previously written are unaffected. The values of data in the new regions in the file (those locations with displacements between old file size and `size`) are undefined. It is implementation dependent whether the `MPI_FILE_SET_SIZE` routine allocates file space—use `MPI_FILE_PREALLOCATE` to force file space to be reserved.

`MPI_FILE_SET_SIZE` does not affect the individual file pointers or the shared file pointer. If `MPI_MODE_SEQUENTIAL` mode was specified when the file was opened, it is erroneous to call this routine.

Advice to users. It is possible for the file pointers to point beyond the end of file after a `MPI_FILE_SET_SIZE` operation truncates a file. This is valid, and equivalent to seeking beyond the current end of file. (*End of advice to users.*)

All nonblocking requests and split collective operations on `fh` must be completed before calling `MPI_FILE_SET_SIZE`. Otherwise, calling `MPI_FILE_SET_SIZE` is erroneous. As far as consistency semantics are concerned, `MPI_FILE_SET_SIZE` is a write operation that conflicts with operations that access bytes at displacements between the old and new file sizes (see Section 14.6.1).

14.2.5 Preallocating Space for a File

MPI_FILE_PREALLOCATE(fh, size)

INOUT	fh	file handle (handle)
IN	size	size to preallocate file (integer)

C binding

```
int MPI_File_preallocate(MPI_File fh, MPI_Offset size)
```

Fortran 2008 binding

```
MPI_File_preallocate(fh, size, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_PREALLOCATE(FH, SIZE, IERROR)
  INTEGER FH, IERROR
  INTEGER(KIND=MPI_OFFSET_KIND) SIZE
```

MPI_FILE_PREALLOCATE ensures that storage space is allocated for the first `size` bytes of the file associated with `fh`. MPI_FILE_PREALLOCATE is collective; all processes in the group must pass identical values for `size`. Regions of the file that have previously been written are unaffected. For newly allocated regions of the file, MPI_FILE_PREALLOCATE has the same effect as writing undefined data. If `size` is larger than the current file size, the file size increases to `size`. If `size` is less than or equal to the current file size, the file size is unchanged.

The treatment of file pointers, nonblocking accesses, and file consistency is the same as with MPI_FILE_SET_SIZE. If MPI_MODE_SEQUENTIAL mode was specified when the file was opened, it is erroneous to call this routine.

Advice to users. In some implementations, file preallocation may be time-consuming.
(*End of advice to users.*)

14.2.6 Querying the Size of a File

MPI_FILE_GET_SIZE(fh, size)

IN	fh	file handle (handle)
OUT	size	size of the file in bytes (integer)

C binding

```
int MPI_File_get_size(MPI_File fh, MPI_Offset *size)
```

Fortran 2008 binding

```
MPI_File_get_size(fh, size, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
```

```

    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: size
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_GET_SIZE(FH, SIZE, IERROR)
    INTEGER FH, IERROR
    INTEGER(KIND=MPI_OFFSET_KIND) SIZE

```

MPI_FILE_GET_SIZE returns, in `size`, the current size in bytes of the file associated with the file handle `fh`. As far as consistency semantics are concerned, MPI_FILE_GET_SIZE is a data access operation (see Section 14.6.1).

14.2.7 Querying File Parameters

```

MPI_FILE_GET_GROUP(fh, group)

```

IN	fh	file handle (handle)
OUT	group	group that opened the file (handle)

C binding

```

int MPI_File_get_group(MPI_File fh, MPI_Group *group)

```

Fortran 2008 binding

```

MPI_File_get_group(fh, group, ierror)
    TYPE(MPI_File), INTENT(IN) :: fh
    TYPE(MPI_Group), INTENT(OUT) :: group
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_GET_GROUP(FH, GROUP, IERROR)
    INTEGER FH, GROUP, IERROR

```

MPI_FILE_GET_GROUP returns a duplicate of the group of the communicator used to open the file associated with `fh`. The group is returned in `group`. The user is responsible for freeing `group`.

```

MPI_FILE_GET_AMODE(fh, amode)

```

IN	fh	file handle (handle)
OUT	amode	file access mode used to open the file (integer)

C binding

```

int MPI_File_get_amode(MPI_File fh, int *amode)

```

Fortran 2008 binding

```

MPI_File_get_amode(fh, amode, ierror)
    TYPE(MPI_File), INTENT(IN) :: fh
    INTEGER, INTENT(OUT) :: amode
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_GET_AMODE(FH, AMODE, IERROR)
  INTEGER FH, AMODE, IERROR

```

MPI_FILE_GET_AMODE returns, in `amode`, the access mode of the file associated with `fh`.

Example 14.1. In Fortran 77, decoding an `amode` bit vector will require a routine such as the following:

```

SUBROUTINE BIT_QUERY(TEST_BIT, MAX_BIT, AMODE, BIT_FOUND)
!
! TEST IF THE INPUT TEST_BIT IS SET IN THE INPUT AMODE
! IF SET, RETURN 1 IN BIT_FOUND, 0 OTHERWISE
!
  INTEGER TEST_BIT, AMODE, BIT_FOUND, CP_AMODE, HIFOUND
  INTEGER L, LBIT, MATCHER, MAX_BIT
  BIT_FOUND = 0
  CP_AMODE = AMODE
100 CONTINUE
  LBIT = 0
  HIFOUND = 0
  DO L = MAX_BIT, 0, -1
    MATCHER = 2**L
    IF (CP_AMODE .GE. MATCHER .AND. HIFOUND .EQ. 0) THEN
      HIFOUND = 1
      LBIT = MATCHER
      CP_AMODE = CP_AMODE - MATCHER
    END IF
  END DO
  IF (HIFOUND .EQ. 1 .AND. LBIT .EQ. TEST_BIT) BIT_FOUND = 1
  IF (BIT_FOUND .EQ. 0 .AND. HIFOUND .EQ. 1 .AND. &
    CP_AMODE .GT. 0) GO TO 100
END

```

This routine could be called successively to decode `amode`, one bit at a time. For example, the following code fragment would check for `MPI_MODE_RDONLY`.

```

CALL BIT_QUERY(MPI_MODE_RDONLY, 30, AMODE, BIT_FOUND)
IF (BIT_FOUND .EQ. 1) THEN
  PRINT *, ' FOUND READ-ONLY BIT IN AMODE=', AMODE
ELSE
  PRINT *, ' READ-ONLY BIT NOT FOUND IN AMODE=', AMODE
END IF

```

14.2.8 File Info

Hints specified via `info` (see Chapter 10) allow a user to provide information such as file access patterns and file system specifics to direct optimization. Providing hints may enable an implementation to deliver increased I/O performance or minimize the use of system resources. As described in Section 10, an implementation is free to ignore all hints; however, applications must comply with any `info` hints they provide that are used by the MPI implementation (i.e., are returned by a call to `MPI_FILE_GET_INFO`) and that place a restriction on the behavior of the application. Hints are specified on a per file basis, in

MPI_FILE_OPEN, MPI_FILE_DELETE, MPI_FILE_SET_VIEW, and MPI_FILE_SET_INFO, via the opaque info object. When an info object that specifies a subset of valid hints is passed to MPI_FILE_SET_VIEW or MPI_FILE_SET_INFO, there will be no effect on previously set or defaulted hints that the info does not specify.

Advice to implementors. It may happen that a program is coded with hints for one system, and later executes on another system that does not support these hints. In general, unsupported hints should simply be ignored.

However, for each hint used by a specific implementation, a default value must be provided when the user does not specify a value for this hint. (*End of advice to implementors.*)

MPI_FILE_SET_INFO(fh, info)

INOUT	fh	file handle (handle)
IN	info	info object (handle)

C binding

```
int MPI_File_set_info(MPI_File fh, MPI_Info info)
```

Fortran 2008 binding

```
MPI_File_set_info(fh, info, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  TYPE(MPI_Info), INTENT(IN) :: info
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_SET_INFO(FH, INFO, IERROR)
  INTEGER FH, INFO, IERROR
```

MPI_FILE_SET_INFO updates the hints of the file associated with fh using the hints provided in info. This operation has no effect on previously set or defaulted hints that are not specified by info. It also has no effect on previously set or defaulted hints that are specified by info, but are ignored by the MPI implementation in this call to MPI_FILE_SET_INFO. MPI_FILE_SET_INFO is a collective routine. The info object may be different on each process, but any info entries that an implementation requires to be the same on all processes must appear with the same value in each process's info object.

Advice to users. Many info items that an implementation can use when it creates or opens a file cannot easily be changed once the file has been created or opened. Thus, an implementation may ignore hints issued in this call that it would have accepted in an open call. An implementation may also be unable to update certain info hints in a call to MPI_FILE_SET_VIEW or MPI_FILE_SET_INFO. MPI_FILE_GET_INFO can be used to determine whether info changes were ignored by the implementation. (*End of advice to users.*)

```

1 MPI_FILE_GET_INFO(fh, info_used)
2     IN      fh                file handle (handle)
3
4     OUT    info_used          new info object (handle)
5

```

C binding

```

7 int MPI_File_get_info(MPI_File fh, MPI_Info *info_used)
8

```

Fortran 2008 binding

```

9 MPI_File_get_info(fh, info_used, ierror)
10     TYPE(MPI_File), INTENT(IN) :: fh
11     TYPE(MPI_Info), INTENT(OUT) :: info_used
12     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13

```

Fortran binding

```

14 MPI_FILE_GET_INFO(FH, INFO_USED, IERROR)
15     INTEGER FH, INFO_USED, IERROR
16

```

MPI_FILE_GET_INFO returns a new info object containing the hints of the file associated with `fh`. The current setting of all hints related to this file is returned in `info_used`. An MPI implementation is required to return all hints that are supported by the implementation and have default values specified; any user-supplied hints that were not ignored by the implementation; and any additional hints that were set by the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no (key,value) pairs. The user is responsible for freeing `info_used` via `MPI_INFO_FREE`.

Reserved File Hints

Some potentially useful hints (info key values) are outlined below. The following key values are reserved. An implementation is not required to interpret these key values, but if it does interpret the key value, it must provide the functionality described. (For more details on “info,” see Chapter 10.)

These hints mainly affect access patterns and the layout of data on parallel I/O devices. For each hint name introduced, we describe the purpose of the hint, and the type of the hint value. The “[**SAME**]” annotation specifies that the hint values provided by all participating processes must be identical; otherwise the program is erroneous. In addition, some hints are context dependent, and are only used by an implementation at specific times (e.g., “file_perm” is only useful during file creation).

“access_style” (comma separated list of strings): This hint specifies the manner in which the file will be accessed until the file is closed or until the “access_style” key value is altered. The hint value is a comma separated list of the following: “read_once”, “write_once”, “read_mostly”, “write_mostly”, “sequential”, “reverse_sequential”, and “random”.

“collective_buffering” (boolean) [SAME**]:** This hint specifies whether the application may benefit from collective buffering. Collective buffering is an optimization performed on collective accesses. Accesses to the file are performed on behalf of all processes in the group by a number of target nodes. These target nodes coalesce small requests into large disk accesses. Valid values for this key are “true” and “false”. Collective buffering

parameters are further directed via additional hints: "cb_block_size", "cb_buffer_size", and "cb_nodes".

"cb_block_size" (integer) [SAME]: This hint specifies the block size to be used for collective buffering file access. **Target nodes** access data in chunks of this size. The chunks are distributed among target nodes in a round-robin (cyclic) pattern.

"cb_buffer_size" (integer) [SAME]: This hint specifies the total buffer space that can be used for collective buffering on each target node, usually a multiple of "cb_block_size".

"cb_nodes" (integer) [SAME]: This hint specifies the number of target nodes to be used for collective buffering.

"chunked" (comma separated list of integers) [SAME]: This hint specifies that the file consists of a multidimensional array that is often accessed by subarrays. The value for this hint is a comma separated list of array dimensions, starting from the most significant one (for an array stored in row-major order, as in C, the most significant dimension is the first one; for an array stored in column-major order, as in Fortran, the most significant dimension is the last one, and array dimensions should be reversed).

"chunked_item" (comma separated list of integers) [SAME]: This hint specifies the size of each array entry, in bytes.

"chunked_size" (comma separated list of integers) [SAME]: This hint specifies the dimensions of the subarrays. This is a comma separated list of array dimensions, starting from the most significant one.

"filename" (string): This hint specifies the file name used when the file was opened. If the implementation is capable of returning the file name of an open file, it will be returned using this key by MPI_FILE_GET_INFO. This key is ignored when passed to MPI_FILE_OPEN, MPI_FILE_SET_VIEW, MPI_FILE_SET_INFO, and MPI_FILE_DELETE.

"file_perm" (string) [SAME]: This hint specifies the file permissions to use for file creation. Setting this hint is only useful when passed to MPI_FILE_OPEN with an amode that includes MPI_MODE_CREATE. The set of valid values for this key is implementation dependent.

"io_node_list" (comma separated list of strings) [SAME]: This hint specifies the list of I/O devices that should be used to store the file. This hint is most relevant when the file is created.

"nb_proc" (integer) [SAME]: This hint specifies the number of parallel processes that will typically be assigned to access this file. This hint is most relevant when the file is created.

"num_io_nodes" (integer) [SAME]: This hint specifies the number of I/O devices in the system. This hint is most relevant when the file is created.

"striping_factor" (integer) [SAME]: This hint specifies the number of I/O devices that the file should be striped across, and is relevant only when the file is created.

1 **"striping_unit" (integer) [SAME]:** This hint specifies the suggested striping unit to be
 2 used for this file. The striping unit is the amount of consecutive data assigned to one
 3 I/O device before progressing to the next device, when striping across a number of
 4 devices. It is expressed in bytes. This hint is relevant only when the file is created.

5 **"mpi_assert_memory_alloc_kinds" (string, not set by default):** If set, the implementation
 6 may assume that the memory for all data buffers passed to MPI operations performed
 7 by the calling MPI process on the given file will use only the memory allocation kinds
 8 listed in the value string. See Section 11.4.3.
 9

11 14.3 File Views

12
 13
 14 MPI_FILE_SET_VIEW(fh, disp, etype, filetype, datarep, info)

15	INOUT	fh	file handle (handle)
16	IN	disp	displacement (integer)
17	IN	etype	elementary datatype (handle)
18	IN	filetype	filetype (handle)
19	IN	datarep	data representation (string)
20	IN	info	info object (handle)
21			
22			
23			
24			

25 C binding

26 int MPI_File_set_view(MPI_File fh, MPI_Offset disp, MPI_Datatype etype,
 27 MPI_Datatype filetype, const char *datarep, MPI_Info info)

28 Fortran 2008 binding

29 MPI_File_set_view(fh, disp, etype, filetype, datarep, info, ierror)
 30 TYPE(MPI_File), INTENT(IN) :: fh
 31 INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: disp
 32 TYPE(MPI_Datatype), INTENT(IN) :: etype, filetype
 33 CHARACTER(LEN=*), INTENT(IN) :: datarep
 34 TYPE(MPI_Info), INTENT(IN) :: info
 35 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
 36

37 Fortran binding

38 MPI_FILE_SET_VIEW(FH, DISP, ETYPE, FILETYPE, DATAREP, INFO, IERROR)
 39 INTEGER FH, ETYPE, FILETYPE, INFO, IERROR
 40 INTEGER(KIND=MPI_OFFSET_KIND) DISP
 41 CHARACTER*(*) DATAREP

42 The MPI_FILE_SET_VIEW routine changes the process's view of the data in the file.
 43 The start of the view is set to **disp**; the type of data is set to **etype**; the distribution of data
 44 to processes is set to **filetype**; and the representation of data in the file is set to **datarep**. In
 45 addition, MPI_FILE_SET_VIEW resets the individual file pointers and the shared file pointer
 46 to zero. MPI_FILE_SET_VIEW is collective; the values for **datarep** and the extents of **etype**
 47 in the file data representation must be identical on all processes in the group; values for **disp**,
 48 **filetype**, and **info** may vary. The datatypes passed in **etype** and **filetype** must be committed.

The `etype` always specifies the data layout in the file. If `etype` is a portable datatype (see Section 2.4), the extent of `etype` is computed by scaling any displacements in the datatype to match the file data representation. If `etype` is not a portable datatype, no scaling is done when computing the extent of `etype`. The user must be careful when using nonportable `etypes` in heterogeneous environments; see Section 14.5.1 for further details.

If `MPI_MODE_SEQUENTIAL` mode was specified when the file was opened, the special displacement `MPI_DISPLACEMENT_CURRENT` must be passed in `disp`. This sets the displacement to the current position of the shared file pointer. `MPI_DISPLACEMENT_CURRENT` is invalid unless the `amode` for the file has `MPI_MODE_SEQUENTIAL` set.

Rationale. For some sequential files, such as those corresponding to magnetic tapes or streaming network connections, the *displacement* may not be meaningful. `MPI_DISPLACEMENT_CURRENT` allows the view to be changed for these types of files. (*End of rationale.*)

Advice to implementors. It is expected that a call to `MPI_FILE_SET_VIEW` will immediately follow `MPI_FILE_OPEN` in numerous instances. A high-quality implementation will ensure that this behavior is efficient. (*End of advice to implementors.*)

The `disp` displacement argument specifies the position (absolute offset in bytes from the beginning of the file) where the view begins.

Advice to users. `disp` can be used to skip headers or when the file includes a sequence of data segments that are to be accessed in different patterns (see Figure 14.3). Separate views, each using a different displacement and `filetype`, can be used to access each segment.




Figure 14.3: Displacements

(*End of advice to users.*)

An *etype* (*elementary datatype*) is the unit of data access and positioning. It can be any MPI predefined or derived datatype. Derived `etypes` can be constructed by using any of the MPI datatype constructor routines, provided all resulting typemap displacements are nonnegative and monotonically nondecreasing. Data access is performed in `etype` units, reading or writing whole data items of type `etype`. Offsets are expressed as a count of `etypes`; file pointers point to the beginning of `etypes`.

Advice to users. In order to ensure interoperability in a heterogeneous environment, additional restrictions must be observed when constructing the `etype` (see Section 14.5). (*End of advice to users.*)

1 A filetype is either a single etype or a derived MPI datatype constructed from multiple
 2 instances of the same etype. In addition, the extent of any hole in the filetype must be
 3 a multiple of the etype’s extent. These displacements are not required to be distinct, but
 4 they cannot be negative, and they must be monotonically nondecreasing.

5 If the file is opened for writing, neither the etype nor the filetype is permitted to
 6 contain overlapping regions. This restriction is equivalent to the “datatype used in a receive
 7 cannot specify overlapping regions” restriction for communication. Note that filetypes from
 8 different processes may still overlap each other.

9 If a filetype has holes in it, then the data in the holes is inaccessible to the calling
 10 process. However, the `disp`, `etype`, and `filetype` arguments can be changed via future calls to
 11 `MPI_FILE_SET_VIEW` to access a different part of the file.

12 It is erroneous to use absolute addresses in the construction of the etype and filetype.

13 The `info` argument is used to provide information regarding file access patterns and file
 14 system specifics to direct optimization (see Section 14.2.8). The constant `MPI_INFO_NULL`
 15 refers to the null info and can be used when no info needs to be specified.

16 The `datarep` argument is a string that specifies the representation of data in the file.
 17 See the file interoperability section (Section 14.5) for details and a discussion of valid values.

18 The user is responsible for ensuring that all nonblocking requests and split collective
 19 operations on `fh` have been completed before calling `MPI_FILE_SET_VIEW`—otherwise, the
 20 call to `MPI_FILE_SET_VIEW` is erroneous.

21
 22
 23 `MPI_FILE_GET_VIEW(fh, disp, etype, filetype, datarep)`

24	IN	<code>fh</code>	file handle (handle)
25	OUT	<code>disp</code>	displacement (integer)
26	OUT	<code>etype</code>	elementary datatype (handle)
27	OUT	<code>filetype</code>	filetype (handle)
28	OUT	<code>datarep</code>	data representation (string)

31 C binding

```
32 int MPI_File_get_view(MPI_File fh, MPI_Offset *disp, MPI_Datatype *etype,
33 MPI_Datatype *filetype, char *datarep)
```

34 Fortran 2008 binding

```
35 MPI_File_get_view(fh, disp, etype, filetype, datarep, ierror)
36 TYPE(MPI_File), INTENT(IN) :: fh
37 INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: disp
38 TYPE(MPI_Datatype), INTENT(OUT) :: etype, filetype
39 CHARACTER(LEN=*), INTENT(OUT) :: datarep
40 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

41 Fortran binding

```
42 MPI_FILE_GET_VIEW(FH, DISP, ETYPER, FILETYPE, DATAREP, IERROR)
43 INTEGER FH, ETYPER, FILETYPE, IERROR
44 INTEGER(KIND=MPI_OFFSET_KIND) DISP
45 CHARACTER*(*) DATAREP
```

MPI_FILE_GET_VIEW returns the process's view of the data in the file. The current value of the displacement is returned in `disp`. The `etype` and `filetype` are new datatypes with typemaps equal to the typemaps of the current `etype` and `filetype`, respectively.

The data representation is returned in `datarep`. The user is responsible for ensuring that `datarep` is large enough to hold the returned data representation string. The length of a data representation string is limited to the value of `MPI_MAX_DATAREP_STRING`.

In addition, if a portable datatype was used to set the current view, then the corresponding datatype returned by `MPI_FILE_GET_VIEW` is also a portable datatype. If `etype` or `filetype` are derived datatypes, the user is responsible for freeing them. The `etype` and `filetype` returned are both in a committed state.

14.4 Data Access

14.4.1 Data Access Routines

Data is moved between files and processes by issuing read and write calls. There are three orthogonal aspects to data access: **positioning** (explicit offset vs. implicit file pointer), **synchronism** (blocking vs. nonblocking and split collective), and **coordination** (noncollective vs. collective). The following combinations of these data access routines, including two types of file pointers (individual and shared) are provided in Table 14.1.

Table 14.1: Data access routines

positioning	synchronism	coordination	
		noncollective	collective
explicit offsets	<i>blocking</i>	MPI_FILE_READ_AT MPI_FILE_WRITE_AT	MPI_FILE_READ_AT_ALL MPI_FILE_WRITE_AT_ALL
	<i>nonblocking</i>	MPI_FILE_IREAD_AT MPI_FILE_IWRITE_AT	MPI_FILE_IREAD_AT_ALL MPI_FILE_IWRITE_AT_ALL
	<i>split collective</i>	N/A	MPI_FILE_READ_AT_ALL_BEGIN MPI_FILE_READ_AT_ALL_END MPI_FILE_WRITE_AT_ALL_BEGIN MPI_FILE_WRITE_AT_ALL_END
individual file pointers	<i>blocking</i>	MPI_FILE_READ MPI_FILE_WRITE	MPI_FILE_READ_ALL MPI_FILE_WRITE_ALL
	<i>nonblocking</i>	MPI_FILE_IREAD MPI_FILE_IWRITE	MPI_FILE_IREAD_ALL MPI_FILE_IWRITE_ALL
	<i>split collective</i>	N/A	MPI_FILE_READ_ALL_BEGIN MPI_FILE_READ_ALL_END MPI_FILE_WRITE_ALL_BEGIN MPI_FILE_WRITE_ALL_END
shared file pointer	<i>blocking</i>	MPI_FILE_READ_SHARED MPI_FILE_WRITE_SHARED	MPI_FILE_READ_ORDERED MPI_FILE_WRITE_ORDERED
	<i>nonblocking</i>	MPI_FILE_IREAD_SHARED MPI_FILE_IWRITE_SHARED	N/A
	<i>split collective</i>	N/A	MPI_FILE_READ_ORDERED_BEGIN MPI_FILE_READ_ORDERED_END MPI_FILE_WRITE_ORDERED_BEGIN MPI_FILE_WRITE_ORDERED_END

POSIX `read()/pread()` and `write()/fwrite()` are blocking, noncollective operations and use individual file pointers. The MPI equivalents are `MPI_FILE_READ` and

1 MPI_FILE_WRITE.

2 Implementations of data access routines may buffer data to improve performance. This
 3 does not affect reads, as the data is always available in the user’s buffer after a read operation
 4 completes. For writes, however, the MPI_FILE_SYNC routine provides the only guarantee
 5 that data has been transferred to the storage device.

7 Positioning

8
 9 MPI provides three types of positioning for data access routines: **explicit offsets**, **indi-**
 10 **vidual file pointers**, and **shared file pointers**. The different positioning methods may
 11 be mixed within the same program and do not affect each other.

12 The data access routines that accept explicit offsets contain `_AT` in their name (e.g.,
 13 `MPI_FILE_WRITE_AT`). Explicit offset operations perform data access at the file position
 14 given directly as an argument—no file pointer is used nor updated. Note that this is not
 15 equivalent to an atomic seek-and-read or seek-and-write operation, as no “seek” is issued.
 16 Operations with explicit offsets are described in Section 14.4.2.

17 The names of the individual file pointer routines contain no positional qualifier (e.g.,
 18 `MPI_FILE_WRITE`). Operations with individual file pointers are described in Section 14.4.3.
 19 The data access routines that use shared file pointers contain `_SHARED` or `_ORDERED`
 20 in their name (e.g., `MPI_FILE_WRITE_SHARED`). Operations with shared file pointers are
 21 described in Section 14.4.4.

22 The main semantic issues with MPI-maintained file pointers are how and when they are
 23 updated by I/O operations. In general, each I/O operation leaves the file pointer pointing to
 24 the next data item after the last one that is accessed by the operation. In a nonblocking or
 25 split collective operation, the pointer is updated by the call that initiates the I/O, possibly
 26 before the access completes.

27 More formally,

$$28 \quad new_file_offset = old_file_offset + \frac{elements(datatype)}{elements(etype)} \times count$$

29
 30 where *count* is the number of *datatype* items to be accessed, *elements*(*X*) is the number of
 31 predefined datatypes in the typemap of *X*, and *old_file_offset* is the value of the implicit
 32 offset before the call. The file position, *new_file_offset*, is in terms of a count of etypes
 33 relative to the current view.

35 Synchronism

36
 37 MPI supports blocking and nonblocking I/O routines.

38 A *blocking* I/O call will not return until the I/O request is completed.

39 A *nonblocking* I/O call initiates an I/O operation, but does not wait for it to complete.
 40 Given suitable hardware, this allows the transfer of data out of and into the user’s buffer
 41 to proceed concurrently with computation. A separate *request complete* call (`MPI_WAIT`,
 42 `MPI_TEST`, or any of their variants) is needed to complete the I/O request, i.e., to confirm
 43 that the data has been read or written and that it is safe for the user to reuse the buffer.
 44 The nonblocking versions of the routines are named `MPI_FILE_IXXX`, where the *I* stands for
 45 immediate.

46 It is erroneous to access the local buffer of a nonblocking data access operation, or to
 47 use that buffer as the source or target of other communications, between the initiation and
 48 completion of the operation.

The split collective routines support a restricted form of “nonblocking” operations for collective data access (see Section 14.4.5).

Coordination

Every noncollective data access routine `MPI_FILE_XXX` has a collective counterpart. For most routines, this counterpart is `MPI_FILE_XXX_ALL` or a pair of `MPI_FILE_XXX_BEGIN` and `MPI_FILE_XXX_END`. The counterparts to the `MPI_FILE_XXX_SHARED` routines are `MPI_FILE_XXX_ORDERED`.

The completion of a noncollective call only depends on the activity of the calling process. However, the completion of a collective call (which must be called by all members of the process group) may depend on the activity of the other processes participating in the collective call. See Section 14.6.4 for rules on semantics of collective calls.

Collective operations may perform much better than their noncollective counterparts, as global data accesses have significant potential for automatic optimization.

Data Access Conventions

Data is moved between files and processes by calling read and write routines. Read routines move data from a file into memory. Write routines move data from memory into a file. The file is designated by a file handle, `fh`. The location of the file data is specified by an offset into the current view. The data in memory is specified by a triple: `buf`, `count`, and `datatype`. Upon completion, the amount of data accessed by the calling process is returned in a `status`.

An offset designates the starting position in the file for an access. The offset is always in `etype` units relative to the current view. Explicit offset routines pass `offset` as an argument (negative values are erroneous). The file pointer routines use implicit offsets maintained by MPI.

A data access routine attempts to transfer (read or write) `count` data items of type `datatype` between the user’s buffer `buf` and the file. The `datatype` passed to the routine must be a committed datatype. The layout of data in memory corresponding to `buf`, `count`, `datatype` is interpreted the same way as in MPI communication functions; see Section 3.2.2 and Section 5.1.11. The data is accessed from those parts of the file specified by the current view (Section 14.3). The type signature of `datatype` must match the type signature of some number of contiguous copies of the `etype` of the current view. As in a receive, it is erroneous to specify a `datatype` for reading that contains overlapping regions (areas of memory that would be stored into more than once).

The nonblocking data access routines indicate that MPI can start a data access and associate a request handle, `request`, with the I/O operation. Nonblocking operations are completed via `MPI_TEST`, `MPI_WAIT`, or any of their variants.

Data access operations, when completed, return the amount of data accessed in `status`.

Advice to users. To prevent problems with the argument copying and register optimization done by Fortran compilers, please note the hints in Sections 19.1.10–19.1.20. (*End of advice to users.*)

For blocking routines, `status` is returned directly. For nonblocking routines and split collective routines, `status` is returned when the operation is completed. The number of `datatype` entries and predefined elements accessed by the calling process can be extracted

1 from `status` by using `MPI_GET_COUNT` and `MPI_GET_ELEMENTS`, respectively. The inter-
 2 pretation of the `MPI_ERROR` field is the same as for other operations—normally undefined,
 3 but meaningful if an MPI routine returns `MPI_ERR_IN_STATUS`. The user can pass (in C and
 4 Fortran) `MPI_STATUS_IGNORE` in the `status` argument if the return value of this argument
 5 is not needed. The `status` can be passed to `MPI_TEST_CANCELLED` to determine if the
 6 operation was cancelled. All other fields of `status` are undefined.

7 When reading, a program can detect the end of file by noting that the amount of data
 8 read is less than the amount requested. Writing past the end of file increases the file size.
 9 The amount of data accessed will be the amount requested, unless an error is raised (or a
 10 read reaches the end of file).

12 14.4.2 Data Access with Explicit Offsets

13 If `MPI_MODE_SEQUENTIAL` mode was specified when the file was opened, it is erroneous to
 14 call the routines in this section.

17 `MPI_FILE_READ_AT(fh, offset, buf, count, datatype, status)`

19	IN	fh	file handle (handle)
20	IN	offset	file offset (integer)
21	OUT	buf	initial address of buffer (choice)
22	IN	count	number of elements in buffer (integer)
23	IN	datatype	datatype of each buffer element (handle)
24	IN	datatype	datatype of each buffer element (handle)
25	OUT	status	status object (status)

27 C binding

28 `int MPI_File_read_at(MPI_File fh, MPI_Offset offset, void *buf, int count,`
 29 `MPI_Datatype datatype, MPI_Status *status)`

31 `int MPI_File_read_at_c(MPI_File fh, MPI_Offset offset, void *buf,`
 32 `MPI_Count count, MPI_Datatype datatype, MPI_Status *status)`

33 Fortran 2008 binding

34 `MPI_File_read_at(fh, offset, buf, count, datatype, status, ierror)`

35 `TYPE(MPI_File), INTENT(IN) :: fh`
 36 `INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset`
 37 `TYPE(*), DIMENSION(..) :: buf`
 38 `INTEGER, INTENT(IN) :: count`
 39 `TYPE(MPI_Datatype), INTENT(IN) :: datatype`
 40 `TYPE(MPI_Status) :: status`
 41 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

42 `MPI_File_read_at(fh, offset, buf, count, datatype, status, ierror) !(_c)`

43 `TYPE(MPI_File), INTENT(IN) :: fh`
 44 `INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset`
 45 `TYPE(*), DIMENSION(..) :: buf`
 46 `INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count`
 47 `TYPE(MPI_Datatype), INTENT(IN) :: datatype`
 48

```

TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
1
2
3
Fortran binding
4
MPI_FILE_READ_AT(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)
5
INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
6
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
7
<type> BUF(*)
8
MPI_FILE_READ_AT reads a file beginning at the position specified by offset.
9
10
11
MPI_FILE_READ_AT_ALL(fh, offset, buf, count, datatype, status)
12
IN fh file handle (handle)
13
IN offset file offset (integer)
14
OUT buf initial address of buffer (choice)
15
IN count number of elements in buffer (integer)
16
IN datatype datatype of each buffer element (handle)
17
OUT status status object (status)
18
19
20
21
C binding
22
int MPI_File_read_at_all(MPI_File fh, MPI_Offset offset, void *buf, int count,
23 MPI_Datatype datatype, MPI_Status *status)
24
int MPI_File_read_at_all_c(MPI_File fh, MPI_Offset offset, void *buf,
25 MPI_Count count, MPI_Datatype datatype, MPI_Status *status)
26
27
Fortran 2008 binding
28
MPI_File_read_at_all(fh, offset, buf, count, datatype, status, ierror)
29
TYPE(MPI_File), INTENT(IN) :: fh
30
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
31
TYPE(*), DIMENSION(..) :: buf
32
INTEGER, INTENT(IN) :: count
33
TYPE(MPI_Datatype), INTENT(IN) :: datatype
34
TYPE(MPI_Status) :: status
35
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
MPI_File_read_at_all(fh, offset, buf, count, datatype, status, ierror) !(_c)
37
TYPE(MPI_File), INTENT(IN) :: fh
38
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
39
TYPE(*), DIMENSION(..) :: buf
40
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
41
TYPE(MPI_Datatype), INTENT(IN) :: datatype
42
TYPE(MPI_Status) :: status
43
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45
Fortran binding
46
MPI_FILE_READ_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)
47
INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
48

```

```

1     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
2     <type> BUF(*)

```

3 MPI_FILE_READ_AT_ALL is a collective version of the blocking MPI_FILE_READ_AT
4 interface.

```

5
6
7 MPI_FILE_WRITE_AT(fh, offset, buf, count, datatype, status)

```

```

8     INOUT   fh                file handle (handle)
9
10    IN      offset            file offset (integer)
11
12    IN      buf                initial address of buffer (choice)
13
14    IN      count              number of elements in buffer (integer)
15
16    IN      datatype           datatype of each buffer element (handle)
17
18    OUT     status             status object (status)

```

17 C binding

```

18 int MPI_File_write_at(MPI_File fh, MPI_Offset offset, const void *buf,
19                      int count, MPI_Datatype datatype, MPI_Status *status)
20
21 int MPI_File_write_at_c(MPI_File fh, MPI_Offset offset, const void *buf,
22                        MPI_Count count, MPI_Datatype datatype, MPI_Status *status)

```

23 Fortran 2008 binding

```

24 MPI_File_write_at(fh, offset, buf, count, datatype, status, ierror)
25     TYPE(MPI_File), INTENT(IN) :: fh
26     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
27     TYPE(*), DIMENSION(..), INTENT(IN) :: buf
28     INTEGER, INTENT(IN) :: count
29     TYPE(MPI_Datatype), INTENT(IN) :: datatype
30     TYPE(MPI_Status) :: status
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_File_write_at(fh, offset, buf, count, datatype, status, ierror) !(_c)
34     TYPE(MPI_File), INTENT(IN) :: fh
35     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
36     TYPE(*), DIMENSION(..), INTENT(IN) :: buf
37     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
38     TYPE(MPI_Datatype), INTENT(IN) :: datatype
39     TYPE(MPI_Status) :: status
40     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

41 Fortran binding

```

42 MPI_FILE_WRITE_AT(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)
43     INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
44     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
45     <type> BUF(*)
46
47 MPI_FILE_WRITE_AT writes a file beginning at the position specified by offset.

```

48

MPI_FILE_WRITE_AT_ALL(fh, offset, buf, count, datatype, status)			1
INOUT	fh	file handle (handle)	2
			3
IN	offset	file offset (integer)	4
IN	buf	initial address of buffer (choice)	5
			6
IN	count	number of elements in buffer (integer)	7
IN	datatype	datatype of each buffer element (handle)	8
OUT	status	status object (status)	9

C binding

```

int MPI_File_write_at_all(MPI_File fh, MPI_Offset offset, const void *buf,
                          int count, MPI_Datatype datatype, MPI_Status *status)
int MPI_File_write_at_all_c(MPI_File fh, MPI_Offset offset, const void *buf,
                             MPI_Count count, MPI_Datatype datatype, MPI_Status *status)

```

Fortran 2008 binding

```

MPI_File_write_at_all(fh, offset, buf, count, datatype, status, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_at_all(fh, offset, buf, count, datatype, status, ierror) !(_c)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)
  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
  INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
  <type> BUF(*)

```

MPI_FILE_WRITE_AT_ALL is a collective version of the blocking MPI_FILE_WRITE_AT interface.

```

1 MPI_FILE_IREAD_AT(fh, offset, buf, count, datatype, request)
2   IN      fh                file handle (handle)
3
4   IN      offset            file offset (integer)
5
6   OUT     buf                initial address of buffer (choice)
7
8   IN      count             number of elements in buffer (integer)
9
10  IN      datatype           datatype of each buffer element (handle)
11
12  OUT     request            request object (handle)

```

C binding

```

13 int MPI_File_iread_at(MPI_File fh, MPI_Offset offset, void *buf, int count,
14                      MPI_Datatype datatype, MPI_Request *request)
15
16 int MPI_File_iread_at_c(MPI_File fh, MPI_Offset offset, void *buf,
17                         MPI_Count count, MPI_Datatype datatype, MPI_Request *request)

```

Fortran 2008 binding

```

18 MPI_File_iread_at(fh, offset, buf, count, datatype, request, ierror)
19   TYPE(MPI_File), INTENT(IN) :: fh
20   INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
21   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
22   INTEGER, INTENT(IN) :: count
23   TYPE(MPI_Datatype), INTENT(IN) :: datatype
24   TYPE(MPI_Request), INTENT(OUT) :: request
25   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27 MPI_File_iread_at(fh, offset, buf, count, datatype, request, ierror) !(_c)
28   TYPE(MPI_File), INTENT(IN) :: fh
29   INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
30   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
31   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
32   TYPE(MPI_Datatype), INTENT(IN) :: datatype
33   TYPE(MPI_Request), INTENT(OUT) :: request
34   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

35 MPI_FILE_IREAD_AT(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)
36   INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
37   INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
38   <type> BUF(*)

```

MPI_FILE_IREAD_AT is a nonblocking version of the MPI_FILE_READ_AT interface.

MPI_FILE_IREAD_AT_ALL(fh, offset, buf, count, datatype, request)			1
IN	fh	file handle (handle)	2
IN	offset	file offset (integer)	3
OUT	buf	initial address of buffer (choice)	4
IN	count	number of elements in buffer (integer)	5
IN	datatype	datatype of each buffer element (handle)	6
OUT	request	request object (handle)	7

C binding

```
int MPI_File_iread_at_all(MPI_File fh, MPI_Offset offset, void *buf, int count,
    MPI_Datatype datatype, MPI_Request *request)
```

```
int MPI_File_iread_at_all_c(MPI_File fh, MPI_Offset offset, void *buf,
    MPI_Count count, MPI_Datatype datatype, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_File_iread_at_all(fh, offset, buf, count, datatype, request, ierror)
```

```
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_File_iread_at_all(fh, offset, buf, count, datatype, request, ierror) !(_c)
```

```
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_IREAD_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)
```

```
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
<type> BUF(*)
```

MPI_FILE_IREAD_AT_ALL is a nonblocking version of MPI_FILE_READ_AT_ALL. See Section 14.6.5 for semantics of nonblocking collective file operations.

```

1 MPI_FILE_IWRITE_AT(fh, offset, buf, count, datatype, request)
2     INOUT   fh                file handle (handle)
3
4     IN      offset            file offset (integer)
5
6     IN      buf               initial address of buffer (choice)
7
8     IN      count             number of elements in buffer (integer)
9
10    IN      datatype          datatype of each buffer element (handle)
11
12    OUT     request           request object (handle)

```

C binding

```

12 int MPI_File_iwrite_at(MPI_File fh, MPI_Offset offset, const void *buf,
13                       int count, MPI_Datatype datatype, MPI_Request *request)
14
15 int MPI_File_iwrite_at_c(MPI_File fh, MPI_Offset offset, const void *buf,
16                          MPI_Count count, MPI_Datatype datatype, MPI_Request *request)

```

Fortran 2008 binding

```

17 MPI_File_iwrite_at(fh, offset, buf, count, datatype, request, ierror)
18     TYPE(MPI_File), INTENT(IN) :: fh
19     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
20     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
21     INTEGER, INTENT(IN) :: count
22     TYPE(MPI_Datatype), INTENT(IN) :: datatype
23     TYPE(MPI_Request), INTENT(OUT) :: request
24     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
25
26 MPI_File_iwrite_at(fh, offset, buf, count, datatype, request, ierror) !(_c)
27     TYPE(MPI_File), INTENT(IN) :: fh
28     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
29     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
30     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
31     TYPE(MPI_Datatype), INTENT(IN) :: datatype
32     TYPE(MPI_Request), INTENT(OUT) :: request
33     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
34

```

Fortran binding

```

35 MPI_FILE_IWRITE_AT(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)
36     INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
37     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
38     <type> BUF(*)
39

```

MPI_FILE_IWRITE_AT is a nonblocking version of the MPI_FILE_WRITE_AT interface.

MPI_FILE_IWRITE_AT_ALL(fh, offset, buf, count, datatype, request)			1
INOUT	fh	file handle (handle)	2
IN	offset	file offset (integer)	3
IN	buf	initial address of buffer (choice)	4
IN	count	number of elements in buffer (integer)	5
IN	datatype	datatype of each buffer element (handle)	6
OUT	request	request object (handle)	7

C binding

```

int MPI_File_iwrite_at_all(MPI_File fh, MPI_Offset offset, const void *buf,
                           int count, MPI_Datatype datatype, MPI_Request *request)
int MPI_File_iwrite_at_all_c(MPI_File fh, MPI_Offset offset, const void *buf,
                             MPI_Count count, MPI_Datatype datatype, MPI_Request *request)

```

Fortran 2008 binding

```

MPI_File_iwrite_at_all(fh, offset, buf, count, datatype, request, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_iwrite_at_all(fh, offset, buf, count, datatype, request, ierror) !(_c)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_IWRITE_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)
  INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
  INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
  <type> BUF(*)

```

MPI_FILE_IWRITE_AT_ALL is a nonblocking version of MPI_FILE_WRITE_AT_ALL.

14.4.3 Data Access with Individual File Pointers

MPI maintains one individual file pointer per process per file handle. The current value of this pointer implicitly specifies the offset in the data access routines described in this section. These routines only use and update the individual file pointers maintained by MPI. The shared file pointer is not used nor updated.

The individual file pointer routines have the same semantics as the data access with explicit offset routines described in Section 14.4.2, with the following modification:

- the offset is defined to be the current value of the MPI-maintained individual file pointer.

After an individual file pointer operation is initiated, the individual file pointer is updated to point to the next etype after the last one that will be accessed. The file pointer is updated relative to the current view of the file.

If MPI_MODE_SEQUENTIAL mode was specified when the file was opened, it is erroneous to call the routines in this section, with the exception of MPI_FILE_GET_BYTE_OFFSET.

MPI_FILE_READ(fh, buf, count, datatype, status)

INOUT	fh	file handle (handle)
OUT	buf	initial address of buffer (choice)
IN	count	number of elements in buffer (integer)
IN	datatype	datatype of each buffer element (handle)
OUT	status	status object (status)

C binding

```
int MPI_File_read(MPI_File fh, void *buf, int count, MPI_Datatype datatype,
                 MPI_Status *status)
```

```
int MPI_File_read_c(MPI_File fh, void *buf, MPI_Count count,
                   MPI_Datatype datatype, MPI_Status *status)
```

Fortran 2008 binding

```
MPI_File_read(fh, buf, count, datatype, status, ierror)
```

```
TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..) :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

```
MPI_File_read(fh, buf, count, datatype, status, ierror) !(_c)
```

```
TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_READ(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
```

```
INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
<type> BUF(*)
```

MPI_FILE_READ reads a file using the individual file pointer.

Example 14.2. The following Fortran code fragment is an example of reading a file until the end of file is reached:

```

!   Read a preexisting input file until all data has been read.
!   Call routine "process_input" if all requested data is read.
!   The Fortran 90 "exit" statement exits the loop.

integer    bufsize, numread, totprocessed, status(MPI_STATUS_SIZE)
parameter (bufsize=100)
real      localbuffer(bufsize)
integer(kind=MPI_OFFSET_KIND) zero

zero = 0

call MPI_FILE_OPEN(MPI_COMM_WORLD, "myoldfile", &
                   MPI_MODE_RDONLY, MPI_INFO_NULL, myfh, ierr)
call MPI_FILE_SET_VIEW(myfh, zero, MPI_REAL, MPI_REAL, 'native', &
                      MPI_INFO_NULL, ierr)

totprocessed = 0
do
  call MPI_FILE_READ(myfh, localbuffer, bufsize, MPI_REAL, &
                    status, ierr)
  call MPI_GET_COUNT(status, MPI_REAL, numread, ierr)
  call process_input(localbuffer, numread)
  totprocessed = totprocessed + numread
  if (numread < bufsize) exit
end do

write(6, 1001) numread, bufsize, totprocessed
1001 format("No more data: read", I3, "and expected", I3, &
           "Processed total of", I6, "before terminating job.")

call MPI_FILE_CLOSE(myfh, ierr)

```

`MPI_FILE_READ_ALL(fh, buf, count, datatype, status)`

INOUT	fh	file handle (handle)
OUT	buf	initial address of buffer (choice)
IN	count	number of elements in buffer (integer)
IN	datatype	datatype of each buffer element (handle)
OUT	status	status object (status)

C binding

```
int MPI_File_read_all(MPI_File fh, void *buf, int count, MPI_Datatype datatype,
                    MPI_Status *status)
```

```
int MPI_File_read_all_c(MPI_File fh, void *buf, MPI_Count count,
                      MPI_Datatype datatype, MPI_Status *status)
```

Fortran 2008 binding

```

1 MPI_File_read_all(fh, buf, count, datatype, status, ierror)
2   TYPE(MPI_File), INTENT(IN) :: fh
3   TYPE(*), DIMENSION(..) :: buf
4   INTEGER, INTENT(IN) :: count
5   TYPE(MPI_Datatype), INTENT(IN) :: datatype
6   TYPE(MPI_Status) :: status
7   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8
9 MPI_File_read_all(fh, buf, count, datatype, status, ierror) !(_c)
10  TYPE(MPI_File), INTENT(IN) :: fh
11  TYPE(*), DIMENSION(..) :: buf
12  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
13  TYPE(MPI_Datatype), INTENT(IN) :: datatype
14  TYPE(MPI_Status) :: status
15  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
16

```

Fortran binding

```

17 MPI_FILE_READ_ALL(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
18   INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
19   <type> BUF(*)
20
21   MPI_FILE_READ_ALL is a collective version of the blocking MPI_FILE_READ interface.
22
23

```

```

24 MPI_FILE_WRITE(fh, buf, count, datatype, status)
25
26   INOUT   fh           file handle (handle)
27   IN      buf          initial address of buffer (choice)
28   IN      count        number of elements in buffer (integer)
29   IN      datatype     datatype of each buffer element (handle)
30   OUT     status       status object (status)
31
32

```

C binding

```

33
34 int MPI_File_write(MPI_File fh, const void *buf, int count,
35                   MPI_Datatype datatype, MPI_Status *status)
36
37 int MPI_File_write_c(MPI_File fh, const void *buf, MPI_Count count,
38                     MPI_Datatype datatype, MPI_Status *status)
39

```

Fortran 2008 binding

```

40 MPI_File_write(fh, buf, count, datatype, status, ierror)
41   TYPE(MPI_File), INTENT(IN) :: fh
42   TYPE(*), DIMENSION(..), INTENT(IN) :: buf
43   INTEGER, INTENT(IN) :: count
44   TYPE(MPI_Datatype), INTENT(IN) :: datatype
45   TYPE(MPI_Status) :: status
46   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
47
48 MPI_File_write(fh, buf, count, datatype, status, ierror) !(_c)
49   TYPE(MPI_File), INTENT(IN) :: fh

```

```

TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
  <type> BUF(*)

```

MPI_FILE_WRITE writes a file using the individual file pointer.

MPI_FILE_WRITE_ALL(fh, buf, count, datatype, status)

INOUT	fh	file handle (handle)	15
IN	buf	initial address of buffer (choice)	16
IN	count	number of elements in buffer (integer)	17
IN	datatype	datatype of each buffer element (handle)	18
OUT	status	status object (status)	19

C binding

```

int MPI_File_write_all(MPI_File fh, const void *buf, int count,
  MPI_Datatype datatype, MPI_Status *status)
int MPI_File_write_all_c(MPI_File fh, const void *buf, MPI_Count count,
  MPI_Datatype datatype, MPI_Status *status)

```

Fortran 2008 binding

```

MPI_File_write_all(fh, buf, count, datatype, status, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_all(fh, buf, count, datatype, status, ierror) !(_c)
  TYPE(MPI_File), INTENT(IN) :: fh
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE_ALL(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
  <type> BUF(*)

```

1 MPI_FILE_WRITE_ALL is a collective version of the blocking MPI_FILE_WRITE inter-
2 face.

3
4
5 MPI_FILE_IREAD(fh, buf, count, datatype, request)

6 INOUT fh file handle (handle)
7 OUT buf initial address of buffer (choice)
8 IN count number of elements in buffer (integer)
9 IN datatype datatype of each buffer element (handle)
10 OUT request request object (handle)
11
12

13 C binding

14 int MPI_File_iread(MPI_File fh, void *buf, int count, MPI_Datatype datatype,
15 MPI_Request *request)

16
17 int MPI_File_iread_c(MPI_File fh, void *buf, MPI_Count count,
18 MPI_Datatype datatype, MPI_Request *request)

19 Fortran 2008 binding

20 MPI_File_iread(fh, buf, count, datatype, request, ierror)

21 TYPE(MPI_File), INTENT(IN) :: fh
22 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
23 INTEGER, INTENT(IN) :: count
24 TYPE(MPI_Datatype), INTENT(IN) :: datatype
25 TYPE(MPI_Request), INTENT(OUT) :: request
26 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

27 MPI_File_iread(fh, buf, count, datatype, request, ierror) !(_c)

28 TYPE(MPI_File), INTENT(IN) :: fh
29 TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
30 INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
31 TYPE(MPI_Datatype), INTENT(IN) :: datatype
32 TYPE(MPI_Request), INTENT(OUT) :: request
33 INTEGER, OPTIONAL, INTENT(OUT) :: ierror
34

35 Fortran binding

36 MPI_FILE_IREAD(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)

37 INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
38 <type> BUF(*)

39 MPI_FILE_IREAD is a nonblocking version of the MPI_FILE_READ interface.

40
41 **Example 14.3.** The following Fortran code fragment illustrates file pointer update
42 semantics:

```
43 ! Read the first twenty real words in a file into two local
44 ! buffers. Note that when the first MPI_FILE_IREAD returns,
45 ! the file pointer has been updated to point to the
46 ! eleventh real word in the file.
```

```
47 integer bufsize, req1, req2
48
```

```

1 integer, dimension(MPI_STATUS_SIZE) :: status1, status2
2 parameter (bufsize=10)
3 real      buf1(bufsize), buf2(bufsize)
4 integer(kind=MPI_OFFSET_KIND) zero
5
6 zero = 0
7 call MPI_FILE_OPEN(MPI_COMM_WORLD, 'myoldfile', &
8                   MPI_MODE_RDONLY, MPI_INFO_NULL, myfh, ierr)
9 call MPI_FILE_SET_VIEW(myfh, zero, MPI_REAL, MPI_REAL, 'native', &
10                      MPI_INFO_NULL, ierr)
11 call MPI_FILE_IREAD(myfh, buf1, bufsize, MPI_REAL, &
12                   req1, ierr)
13 call MPI_FILE_IREAD(myfh, buf2, bufsize, MPI_REAL, &
14                   req2, ierr)
15
16 call MPI_WAIT(req1, status1, ierr)
17 call MPI_WAIT(req2, status2, ierr)
18
19 call MPI_FILE_CLOSE(myfh, ierr)

```

MPI_FILE_IREAD_ALL(fh, buf, count, datatype, request)

INOUT	fh	file handle (handle)
OUT	buf	initial address of buffer (choice)
IN	count	number of elements in buffer (integer)
IN	datatype	datatype of each buffer element (handle)
OUT	request	request object (handle)

C binding

```
int MPI_File_iread_all(MPI_File fh, void *buf, int count,
MPI_Datatype datatype, MPI_Request *request)
```

```
int MPI_File_iread_all_c(MPI_File fh, void *buf, MPI_Count count,
MPI_Datatype datatype, MPI_Request *request)
```

Fortran 2008 binding

```
MPI_File_iread_all(fh, buf, count, datatype, request, ierror)
```

```

TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```
MPI_File_iread_all(fh, buf, count, datatype, request, ierror) !(_c)
```

```

TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype

```

```

1     TYPE(MPI_Request), INTENT(OUT) :: request
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

4 MPI_FILE_IREAD_ALL(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)
5     INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
6     <type> BUF(*)

```

8 MPI_FILE_IREAD_ALL is a nonblocking version of MPI_FILE_READ_ALL.

```

10 MPI_FILE_IWRITE(fh, buf, count, datatype, request)

```

12	INOUT	fh	file handle (handle)
13	IN	buf	initial address of buffer (choice)
14	IN	count	number of elements in buffer (integer)
15	IN	datatype	datatype of each buffer element (handle)
16	IN	datatype	datatype of each buffer element (handle)
17	OUT	request	request object (handle)

C binding

```

20 int MPI_File_irewrite(MPI_File fh, const void *buf, int count,
21                      MPI_Datatype datatype, MPI_Request *request)
22
23 int MPI_File_irewrite_c(MPI_File fh, const void *buf, MPI_Count count,
24                        MPI_Datatype datatype, MPI_Request *request)

```

Fortran 2008 binding

```

26 MPI_File_irewrite(fh, buf, count, datatype, request, ierror)
27     TYPE(MPI_File), INTENT(IN) :: fh
28     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
29     INTEGER, INTENT(IN) :: count
30     TYPE(MPI_Datatype), INTENT(IN) :: datatype
31     TYPE(MPI_Request), INTENT(OUT) :: request
32     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
33
34 MPI_File_irewrite(fh, buf, count, datatype, request, ierror) !(_c)
35     TYPE(MPI_File), INTENT(IN) :: fh
36     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
37     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
38     TYPE(MPI_Datatype), INTENT(IN) :: datatype
39     TYPE(MPI_Request), INTENT(OUT) :: request
40     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

42 MPI_FILE_IWRITE(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)
43     INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
44     <type> BUF(*)

```

45 MPI_FILE_IWRITE is a nonblocking version of MPI_FILE_WRITE.

48

MPI_FILE_IWRITE_ALL(fh, buf, count, datatype, request)	1
INOUT fh file handle (handle)	2
IN buf initial address of buffer (choice)	3
IN count number of elements in buffer (integer)	4
IN datatype datatype of each buffer element (handle)	5
OUT request request object (handle)	6
	7
	8
	9
C binding	10
int MPI_File_ fwrite_all(MPI_File fh, const void *buf, int count,	11
MPI_Datatype datatype, MPI_Request *request)	12
int MPI_File_ fwrite_all_c(MPI_File fh, const void *buf, MPI_Count count,	13
MPI_Datatype datatype, MPI_Request *request)	14
	15
Fortran 2008 binding	16
MPI_File_ fwrite_all(fh, buf, count, datatype, request, ierror)	17
TYPE(MPI_File), INTENT(IN) :: fh	18
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf	19
INTEGER, INTENT(IN) :: count	20
TYPE(MPI_Datatype), INTENT(IN) :: datatype	21
TYPE(MPI_Request), INTENT(OUT) :: request	22
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	23
MPI_File_ fwrite_all(fh, buf, count, datatype, request, ierror) !(_c)	24
TYPE(MPI_File), INTENT(IN) :: fh	25
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf	26
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count	27
TYPE(MPI_Datatype), INTENT(IN) :: datatype	28
TYPE(MPI_Request), INTENT(OUT) :: request	29
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	30
	31
Fortran binding	32
MPI_FILE_IWRITE_ALL(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)	33
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	34
<type> BUF(*)	35
MPI_FILE_IWRITE_ALL is a nonblocking version of MPI_FILE_WRITE_ALL.	36
	37
	38
MPI_FILE_SEEK(fh, offset, whence)	39
INOUT fh file handle (handle)	40
IN offset file offset (integer)	41
IN whence update mode (state)	42
	43
	44
C binding	45
int MPI_File_seek(MPI_File fh, MPI_Offset offset, int whence)	46
	47
Fortran 2008 binding	48
MPI_File_seek(fh, offset, whence, ierror)	48

```

1      TYPE(MPI_File), INTENT(IN) :: fh
2      INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
3      INTEGER, INTENT(IN) :: whence
4      INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

6      MPI_FILE_SEEK(FH, OFFSET, WHENCE, IERROR)
7      INTEGER FH, WHENCE, IERROR
8      INTEGER(KIND=MPI_OFFSET_KIND) OFFSET

```

10 MPI_FILE_SEEK updates the individual file pointer according to whence, which has the
11 following possible values:

12	MPI_SEEK_SET	the pointer is set to <i>offset</i>
13	MPI_SEEK_CUR	the pointer is set to the current pointer position 14 plus <i>offset</i>
15	MPI_SEEK_END	the pointer is set to the end of file plus <i>offset</i>

16 The *offset* can be negative, which allows seeking backwards. It is erroneous to seek to
17 a negative position in the view.

```

20 MPI_FILE_GET_POSITION(fh, offset)

```

21	IN	fh	file handle (handle)
22	OUT	offset	offset of individual pointer (integer)

C binding

```

26 int MPI_File_get_position(MPI_File fh, MPI_Offset *offset)

```

Fortran 2008 binding

```

28 MPI_File_get_position(fh, offset, ierror)
29     TYPE(MPI_File), INTENT(IN) :: fh
30     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: offset
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

34 MPI_FILE_GET_POSITION(FH, OFFSET, IERROR)
35     INTEGER FH, IERROR
36     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET

```

37 MPI_FILE_GET_POSITION returns, in *offset*, the current position of the individual file
38 pointer in etype units relative to the current view.

40 *Advice to users.* The *offset* can be used in a future call to MPI_FILE_SEEK using
41 *whence* = MPI_SEEK_SET to return to the current position. To set the displacement to
42 the current file pointer position, first convert *offset* into an absolute byte position using
43 MPI_FILE_GET_BYTE_OFFSET, then call MPI_FILE_SET_VIEW with the resulting
44 displacement. (*End of advice to users.*)

MPI_FILE_GET_BYTE_OFFSET(fh, offset, disp)			1
IN	fh	file handle (handle)	2
IN	offset	offset (integer)	3
OUT	disp	absolute byte position of offset (integer)	4
			5
			6

C binding

```
int MPI_File_get_byte_offset(MPI_File fh, MPI_Offset offset, MPI_Offset *disp)
```

Fortran 2008 binding

```
MPI_File_get_byte_offset(fh, offset, disp, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: disp
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_GET_BYTE_OFFSET(FH, OFFSET, DISP, IERROR)
  INTEGER FH, IERROR
  INTEGER(KIND=MPI_OFFSET_KIND) OFFSET, DISP
```

MPI_FILE_GET_BYTE_OFFSET converts a view-relative offset into an absolute byte position. The absolute byte position (from the beginning of the file) of `offset` relative to the current view of `fh` is returned in `disp`.

14.4.4 Data Access with Shared File Pointers

MPI maintains exactly one shared file pointer per collective MPI_FILE_OPEN (shared among processes in the communicator group). The current value of this pointer implicitly specifies the offset in the data access routines described in this section. These routines only use and update the shared file pointer maintained by MPI. The individual file pointers are not used nor updated.

The shared file pointer routines have the same semantics as the data access with explicit offset routines described in Section 14.4.2, with the following modifications:

- the `offset` is defined to be the current value of the MPI-maintained shared file pointer,
- the effect of multiple calls to shared file pointer routines is defined to behave as if the calls were serialized, and
- the use of shared file pointer routines is erroneous unless all processes use the same file view.

For the noncollective shared file pointer routines, the serialization ordering is not deterministic. The user needs to use other synchronization means to enforce a specific order.

After a shared file pointer operation is initiated, the shared file pointer is updated to point to the next etype after the last one that will be accessed. The file pointer is updated relative to the current view of the file.

1 *Noncollective Operations*

2
3
4 **MPI_FILE_READ_SHARED**(fh, buf, count, datatype, status)
5 INOUT fh file handle (handle)
6 OUT buf initial address of buffer (choice)
7 IN count number of elements in buffer (integer)
8 IN datatype datatype of each buffer element (handle)
9 OUT status status object (status)
10
11

12
13 **C binding**

14 `int MPI_File_read_shared(MPI_File fh, void *buf, int count,`
15 `MPI_Datatype datatype, MPI_Status *status)`
16
17 `int MPI_File_read_shared_c(MPI_File fh, void *buf, MPI_Count count,`
18 `MPI_Datatype datatype, MPI_Status *status)`

19 **Fortran 2008 binding**

20 `MPI_File_read_shared(fh, buf, count, datatype, status, ierror)`
21 `TYPE(MPI_File), INTENT(IN) :: fh`
22 `TYPE(*), DIMENSION(..) :: buf`
23 `INTEGER, INTENT(IN) :: count`
24 `TYPE(MPI_Datatype), INTENT(IN) :: datatype`
25 `TYPE(MPI_Status) :: status`
26 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`
27
28 `MPI_File_read_shared(fh, buf, count, datatype, status, ierror) !(_c)`
29 `TYPE(MPI_File), INTENT(IN) :: fh`
30 `TYPE(*), DIMENSION(..) :: buf`
31 `INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count`
32 `TYPE(MPI_Datatype), INTENT(IN) :: datatype`
33 `TYPE(MPI_Status) :: status`
34 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

35 **Fortran binding**

36 `MPI_FILE_READ_SHARED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)`
37 `INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR`
38 `<type> BUF(*)`
39
40 `MPI_FILE_READ_SHARED reads a file using the shared file pointer.`
41
42
43
44
45
46
47
48

MPI_FILE_WRITE_SHARED(fh, buf, count, datatype, status)			1
INOUT	fh	file handle (handle)	2
IN	buf	initial address of buffer (choice)	3
IN	count	number of elements in buffer (integer)	4
IN	datatype	datatype of each buffer element (handle)	5
OUT	status	status object (status)	6

C binding

```

int MPI_File_write_shared(MPI_File fh, const void *buf, int count,
                          MPI_Datatype datatype, MPI_Status *status)
int MPI_File_write_shared_c(MPI_File fh, const void *buf, MPI_Count count,
                            MPI_Datatype datatype, MPI_Status *status)

```

Fortran 2008 binding

```

MPI_File_write_shared(fh, buf, count, datatype, status, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_shared(fh, buf, count, datatype, status, ierror) !(_c)
  TYPE(MPI_File), INTENT(IN) :: fh
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE_SHARED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
  <type> BUF(*)
MPI_FILE_WRITE_SHARED writes a file using the shared file pointer.

```

MPI_FILE_IREAD_SHARED(fh, buf, count, datatype, request)

INOUT	fh	file handle (handle)	39
OUT	buf	initial address of buffer (choice)	40
IN	count	number of elements in buffer (integer)	41
IN	datatype	datatype of each buffer element (handle)	42
OUT	request	request object (handle)	43

C binding

```

int MPI_File_iread_shared(MPI_File fh, void *buf, int count,

```

```

1         MPI_Datatype datatype, MPI_Request *request)
2
3     int MPI_File_iread_shared_c(MPI_File fh, void *buf, MPI_Count count,
4         MPI_Datatype datatype, MPI_Request *request)
5
6 Fortran 2008 binding
7 MPI_File_iread_shared(fh, buf, count, datatype, request, ierror)
8     TYPE(MPI_File), INTENT(IN) :: fh
9     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
10    INTEGER, INTENT(IN) :: count
11    TYPE(MPI_Datatype), INTENT(IN) :: datatype
12    TYPE(MPI_Request), INTENT(OUT) :: request
13    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14
15 MPI_File_iread_shared(fh, buf, count, datatype, request, ierror) !(_c)
16    TYPE(MPI_File), INTENT(IN) :: fh
17    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
18    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
19    TYPE(MPI_Datatype), INTENT(IN) :: datatype
20    TYPE(MPI_Request), INTENT(OUT) :: request
21    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22
23 Fortran binding
24 MPI_FILE_IREAD_SHARED(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)
25     INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR
26     <type> BUF(*)
27
28     MPI_FILE_IREAD_SHARED is a nonblocking version of MPI_FILE_READ_SHARED.
29
30 MPI_FILE_IWRITE_SHARED(fh, buf, count, datatype, request)
31
32     INOUT   fh                file handle (handle)
33     IN      buf                initial address of buffer (choice)
34     IN      count              number of elements in buffer (integer)
35     IN      datatype           datatype of each buffer element (handle)
36     OUT     request            request object (handle)
37
38 C binding
39 int MPI_File_iwrite_shared(MPI_File fh, const void *buf, int count,
40     MPI_Datatype datatype, MPI_Request *request)
41
42 int MPI_File_iwrite_shared_c(MPI_File fh, const void *buf, MPI_Count count,
43     MPI_Datatype datatype, MPI_Request *request)
44
45 Fortran 2008 binding
46 MPI_File_iwrite_shared(fh, buf, count, datatype, request, ierror)
47     TYPE(MPI_File), INTENT(IN) :: fh
48     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
49     INTEGER, INTENT(IN) :: count
50     TYPE(MPI_Datatype), INTENT(IN) :: datatype

```

```

TYPE(MPI_Request), INTENT(OUT) :: request      1
INTEGER, OPTIONAL, INTENT(OUT) :: ierror      2
MPI_File_iwrite_shared(fh, buf, count, datatype, request, ierror) !(_c)      3
TYPE(MPI_File), INTENT(IN) :: fh              4
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf      5
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count      6
TYPE(MPI_Datatype), INTENT(IN) :: datatype      7
TYPE(MPI_Request), INTENT(OUT) :: request      8
INTEGER, OPTIONAL, INTENT(OUT) :: ierror      9

```

Fortran binding

```

MPI_FILE_IWRITE_SHARED(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)      11
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR      12
<type> BUF(*)      13

```

MPI_FILE_IWRITE_SHARED is a nonblocking version of the MPI_FILE_WRITE_SHARED interface. 14 15 16

Collective Operations

The semantics of a collective access using a shared file pointer are that the accesses to the file will be in the order determined by the ranks of the processes within the group. For each process, the location in the file at which data is accessed is the position at which the shared file pointer would be after all processes whose ranks within the group less than that of this process had accessed their data. In addition, in order to prevent subsequent shared offset accesses by the same processes from interfering with this collective access, the call might return only after all the processes within the group have initiated their accesses. When the call returns, the shared file pointer points to the next etype accessible, according to the file view used by all processes, after the last etype requested. 17 18 19 20 21 22 23 24 25 26 27 28

Advice to users. There may be some programs in which all processes in the group need to access the file using the shared file pointer, but the program may not *require* that data be accessed in order of process rank. In such programs, using the shared ordered routines (e.g., MPI_FILE_WRITE_ORDERED rather than MPI_FILE_WRITE_SHARED) may enable an implementation to optimize access, improving performance. (*End of advice to users.*) 29 30 31 32 33 34 35 36

Advice to implementors. Accesses to the data requested by all processes do not have to be serialized. Once all processes have issued their requests, locations within the file for all accesses can be computed, and accesses can proceed independently from each other, possibly in parallel. (*End of advice to implementors.*) 37 38 39 40 41 42 43 44 45 46 47 48

```

1 MPI_FILE_READ_ORDERED(fh, buf, count, datatype, status)
2     INOUT   fh                file handle (handle)
3
4     OUT     buf                initial address of buffer (choice)
5
6     IN      count              number of elements in buffer (integer)
7
8     IN      datatype           datatype of each buffer element (handle)
9
10    OUT     status              status object (status)

```

10 C binding

```

11 int MPI_File_read_ordered(MPI_File fh, void *buf, int count,
12     MPI_Datatype datatype, MPI_Status *status)
13
14 int MPI_File_read_ordered_c(MPI_File fh, void *buf, MPI_Count count,
15     MPI_Datatype datatype, MPI_Status *status)

```

16 Fortran 2008 binding

```

17 MPI_File_read_ordered(fh, buf, count, datatype, status, ierror)
18     TYPE(MPI_File), INTENT(IN) :: fh
19     TYPE(*), DIMENSION(..) :: buf
20     INTEGER, INTENT(IN) :: count
21     TYPE(MPI_Datatype), INTENT(IN) :: datatype
22     TYPE(MPI_Status) :: status
23     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
24
25 MPI_File_read_ordered(fh, buf, count, datatype, status, ierror) !(_c)
26     TYPE(MPI_File), INTENT(IN) :: fh
27     TYPE(*), DIMENSION(..) :: buf
28     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
29     TYPE(MPI_Datatype), INTENT(IN) :: datatype
30     TYPE(MPI_Status) :: status
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

32 Fortran binding

```

33 MPI_FILE_READ_ORDERED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
34     INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
35     <type> BUF(*)

```

36 MPI_FILE_READ_ORDERED is a collective version of the MPI_FILE_READ_SHARED
37 interface.

```

40 MPI_FILE_WRITE_ORDERED(fh, buf, count, datatype, status)

```

```

41     INOUT   fh                file handle (handle)
42
43     IN      buf                initial address of buffer (choice)
44
45     IN      count              number of elements in buffer (integer)
46
47     IN      datatype           datatype of each buffer element (handle)
48
49     OUT     status              status object (status)

```


C binding

```

int MPI_File_write_ordered(MPI_File fh, const void *buf, int count,
                           MPI_Datatype datatype, MPI_Status *status)
int MPI_File_write_ordered_c(MPI_File fh, const void *buf, MPI_Count count,
                             MPI_Datatype datatype, MPI_Status *status)

```

Fortran 2008 binding

```

MPI_File_write_ordered(fh, buf, count, datatype, status, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_ordered(fh, buf, count, datatype, status, ierror) !(_c)
  TYPE(MPI_File), INTENT(IN) :: fh
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Status) :: status
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE_ORDERED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
  <type> BUF(*)
  MPI_FILE_WRITE_ORDERED is a collective version of the MPI_FILE_WRITE_SHARED
  interface.

```

Seek

If MPI_MODE_SEQUENTIAL mode was specified when the file was opened, it is erroneous to call the following two routines (MPI_FILE_SEEK_SHARED and MPI_FILE_GET_POSITION_SHARED).

```

MPI_FILE_SEEK_SHARED(fh, offset, whence)
  INOUT fh file handle (handle)
  IN offset file offset (integer)
  IN whence update mode (state)

```

C binding

```

int MPI_File_seek_shared(MPI_File fh, MPI_Offset offset, int whence)

```

Fortran 2008 binding

```

MPI_File_seek_shared(fh, offset, whence, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset

```

```

1      INTEGER, INTENT(IN) :: whence
2      INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

4      MPI_FILE_SEEK_SHARED(FH, OFFSET, WHENCE, IERROR)
5      INTEGER FH, WHENCE, IERROR
6      INTEGER(KIND=MPI_OFFSET_KIND) OFFSET

```

8 `MPI_FILE_SEEK_SHARED` updates the shared file pointer according to `whence`, which has the following possible values:

10	<code>MPI_SEEK_SET</code>	the pointer is set to <code>offset</code>
11	<code>MPI_SEEK_CUR</code>	the pointer is set to the current pointer position plus <code>offset</code>
12		
13	<code>MPI_SEEK_END</code>	the pointer is set to the end of file plus <code>offset</code>

14 `MPI_FILE_SEEK_SHARED` is collective; all the processes in the communicator group associated with the file handle `fh` must call `MPI_FILE_SEEK_SHARED` with the same values for `offset` and `whence`.

17 The `offset` can be negative, which allows seeking backwards. It is erroneous to seek to a negative position in the view.

```

21      MPI_FILE_GET_POSITION_SHARED(fh, offset)

```

22	IN	<code>fh</code>	file handle (handle)
23			
24	OUT	<code>offset</code>	offset of shared pointer (integer)

C binding

```

27      int MPI_File_get_position_shared(MPI_File fh, MPI_Offset *offset)

```

Fortran 2008 binding

```

29      MPI_File_get_position_shared(fh, offset, ierror)
30      TYPE(MPI_File), INTENT(IN) :: fh
31      INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: offset
32      INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

35      MPI_FILE_GET_POSITION_SHARED(FH, OFFSET, IERROR)
36      INTEGER FH, IERROR
37      INTEGER(KIND=MPI_OFFSET_KIND) OFFSET

```

38 `MPI_FILE_GET_POSITION_SHARED` returns, in `offset`, the current position of the shared file pointer in `etype` units relative to the current view.

41 *Advice to users.* The `offset` can be used in a future call to `MPI_FILE_SEEK_SHARED` using `whence = MPI_SEEK_SET` to return to the current position. To set the displacement to the current file pointer position, first convert `offset` into an absolute byte position using `MPI_FILE_GET_BYTE_OFFSET`, then call `MPI_FILE_SET_VIEW` with the resulting displacement. (*End of advice to users.*)

14.4.5 Split Collective Data Access Routines

MPI provides a restricted form of “nonblocking collective” I/O operations for all data accesses using split collective data access routines. These routines are referred to as “split” collective routines, because a single collective operation is split in two: a begin routine and an end routine. The begin routine begins the operation, much like a nonblocking data access (e.g., `MPI_FILE_IREAD`). The end routine completes the operation, much like the matching test or wait (e.g., `MPI_WAIT`). As with nonblocking data access operations, the user must not use the buffer passed to a begin routine while the routine is outstanding; the operation must be completed with an end routine before it is safe to free buffers, etc.

Split collective data access operations on a file handle `fh` are subject to the semantic rules given below.

- On any MPI process, each file handle may have at most one active split collective operation at any time.
- Begin calls are collective over the group of processes that participated in the collective open and follow the ordering rules for collective calls.
- End calls are collective over the group of processes that participated in the collective open and follow the ordering rules for collective calls. Each end call matches the preceding begin call for the same collective operation. When an “end” call is made, exactly one unmatched “begin” call for the same operation must precede it.
- An implementation is free to implement any split collective data access routine using the corresponding blocking collective routine when either the begin call (e.g., `MPI_FILE_READ_ALL_BEGIN`) or the end call (e.g., `MPI_FILE_READ_ALL_END`) is issued. The begin and end calls are provided to allow the user and MPI implementation to optimize the collective operation.

According to the definitions in Section 2.4.2, the begin procedures are incomplete. They are also nonlocal procedures because they may or may not return before they are called in all MPI processes of the process group.

Advice to users. This is one of the exceptions in which incomplete procedures are nonlocal and therefore blocking. (*End of advice to users.*)

- Split collective operations do not match the corresponding regular collective operation. For example, in a single collective read operation, an `MPI_FILE_READ_ALL` on one process does not match an `MPI_FILE_READ_ALL_BEGIN/MPI_FILE_READ_ALL_END` pair on another process.
- Split collective routines must specify a buffer in both the begin and end routines. By specifying the buffer that receives data in the end routine, we can avoid the problems described in [Problems with Code Movement and Register Optimization](#), Section 19.1.17, but not all of the problems, such as those described in Sections 19.1.12, 19.1.13, and 19.1.16.
- No collective I/O operations are permitted on a file handle concurrently with a split collective access on that file handle (i.e., between the begin and end of the access). That is, the following example is erroneous.

Example 14.4. Erroneous example fragment of concurrent split collective access on a file handle:

```
MPI_File_read_all_begin(fh, ...);
...
MPI_File_read_all(fh, ...);
...
MPI_File_read_all_end(fh, ...);
```

- In a multithreaded implementation, any split collective begin and end operation called by a process must be called from the same thread. This restriction is made to simplify the implementation in the multithreaded case. (Note that we have already disallowed having two threads begin a split collective operation on the same file handle since only one split collective operation can be active on a file handle at any time.)

The arguments for these routines have the same meaning as for the equivalent collective versions (e.g., the argument definitions for `MPI_FILE_READ_ALL_BEGIN` and `MPI_FILE_READ_ALL_END` are equivalent to the arguments for `MPI_FILE_READ_ALL`). The begin routine (e.g., `MPI_FILE_READ_ALL_BEGIN`) begins a split collective operation that, when completed with the matching end routine (i.e., `MPI_FILE_READ_ALL_END`) produces the result as defined for the equivalent collective routine (i.e., `MPI_FILE_READ_ALL`).

For the purpose of consistency semantics (Section 14.6.1), a matched pair of split collective data access operations (e.g., `MPI_FILE_READ_ALL_BEGIN` and `MPI_FILE_READ_ALL_END`) compose a single data access.

`MPI_FILE_READ_AT_ALL_BEGIN(fh, offset, buf, count, datatype)`

IN	fh	file handle (handle)
IN	offset	file offset (integer)
OUT	buf	initial address of buffer (choice)
IN	count	number of elements in buffer (integer)
IN	datatype	datatype of each buffer element (handle)

C binding

```
int MPI_File_read_at_all_begin(MPI_File fh, MPI_Offset offset, void *buf,
                               int count, MPI_Datatype datatype)
```

```
int MPI_File_read_at_all_begin_c(MPI_File fh, MPI_Offset offset, void *buf,
                                 MPI_Count count, MPI_Datatype datatype)
```

Fortran 2008 binding

```
MPI_File_read_at_all_begin(fh, offset, buf, count, datatype, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
  INTEGER, INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
```

```

    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_read_at_all_begin(fh, offset, buf, count, datatype, ierror) !(_c)
    TYPE(MPI_File), INTENT(IN) :: fh
    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
    TYPE(MPI_Datatype), INTENT(IN) :: datatype
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding
MPI_FILE_READ_AT_ALL_BEGIN(FH, OFFSET, BUF, COUNT, DATATYPE, IERROR)
    INTEGER FH, COUNT, DATATYPE, IERROR
    INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
    <type> BUF(*)

MPI_FILE_READ_AT_ALL_END(fh, buf, status)

    IN      fh                file handle (handle)
    OUT     buf                initial address of buffer (choice)
    OUT     status             status object (status)

C binding
int MPI_File_read_at_all_end(MPI_File fh, void *buf, MPI_Status *status)

Fortran 2008 binding
MPI_File_read_at_all_end(fh, buf, status, ierror)
    TYPE(MPI_File), INTENT(IN) :: fh
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
    TYPE(MPI_Status) :: status
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

Fortran binding
MPI_FILE_READ_AT_ALL_END(FH, BUF, STATUS, IERROR)
    INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
    <type> BUF(*)

MPI_FILE_WRITE_AT_ALL_BEGIN(fh, offset, buf, count, datatype)

    INOUT   fh                file handle (handle)
    IN      offset             file offset (integer)
    IN      buf                initial address of buffer (choice)
    IN      count              number of elements in buffer (integer)
    IN      datatype           datatype of each buffer element (handle)

C binding
int MPI_File_write_at_all_begin(MPI_File fh, MPI_Offset offset,
    const void *buf, int count, MPI_Datatype datatype)

```

```

1  int MPI_File_write_at_all_begin_c(MPI_File fh, MPI_Offset offset,
2      const void *buf, MPI_Count count, MPI_Datatype datatype)
3
4  Fortran 2008 binding
5  MPI_File_write_at_all_begin(fh, offset, buf, count, datatype, ierror)
6      TYPE(MPI_File), INTENT(IN) :: fh
7      INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
8      TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
9      INTEGER, INTENT(IN) :: count
10     TYPE(MPI_Datatype), INTENT(IN) :: datatype
11     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_File_write_at_all_begin(fh, offset, buf, count, datatype, ierror) !(_c)
14     TYPE(MPI_File), INTENT(IN) :: fh
15     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
16     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
17     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
18     TYPE(MPI_Datatype), INTENT(IN) :: datatype
19     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21 Fortran binding
22 MPI_FILE_WRITE_AT_ALL_BEGIN(FH, OFFSET, BUF, COUNT, DATATYPE, IERROR)
23     INTEGER FH, COUNT, DATATYPE, IERROR
24     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
25     <type> BUF(*)
26
27 MPI_FILE_WRITE_AT_ALL_END(fh, buf, status)
28     INOUT   fh           file handle (handle)
29     IN      buf          initial address of buffer (choice)
30     OUT     status       status object (status)
31
32
33 C binding
34 int MPI_File_write_at_all_end(MPI_File fh, const void *buf, MPI_Status *status)
35
36 Fortran 2008 binding
37 MPI_File_write_at_all_end(fh, buf, status, ierror)
38     TYPE(MPI_File), INTENT(IN) :: fh
39     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
40     TYPE(MPI_Status) :: status
41     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 Fortran binding
44 MPI_FILE_WRITE_AT_ALL_END(FH, BUF, STATUS, IERROR)
45     INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
46     <type> BUF(*)
47
48

```

MPI_FILE_READ_ALL_BEGIN(fh, buf, count, datatype)	1
INOUT fh file handle (handle)	2
OUT buf initial address of buffer (choice)	3
IN count number of elements in buffer (integer)	4
IN datatype datatype of each buffer element (handle)	5
	6
	7
	8
C binding	9
int MPI_File_read_all_begin(MPI_File fh, void *buf, int count,	10
MPI_Datatype datatype)	11
int MPI_File_read_all_begin_c(MPI_File fh, void *buf, MPI_Count count,	12
MPI_Datatype datatype)	13
	14
Fortran 2008 binding	15
MPI_File_read_all_begin(fh, buf, count, datatype, ierror)	16
TYPE(MPI_File), INTENT(IN) :: fh	17
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf	18
INTEGER, INTENT(IN) :: count	19
TYPE(MPI_Datatype), INTENT(IN) :: datatype	20
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	21
MPI_File_read_all_begin(fh, buf, count, datatype, ierror) !(_c)	22
TYPE(MPI_File), INTENT(IN) :: fh	23
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf	24
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count	25
TYPE(MPI_Datatype), INTENT(IN) :: datatype	26
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	27
	28
Fortran binding	29
MPI_FILE_READ_ALL_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)	30
INTEGER FH, COUNT, DATATYPE, IERROR	31
<type> BUF(*)	32
	33
	34
MPI_FILE_READ_ALL_END(fh, buf, status)	35
INOUT fh file handle (handle)	36
OUT buf initial address of buffer (choice)	37
OUT status status object (status)	38
	39
	40
C binding	41
int MPI_File_read_all_end(MPI_File fh, void *buf, MPI_Status *status)	42
	43
Fortran 2008 binding	44
MPI_File_read_all_end(fh, buf, status, ierror)	45
TYPE(MPI_File), INTENT(IN) :: fh	46
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf	47
TYPE(MPI_Status) :: status	48
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	49

```

1 Fortran binding
2 MPI_FILE_READ_ALL_END(FH, BUF, STATUS, IERROR)
3     INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
4     <type> BUF(*)
5
6
7 MPI_FILE_WRITE_ALL_BEGIN(fh, buf, count, datatype)
8
9     INOUT    fh                file handle (handle)
10    IN       buf                initial address of buffer (choice)
11    IN       count              number of elements in buffer (integer)
12    IN       datatype           datatype of each buffer element (handle)
13
14
15 C binding
16 int MPI_File_write_all_begin(MPI_File fh, const void *buf, int count,
17     MPI_Datatype datatype)
18
19 int MPI_File_write_all_begin_c(MPI_File fh, const void *buf, MPI_Count count,
20     MPI_Datatype datatype)
21
22 Fortran 2008 binding
23 MPI_File_write_all_begin(fh, buf, count, datatype, ierror)
24     TYPE(MPI_File), INTENT(IN) :: fh
25     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
26     INTEGER, INTENT(IN) :: count
27     TYPE(MPI_Datatype), INTENT(IN) :: datatype
28     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
29
30 MPI_File_write_all_begin(fh, buf, count, datatype, ierror) !(_c)
31     TYPE(MPI_File), INTENT(IN) :: fh
32     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
33     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
34     TYPE(MPI_Datatype), INTENT(IN) :: datatype
35     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
37 Fortran binding
38 MPI_FILE_WRITE_ALL_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)
39     INTEGER FH, COUNT, DATATYPE, IERROR
40     <type> BUF(*)
41
42
43 MPI_FILE_WRITE_ALL_END(fh, buf, status)
44
45     INOUT    fh                file handle (handle)
46    IN       buf                initial address of buffer (choice)
47    OUT      status              status object (status)
48
49
50 C binding
51 int MPI_File_write_all_end(MPI_File fh, const void *buf, MPI_Status *status)

```


Fortran 2008 binding

```

MPI_File_write_all_end(fh, buf, status, ierror)
    TYPE(MPI_File), INTENT(IN) :: fh
    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
    TYPE(MPI_Status) :: status
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE_ALL_END(FH, BUF, STATUS, IERROR)
    INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
    <type> BUF(*)

```

```

MPI_FILE_READ_ORDERED_BEGIN(fh, buf, count, datatype)

```

INOUT	fh	file handle (handle)
OUT	buf	initial address of buffer (choice)
IN	count	number of elements in buffer (integer)
IN	datatype	datatype of each buffer element (handle)

C binding

```

int MPI_File_read_ordered_begin(MPI_File fh, void *buf, int count,
    MPI_Datatype datatype)
int MPI_File_read_ordered_begin_c(MPI_File fh, void *buf, MPI_Count count,
    MPI_Datatype datatype)

```

Fortran 2008 binding

```

MPI_File_read_ordered_begin(fh, buf, count, datatype, ierror)
    TYPE(MPI_File), INTENT(IN) :: fh
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
    INTEGER, INTENT(IN) :: count
    TYPE(MPI_Datatype), INTENT(IN) :: datatype
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_File_read_ordered_begin(fh, buf, count, datatype, ierror) !(_c)
    TYPE(MPI_File), INTENT(IN) :: fh
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
    TYPE(MPI_Datatype), INTENT(IN) :: datatype
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_READ_ORDERED_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)
    INTEGER FH, COUNT, DATATYPE, IERROR
    <type> BUF(*)

```

```

1 MPI_FILE_READ_ORDERED_END(fh, buf, status)
2     INOUT   fh                file handle (handle)
3
4     OUT     buf                initial address of buffer (choice)
5
6     OUT     status             status object (status)
7
8 C binding
9 int MPI_File_read_ordered_end(MPI_File fh, void *buf, MPI_Status *status)
10
11 Fortran 2008 binding
12 MPI_File_read_ordered_end(fh, buf, status, ierror)
13     TYPE(MPI_File), INTENT(IN) :: fh
14     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
15     TYPE(MPI_Status) :: status
16     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
17
18 Fortran binding
19 MPI_FILE_READ_ORDERED_END(FH, BUF, STATUS, IERROR)
20     INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
21     <type> BUF(*)
22
23 MPI_FILE_WRITE_ORDERED_BEGIN(fh, buf, count, datatype)
24     INOUT   fh                file handle (handle)
25
26     IN      buf                initial address of buffer (choice)
27
28     IN      count              number of elements in buffer (integer)
29
30     IN      datatype           datatype of each buffer element (handle)
31
32 C binding
33 int MPI_File_write_ordered_begin(MPI_File fh, const void *buf, int count,
34     MPI_Datatype datatype)
35
36 int MPI_File_write_ordered_begin_c(MPI_File fh, const void *buf,
37     MPI_Count count, MPI_Datatype datatype)
38
39 Fortran 2008 binding
40 MPI_File_write_ordered_begin(fh, buf, count, datatype, ierror)
41     TYPE(MPI_File), INTENT(IN) :: fh
42     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
43     INTEGER, INTENT(IN) :: count
44     TYPE(MPI_Datatype), INTENT(IN) :: datatype
45     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
46
47 MPI_File_write_ordered_begin(fh, buf, count, datatype, ierror) !(_c)
48     TYPE(MPI_File), INTENT(IN) :: fh
49     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
50     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
51     TYPE(MPI_Datatype), INTENT(IN) :: datatype
52     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE_ORDERED_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)
    INTEGER FH, COUNT, DATATYPE, IERROR
    <type> BUF(*)

```

```

MPI_FILE_WRITE_ORDERED_END(fh, buf, status)

```

```

INOUT  fh          file handle (handle)
IN     buf         initial address of buffer (choice)
OUT    status      status object (status)

```

C binding

```

int MPI_File_write_ordered_end(MPI_File fh, const void *buf,
                               MPI_Status *status)

```

Fortran 2008 binding

```

MPI_File_write_ordered_end(fh, buf, status, ierror)
    TYPE(MPI_File), INTENT(IN) :: fh
    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
    TYPE(MPI_Status) :: status
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_WRITE_ORDERED_END(FH, BUF, STATUS, IERROR)
    INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
    <type> BUF(*)

```

14.5 File Interoperability

At the most basic level, file interoperability is the ability to read the information previously written to a file—not just the bits of data, but the actual information the bits represent. MPI guarantees full interoperability within a single MPI environment, and supports increased interoperability outside that environment through the external data representation (Section 14.5.2) as well as the data conversion functions (Section 14.5.3).

Interoperability within a single MPI environment (which could be considered “operability”) ensures that file data written by one MPI process can be read by any other MPI process, subject to the consistency constraints (see Section 14.6.1), provided that it would have been possible to start the two processes simultaneously and have them reside in a single MPI_COMM_WORLD. Furthermore, both processes must see the same data values at every absolute byte offset in the file for which data was written.

This single environment file interoperability implies that file data is accessible regardless of the number of processes.

There are three aspects to file interoperability:

- transferring the bits,
- converting between different file structures, and

- converting between different machine representations.

The first two aspects of file interoperability are beyond the scope of this standard, as both are highly machine dependent. However, transferring the bits of a file into and out of the MPI environment (e.g., by writing a file to tape) is required to be supported by all MPI implementations. In particular, an implementation must specify how familiar operations similar to POSIX `cp`, `rm`, and `mv` can be performed on the file. Furthermore, it is expected that the facility provided maintains the correspondence between absolute byte offsets (e.g., after possible file structure conversion, the data bits at byte offset 102 in the MPI environment are at byte offset 102 outside the MPI environment). As an example, a simple off-line conversion utility that transfers and converts files between the native file system and the MPI environment would suffice, provided it maintained the offset coherence mentioned above. In a high-quality implementation of MPI, users will be able to manipulate MPI files using the same or similar tools that the native file system offers for manipulating its files.

The remaining aspect of file interoperability, converting between different machine representations, is supported by the typing information specified in the `etype` and `filetype`. This facility allows the information in files to be shared between any two applications, regardless of whether they use MPI, and regardless of the machine architectures on which they run.

MPI supports multiple data representations: "native", "internal", and "external32". An implementation may support additional data representations. MPI also supports user-defined data representations (see Section 14.5.3). The "native" and "internal" data representations are implementation dependent, while the "external32" representation is common to all MPI implementations and facilitates file interoperability. The data representation is specified in the `datarep` argument to `MPI_FILE_SET_VIEW`.

Advice to users. MPI is not guaranteed to retain knowledge of what data representation was used when a file is written. Therefore, to correctly retrieve file data, an MPI application is responsible for specifying the same data representation as was used to create the file. (*End of advice to users.*)

"native": Data in this representation is stored in a file exactly as it is in memory. The advantage of this data representation is that data precision and I/O performance are not lost in type conversions with a purely homogeneous environment. The disadvantage is the loss of transparent interoperability within a heterogeneous MPI environment.

Advice to users. This data representation should only be used in a homogeneous MPI environment, or when the MPI application is capable of performing the datatype conversions itself. (*End of advice to users.*)

Advice to implementors. When implementing read and write operations on top of MPI message-passing, the message data should be typed as `MPI_BYTE` to ensure that the message routines do not perform any type conversions on the data. (*End of advice to implementors.*)

"internal": This data representation can be used for I/O operations in a homogeneous or heterogeneous environment; the implementation will perform type conversions if necessary. The implementation is free to store data in any format of its choice, with the restriction that it will maintain constant extents for all predefined datatypes in any

one file. The environment in which the resulting file can be reused is implementation-defined and must be documented by the implementation.

Rationale. This data representation allows the implementation to perform I/O efficiently in a heterogeneous environment, though with implementation-defined restrictions on how the file can be reused. (*End of rationale.*)

Advice to implementors. Since "external32" is a superset of the functionality provided by "internal", an implementation may choose to implement "internal" as "external32". (*End of advice to implementors.*)

"external32": This data representation states that read and write operations convert all data from and to the "external32" representation defined in Section 14.5.2. The data conversion rules for communication also apply to these conversions (see Section 3.3.2). The data on the storage medium is always in this canonical representation, and the data in memory is always in the local process's native representation.

This data representation has several advantages. First, all processes reading the file in a heterogeneous MPI environment will automatically have the data converted to their respective native representations. Second, the file can be exported from one MPI environment and imported into any other MPI environment with the guarantee that the second environment will be able to read all the data in the file.

The disadvantage of this data representation is that data precision and I/O performance may be lost in datatype conversions.

Advice to implementors. When implementing read and write operations on top of MPI message-passing, the message data should be converted to and from the "external32" representation in the client, and sent as type MPI_BYTE. This will avoid possible double datatype conversions and the associated further loss of precision and performance. (*End of advice to implementors.*)

14.5.1 Datatypes for File Interoperability

If the file data representation is other than "native", care must be taken in constructing etypes and filetypes. Any of the datatype constructor functions may be used; however, for those functions that accept displacements in bytes, the displacements must be specified in terms of their values in the file for the file data representation being used. MPI will interpret these byte displacements as is; no scaling will be done. The function MPI_FILE_GET_TYPE_EXTENT can be used to calculate the extents of datatypes in the file. For etypes and filetypes that are portable datatypes (see Section 2.4), MPI will scale any displacements in the datatypes to match the file data representation. Datatypes passed as arguments to read/write routines specify the data layout in memory; therefore, they must always be constructed using displacements corresponding to displacements in memory.

Advice to users. One can logically think of the file as if it were stored in the memory of a file server. The etype and filetype are interpreted as if they were defined at this file server, by the same sequence of calls used to define them at the calling process. If the data representation is "native", then this logical file server runs on the same architecture as the calling process, so that these types define the same data layout on the file as they would define in the memory of the calling process. If the

1 etype and filetype are portable datatypes, then the data layout defined in the file is
 2 the same as would be defined in the calling process memory, up to a scaling factor.
 3 The routine `MPI_FILE_GET_TYPE_EXTENT` can be used to calculate this scaling
 4 factor. Thus, two equivalent, portable datatypes will define the same data layout
 5 in the file, even in a heterogeneous environment with "internal", "external32", or user
 6 defined data representations. Otherwise, the etype and filetype must be constructed
 7 so that their typemap and extent are the same on any architecture. This can be
 8 achieved if they have an explicit upper bound and lower bound (defined using
 9 `MPI_TYPE_CREATE_RESIZED`). This condition must also be fulfilled by any datatype
 10 that is used in the construction of the etype and filetype, if this datatype is replicated
 11 contiguously, either explicitly, by a call to `MPI_TYPE_CONTIGUOUS`, or implicitly,
 12 by a blocklength argument that is greater than one. If an etype or filetype is not
 13 portable, and has a typemap or extent that is architecture dependent, then the data
 14 layout specified by it on a file is implementation dependent.

15 File data representations other than "native" may be different from corresponding
 16 data representations in memory. Therefore, for these file data representations, it is
 17 important not to use hardwired byte offsets for file positioning, including the initial
 18 displacement that specifies the view. When a portable datatype (see Section 2.4) is
 19 used in a data access operation, any holes in the datatype are scaled to match the data
 20 representation. However, note that this technique only works when all the processes
 21 that created the file view build their etypes from the same predefined datatypes. For
 22 example, if one process uses an etype built from `MPI_INT` and another uses an etype
 23 built from `MPI_FLOAT`, the resulting views may be nonportable because the relative
 24 sizes of these types may differ from one data representation to another. (*End of advice*
 25 *to users.*)

26
 27
 28
 29 `MPI_FILE_GET_TYPE_EXTENT(fh, datatype, extent)`

30	IN	fh	file handle (handle)
31			
32	IN	datatype	datatype (handle)
33	OUT	extent	datatype extent (integer)

34 C binding

35
 36 `int MPI_File_get_type_extent(MPI_File fh, MPI_Datatype datatype,`
 37 `MPI_Aint *extent)`

38
 39 `int MPI_File_get_type_extent_c(MPI_File fh, MPI_Datatype datatype,`
 40 `MPI_Count *extent)`

41 Fortran 2008 binding

42 `MPI_File_get_type_extent(fh, datatype, extent, ierror)`

43 `TYPE(MPI_File), INTENT(IN) :: fh`

44 `TYPE(MPI_Datatype), INTENT(IN) :: datatype`

45 `INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: extent`

46 `INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

47
 48 `MPI_File_get_type_extent(fh, datatype, extent, ierror) !(_c)`

```

TYPE(MPI_File), INTENT(IN) :: fh
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: extent
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_FILE_GET_TYPE_EXTENT(FH, DATATYPE, EXTENT, IERROR)
  INTEGER FH, DATATYPE, IERROR
  INTEGER(KIND=MPI_ADDRESS_KIND) EXTENT

```

Returns the extent of `datatype` in the file `fh`. This extent will be the same for all processes accessing the file `fh`. If the current view uses a user-defined data representation (see Section 14.5.3), MPI uses the `dtype_file_extent_fn` callback to calculate the extent.

If the `datatype` extent cannot be represented in `extent`, it is set to `MPI_UNDEFINED`.

Advice to implementors. In the case of user-defined data representations, the extent of a derived datatype can be calculated by first determining the extents of the predefined datatypes in this derived datatype using `dtype_file_extent_fn` (see Section 14.5.3). (*End of advice to implementors.*)

14.5.2 External Data Representation: "external32"

All MPI implementations are required to support the data representation defined in this section. Support of optional datatypes (e.g., `MPI_INTEGER2`) is not required.

All floating point values are in big-endian IEEE format [43] of the appropriate size. Floating point values are represented by one of three IEEE formats. These are the IEEE “Single (binary32),” “Double (binary64),” and “Double Extended (binary128)” formats, requiring 4, 8, and 16 bytes of storage, respectively. For the IEEE “Double Extended (binary128)” formats, MPI specifies a format width of 16 bytes, with 15 exponent bits, bias = +16383, 112 fraction bits, and an encoding analogous to the “Double (binary64)” format. All integral values are in two’s complement big-endian format. Big-endian means most significant byte at lowest address byte. For C `_Bool`, Fortran `LOGICAL`, and C++ `bool`, 0 implies false and nonzero implies true. C `float` `_Complex`, `double` `_Complex`, and `long double` `_Complex`, Fortran `COMPLEX` and `DOUBLE COMPLEX`, and other complex types are represented by a pair of floating point format values for the real and imaginary components. Characters are in ISO 8859-1 format [44]. Wide characters (of type `MPI_WCHAR`) are in Unicode format [69].

All signed numerals (e.g., `MPI_INT`, `MPI_REAL`) have the sign bit at the most significant bit. `MPI_COMPLEX` and `MPI_DOUBLE_COMPLEX` have the sign bit of the real and imaginary parts at the most significant bit of each part.

According to IEEE specifications [43], the “NaN” (not a number) is system dependent. It should not be interpreted within MPI as anything other than “NaN.”

Advice to implementors. The MPI treatment of “NaN” is similar to the approach used in XDR [66]. (*End of advice to implementors.*)

All data is byte aligned, regardless of type. All data items are stored contiguously in the file (if the file view is contiguous).

Advice to implementors. All bytes of `LOGICAL` and `bool` must be checked to determine the value. (*End of advice to implementors.*)

Advice to users. The type `MPI_PACKED` is treated as bytes and is not converted. The user should be aware that `MPI_PACK` has the option of placing a header in the beginning of the pack buffer. (*End of advice to users.*)

The sizes of the predefined datatypes returned from `MPI_TYPE_CREATE_F90_REAL`, `MPI_TYPE_CREATE_F90_COMPLEX`, and `MPI_TYPE_CREATE_F90_INTEGER` are defined in Section 19.1.9, page 809.

Advice to implementors. When converting a larger size integer to a smaller size integer, only the least significant bytes are moved. Care must be taken to preserve the sign bit value. This allows no conversion errors if the data range is within the range of the smaller size integer. (*End of advice to implementors.*)

Table 14.2, 14.3, and 14.4 specify the sizes of predefined, optional, and C++ datatypes in "external32" format, respectively.

14.5.3 User-Defined Data Representations

There are two situations that cannot be handled by the required representations:

1. a user wants to write a file in a representation unknown to the implementation, and
2. a user wants to read a file written in a representation unknown to the implementation.

User-defined data representations allow the user to insert a third party converter into the I/O stream to do the data representation conversion.

```
MPI_REGISTER_DATAREP(datarep, read_conversion_fn, write_conversion_fn,
                    dtype_file_extent_fn, extra_state)
```

IN	datarep	data representation identifier (string)
IN	read_conversion_fn	function invoked to convert from file representation to native representation (function)
IN	write_conversion_fn	function invoked to convert from native representation to file representation (function)
IN	dtype_file_extent_fn	function invoked to get the extent of a datatype as represented in the file (function)
IN	extra_state	extra state

C binding

```
int MPI_Register_datarep(const char *datarep,
                       MPI_Datarep_conversion_function *read_conversion_fn,
                       MPI_Datarep_conversion_function *write_conversion_fn,
                       MPI_Datarep_extent_function *dtype_file_extent_fn,
                       void *extra_state)

int MPI_Register_datarep_c(const char *datarep,
                          MPI_Datarep_conversion_function_c *read_conversion_fn,
                          MPI_Datarep_conversion_function_c *write_conversion_fn,
```


Table 14.2: "external32" sizes of predefined datatypes

Predefined Type	Length
MPI_PACKED	1
MPI_BYTE	1
MPI_CHAR	1
MPI_UNSIGNED_CHAR	1
MPI_SIGNED_CHAR	1
MPI_WCHAR	2
MPI_SHORT	2
MPI_UNSIGNED_SHORT	2
MPI_INT	4
MPI_LONG	4
MPI_UNSIGNED	4
MPI_UNSIGNED_LONG	4
MPI_LONG_LONG_INT	8
MPI_UNSIGNED_LONG_LONG	8
MPI_FLOAT	4
MPI_DOUBLE	8
MPI_LONG_DOUBLE	16
MPI_C_BOOL	1
MPI_INT8_T	1
MPI_INT16_T	2
MPI_INT32_T	4
MPI_INT64_T	8
MPI_UINT8_T	1
MPI_UINT16_T	2
MPI_UINT32_T	4
MPI_UINT64_T	8
MPI_AINT	8
MPI_COUNT	8
MPI_OFFSET	8
MPI_C_COMPLEX	2*4
MPI_C_FLOAT_COMPLEX	2*4
MPI_C_DOUBLE_COMPLEX	2*8
MPI_C_LONG_DOUBLE_COMPLEX	2*16
MPI_CHARACTER	1
MPI_LOGICAL	4
MPI_INTEGER	4
MPI_REAL	4
MPI_DOUBLE_PRECISION	8
MPI_COMPLEX	2*4
MPI_DOUBLE_COMPLEX	2*8
MPI_CXX_BOOL	1
MPI_CXX_FLOAT_COMPLEX	2*4
MPI_CXX_DOUBLE_COMPLEX	2*8
MPI_CXX_LONG_DOUBLE_COMPLEX	2*16

Table 14.3: "external32" sizes of optional datatypes

Predefined Type	Length
MPI_INTEGER1	1
MPI_INTEGER2	2
MPI_INTEGER4	4
MPI_INTEGER8	8
MPI_INTEGER16	16
MPI_REAL2	2
MPI_REAL4	4
MPI_REAL8	8
MPI_REAL16	16
MPI_COMPLEX4	2*2
MPI_COMPLEX8	2*4
MPI_COMPLEX16	2*8
MPI_COMPLEX32	2*16

Table 14.4: "external32" sizes of C++ datatypes

C++ Types	Length
MPI_CXX_BOOL	1
MPI_CXX_FLOAT_COMPLEX	2*4
MPI_CXX_DOUBLE_COMPLEX	2*8
MPI_CXX_LONG_DOUBLE_COMPLEX	2*16

```
MPI_Datarep_extent_function *dtype_file_extent_fn,
void *extra_state)
```

Fortran 2008 binding

```
MPI_Register_datarep(datarep, read_conversion_fn, write_conversion_fn,
dtype_file_extent_fn, extra_state, ierror)
CHARACTER(LEN=*), INTENT(IN) :: datarep
PROCEDURE(MPI_Datarep_conversion_function) :: read_conversion_fn,
write_conversion_fn
PROCEDURE(MPI_Datarep_extent_function) :: dtype_file_extent_fn
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Register_datarep_c(datarep, read_conversion_fn, write_conversion_fn,
dtype_file_extent_fn, extra_state, ierror) !(_c)
CHARACTER(LEN=*), INTENT(IN) :: datarep
PROCEDURE(MPI_Datarep_conversion_function_c) :: read_conversion_fn,
write_conversion_fn
PROCEDURE(MPI_Datarep_extent_function) :: dtype_file_extent_fn
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```

MPI_REGISTER_DATAREP(DATAREP, READ_CONVERSION_FN, WRITE_CONVERSION_FN,
                    DTYPE_FILE_EXTENT_FN, EXTRA_STATE, IERROR)
CHARACTER*(*) DATAREP
EXTERNAL READ_CONVERSION_FN, WRITE_CONVERSION_FN, DTYPE_FILE_EXTENT_FN
INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
INTEGER IERROR

```

The call associates `read_conversion_fn`, `write_conversion_fn`, and `dtype_file_extent_fn` with the data representation identifier `datarep`. `datarep` can then be used as an argument to `MPI_FILE_SET_VIEW`, causing subsequent data access operations to call the conversion functions to convert all data items accessed between file data representation and native representation. `MPI_REGISTER_DATAREP` is a local operation and only registers the data representation for the calling MPI process. If `datarep` is already defined, an error in the error class `MPI_ERR_DUP_DATAREP` is raised using the default file error handler (see Section 14.7). The length of a data representation string is limited to the value of `MPI_MAX_DATAREP_STRING`. `MPI_MAX_DATAREP_STRING` must have a value of at least 64. No routines are provided to delete data representations and free the associated resources; it is not expected that an application will generate them in significant numbers.

Extent Callback

```

typedef int MPI_Datarep_extent_function(MPI_Datatype datatype,
                                       MPI_Aint *extent, void *extra_state);
ABSTRACT INTERFACE
  SUBROUTINE MPI_Datarep_extent_function(datatype, extent, extra_state, ierror)
    TYPE(MPI_Datatype) :: datatype
    INTEGER(KIND=MPI_ADDRESS_KIND) :: extent, extra_state
    INTEGER :: ierror
  SUBROUTINE DATAREP_EXTENT_FUNCTION(DATATYPE, EXTENT, EXTRA_STATE, IERROR)
    INTEGER DATATYPE, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) EXTENT, EXTRA_STATE

```

The function `dtype_file_extent_fn` must return, in `file_extent`, the number of bytes required to store `datatype` in the file representation. The function is passed, in `extra_state`, the argument that was passed to the `MPI_REGISTER_DATAREP` call. MPI will only call this routine with predefined datatypes employed by the user.

Rationale. This callback does not have a large count variant because it is anticipated that large counts will not be required to represent the `extent` output value. (*End of rationale.*)

`MPI_Datarep_conversion_function` also supports large count types in separate additional MPI procedures in C (suffixed with the “_c”) and multiple abstract interfaces in Fortran when using `USE mpi_f08`.

If the extent cannot be represented in `extent`, the callback function shall set `extent` to `MPI_UNDEFINED`. The MPI implementation will then raise an error of class `MPI_ERR_VALUE_TOO_LARGE`.

1 *Datarep Conversion Functions*

```

2
3 typedef int MPI_Datarep_conversion_function(void *userbuf,
4         MPI_Datatype datatype, int count, void *filebuf,
5         MPI_Offset position, void *extra_state);
6
7 typedef int MPI_Datarep_conversion_function_c(void *userbuf,
8         MPI_Datatype datatype, MPI_Count count, void *filebuf,
9         MPI_Offset position, void *extra_state);
10
11 ABSTRACT INTERFACE
12     SUBROUTINE MPI_Datarep_conversion_function(userbuf, datatype, count, filebuf,
13         position, extra_state, ierror)
14     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
15     TYPE(C_PTR), VALUE :: userbuf, filebuf
16     TYPE(MPI_Datatype) :: datatype
17     INTEGER :: count, ierror
18     INTEGER(KIND=MPI_OFFSET_KIND) :: position
19     INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
20
21 ABSTRACT INTERFACE
22     SUBROUTINE MPI_Datarep_conversion_function_c(userbuf, datatype, count,
23         filebuf, position, extra_state, ierror) !(_c)
24     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
25     TYPE(C_PTR), VALUE :: userbuf, filebuf
26     TYPE(MPI_Datatype) :: datatype
27     INTEGER(KIND=MPI_COUNT_KIND) :: count
28     INTEGER(KIND=MPI_OFFSET_KIND) :: position
29     INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
30     INTEGER :: ierror
31
32 SUBROUTINE DATAREP_CONVERSION_FUNCTION(USERBUF, DATATYPE, COUNT, FILEBUF,
33     POSITION, EXTRA_STATE, IERROR)
34     <TYPE> USERBUF(*), FILEBUF(*)
35     INTEGER DATATYPE, COUNT, IERROR
36     INTEGER(KIND=MPI_OFFSET_KIND) POSITION
37     INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE

```

36 The function `read_conversion_fn` must convert from file data representation to native representation. Before calling this routine, MPI allocates and fills `filebuf` with `count` contiguous data items. The type of each data item matches the corresponding entry for the predefined `datatype` in the type signature of `datatype`. The function is passed, in `extra_state`, the argument that was passed to the `MPI_REGISTER_DATAREP` call. The function must copy all `count` data items from `filebuf` to `userbuf` in the distribution described by `datatype`, converting each data item from file representation to native representation. `datatype` will be equivalent to the `datatype` that the user passed to the read function. If the size of `datatype` is less than the size of the `count` data items, the conversion function must treat `datatype` as being contiguously tiled over the `userbuf`. The conversion function must begin storing converted data at the location in `userbuf` specified by `position` into the (tiled) `datatype`.

47 *Advice to users.* Although the conversion functions have similarities to `MPI_PACK`
48 and `MPI_UNPACK`, one should note the differences in the use of the arguments `count`

and `position`. In the conversion functions, `count` is a count of data items (i.e., count of `typemap` entries of `datatype`), and `position` is an index into this `typemap`. In `MPI_PACK`, `incount` refers to the number of whole `datatypes`, and `position` is a number of bytes. (*End of advice to users.*)

Advice to implementors. A converted read operation could be implemented as follows:

1. Get file extent of all data items
2. Allocate a filebuf large enough to hold all count data items
3. Read data from file into filebuf
4. Call `read_conversion_fn` to convert data and place it into `userbuf`
5. Deallocate filebuf

(*End of advice to implementors.*)

If MPI cannot allocate a buffer large enough to hold all the data to be converted from a read operation, it may call the conversion function repeatedly using the same `datatype` and `userbuf`, and reading successive chunks of data to be converted in `filebuf`. For the first call (and in the case when all the data to be converted fits into `filebuf`), MPI will call the function with `position` set to zero. Data converted during this call will be stored in the `userbuf` according to the first `count` data items in `datatype`. Then in subsequent calls to the conversion function, MPI will increment the value in `position` by the count of items converted in the previous call, and the `userbuf` pointer will be unchanged.

Rationale. Passing the conversion function a `position` and one `datatype` for the transfer allows the conversion function to decode the `datatype` only once and cache an internal representation of it on the `datatype`. Then on subsequent calls, the conversion function can use the `position` to quickly find its place in the `datatype` and continue storing converted data where it left off at the end of the previous call. (*End of rationale.*)

Advice to users. Although the conversion function may usefully cache an internal representation on the `datatype`, it should not cache any state information specific to an ongoing conversion operation, since it is possible for the same `datatype` to be used concurrently in multiple conversion operations. (*End of advice to users.*)

The function `write_conversion_fn` must convert from native representation to file data representation. Before calling this routine, MPI allocates `filebuf` of a size large enough to hold `count` contiguous data items. The type of each data item matches the corresponding entry for the predefined `datatype` in the type signature of `datatype`. The function must copy `count` data items from `userbuf` in the distribution described by `datatype`, to a contiguous distribution in `filebuf`, converting each data item from native representation to file representation. If the size of `datatype` is less than the size of `count` data items, the conversion function must treat `datatype` as being contiguously tiled over the `userbuf`.

The function must begin copying at the location in `userbuf` specified by `position` into the (tiled) `datatype`. `datatype` will be equivalent to the `datatype` that the user passed to the write function. The function is passed, in `extra_state`, the argument that was passed to the `MPI_REGISTER_DATAREP` call.

1 The predefined constant `MPI_CONVERSION_FN_NULL` may be used as either
2 `write_conversion_fn` or `read_conversion_fn` in bindings of `MPI_REGISTER_DATAREP` with-
3 out large counts in these conversion callbacks, whereas the constant
4 `MPI_CONVERSION_FN_NULL_C` can be used in the large count version (i.e.,
5 `MPI_Register_datarep_c`). In either of these cases, MPI will not attempt to invoke
6 `write_conversion_fn` or `read_conversion_fn`, respectively, but will perform the requested data
7 access using the native data representation.

8 An MPI implementation must ensure that all data accessed is converted, either by
9 using a filebuf large enough to hold all the requested data items or else by making repeated
10 calls to the conversion function with the same `datatype` argument and appropriate values
11 for `position`.

12 An implementation will only invoke the callback routines in this section
13 (`read_conversion_fn`, `write_conversion_fn`, and `dtype_file_extent_fn`) when one of the read or
14 write routines in Section 14.4, or `MPI_FILE_GET_TYPE_EXTENT` is called by the user.
15 `dtype_file_extent_fn` will only be passed predefined datatypes employed by the user. The
16 conversion functions will only be passed datatypes equivalent to those that the user has
17 passed to one of the routines noted above.

18 The conversion functions must be reentrant. User defined data representations are
19 restricted to use byte alignment for all types. Furthermore, it is erroneous for the conversion
20 functions to call any collective routines or to free `datatype`.

21 The conversion functions should return an error code. If the returned error code has
22 a value other than `MPI_SUCCESS`, the implementation will raise an error in the class
23 `MPI_ERR_CONVERSION`.

25 14.5.4 Matching Data Representations

26 It is the user's responsibility to ensure that the data representation used to read data from
27 a file is *compatible* with the data representation that was used to write that data to the file.

28 In general, using the same data representation name when writing and reading a file
29 does not guarantee that the representation is compatible. Similarly, using different repre-
30 sentation names on two different implementations may yield compatible representations.

31 Compatibility can be obtained when "external32" representation is used, although pre-
32 cision may be lost and the performance may be less than when "native" representation is
33 used. Compatibility is guaranteed using "external32" provided at least one of the following
34 conditions is met.

- 35 • The data access routines directly use types enumerated in Section 14.5.2, that are
36 supported by all implementations participating in the I/O. The predefined type used
37 to write a data item must also be used to read a data item.
- 38 • In the case of Fortran 90 programs, the programs participating in the data accesses
39 obtain compatible datatypes using MPI routines that specify precision and/or range
40 (Section 19.1.9).
- 41 • For any given data item, the programs participating in the data accesses use compat-
42 ible predefined types to write and read the data item.

43 User-defined data representations may be used to provide an implementation compat-
44 ibility with another implementation's "native" or "internal" representation.

Advice to users. Section 19.1.9 defines routines that support the use of matching datatypes in heterogeneous environments and contains examples illustrating their use. (*End of advice to users.*)

14.6 Consistency and Semantics

14.6.1 File Consistency

Consistency semantics define the outcome of multiple accesses to a single file. All file accesses in MPI are relative to a specific file handle created from a collective open. MPI provides three levels of consistency: sequential consistency among all accesses using a single file handle, sequential consistency among all accesses using file handles created from a single collective open with atomic mode enabled, and user-imposed consistency among accesses other than the above. Sequential consistency means the behavior of a set of operations will be as if the operations were performed in some serial order consistent with program order; each access appears atomic, although the exact ordering of accesses is unspecified. User-imposed consistency may be obtained using program order and calls to `MPI_FILE_SYNC`.

Let FH_1 be the set of file handles created from one particular collective open of the file FOO , and FH_2 be the set of file handles created from a different collective open of FOO . Note that nothing restrictive is said about FH_1 and FH_2 : the sizes of FH_1 and FH_2 may be different, the groups of processes used for each open may or may not intersect, the file handles in FH_1 may be destroyed before those in FH_2 are created, etc. Consider the following three cases: a single file handle (e.g., $fh_1 \in FH_1$), two file handles created from a single collective open (e.g., $fh_{1a} \in FH_1$ and $fh_{1b} \in FH_1$), and two file handles from different collective opens (e.g., $fh_1 \in FH_1$ and $fh_2 \in FH_2$).

For the purpose of consistency semantics, a matched pair (Section 14.4.5) of split collective data access operations (e.g., `MPI_FILE_READ_ALL_BEGIN` and `MPI_FILE_READ_ALL_END`) compose a single data access operation. Similarly, a non-blocking data access routine (e.g., `MPI_FILE_IREAD`) and the routine that completes the request (e.g., `MPI_WAIT`) also compose a single data access operation. For all cases below, these data access operations are subject to the same constraints as blocking data access operations.

Advice to users. For an `MPI_FILE_IREAD` and `MPI_WAIT` pair, the operation begins when `MPI_FILE_IREAD` is called and ends when `MPI_WAIT` returns. (*End of advice to users.*)

Assume that A_1 and A_2 are two data access operations. Let D_1 (D_2) be the set of absolute byte displacements of every byte accessed in A_1 (A_2). The two data accesses **overlap** if $D_1 \cap D_2 \neq \emptyset$. The two data accesses **conflict** if they overlap and at least one is a write access.

Let SEQ_{fh} be a sequence of file operations on a single file handle, bracketed by `MPI_FILE_SYNC`s on that file handle. (Both opening and closing a file implicitly perform an `MPI_FILE_SYNC`.) SEQ_{fh} is a “write sequence” if any of the data access operations in the sequence are writes or if any of the file manipulation operations in the sequence change the state of the file (e.g., `MPI_FILE_SET_SIZE` or `MPI_FILE_PREALLOCATE`). Given two sequences, SEQ_1 and SEQ_2 , we say they are not **concurrent** if one sequence is guaranteed to completely precede the other (temporally).

The requirements for guaranteeing sequential consistency among all accesses to a particular file are divided into the three cases given below. If any of these requirements are not met, then the value of all data in that file is implementation dependent.

Case 1: $fh_1 \in FH_1$. All operations on fh_1 are sequentially consistent if atomic mode is set. If nonatomic mode is set, then all operations on fh_1 are sequentially consistent if they are either nonconcurrent, nonconflicting, or both.

Case 2: $fh_{1a} \in FH_1$ and $fh_{1b} \in FH_1$. Assume A_1 is a data access operation using fh_{1a} , and A_2 is a data access operation using fh_{1b} . If for any access A_1 , there is no access A_2 that conflicts with A_1 , then MPI guarantees sequential consistency.

However, unlike POSIX semantics, the default MPI semantics for conflicting accesses do not guarantee sequential consistency. If A_1 and A_2 conflict, sequential consistency can be guaranteed by either enabling atomic mode via the `MPI_FILE_SET_ATOMICITY` routine, or meeting the condition described in Case 3 below.

Case 3: $fh_1 \in FH_1$ and $fh_2 \in FH_2$. Consider access to a single file using file handles from distinct collective opens. In order to guarantee sequential consistency, `MPI_FILE_SYNC` must be used (both opening and closing a file implicitly perform an `MPI_FILE_SYNC`).

Sequential consistency is guaranteed among accesses to a single file if for any write sequence SEQ_1 to the file, there is no sequence SEQ_2 to the file that is *concurrent* with SEQ_1 . To guarantee sequential consistency when there are write sequences, `MPI_FILE_SYNC` must be used together with a mechanism that guarantees nonconcurrency of the sequences.

See the examples in Section 14.6.11 for further clarification of some of these consistency semantics.

`MPI_FILE_SET_ATOMICITY(fh, flag)`

INOUT	fh	file handle (handle)
IN	flag	true to set atomic mode, false to set nonatomic mode (logical)

C binding

```
int MPI_File_set_atomicity(MPI_File fh, int flag)
```

Fortran 2008 binding

```
MPI_File_set_atomicity(fh, flag, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  LOGICAL, INTENT(IN) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_SET_ATOMICITY(FH, FLAG, IERROR)
  INTEGER FH, IERROR
  LOGICAL FLAG
```

Let FH be the set of file handles created by one collective open. The consistency semantics for data access operations using FH is set by collectively calling

MPI_FILE_SET_ATOMICITY on *FH*. MPI_FILE_SET_ATOMICITY is collective; all processes in the group must pass identical values for *fh* and *flag*. If *flag* is true, atomic mode is set; if *flag* is false, nonatomic mode is set.

Changing the consistency semantics for an open file only affects new data accesses. All completed data accesses are guaranteed to abide by the consistency semantics in effect during their execution. Nonblocking data accesses and split collective operations that have not completed (e.g., via MPI_WAIT) are only guaranteed to abide by nonatomic mode consistency semantics.

Advice to implementors. Since the semantics guaranteed by atomic mode are stronger than those guaranteed by nonatomic mode, an implementation is free to adhere to the more stringent atomic mode semantics for outstanding requests. (*End of advice to implementors.*)

MPI_FILE_GET_ATOMICITY(*fh*, *flag*)

IN	<i>fh</i>	file handle (handle)
OUT	<i>flag</i>	true if atomic mode, false if nonatomic mode (logical)

C binding

```
int MPI_File_get_atomicity(MPI_File fh, int *flag)
```

Fortran 2008 binding

```
MPI_File_get_atomicity(fh, flag, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_FILE_GET_ATOMICITY(FH, FLAG, IERROR)
  INTEGER FH, IERROR
  LOGICAL FLAG
```

MPI_FILE_GET_ATOMICITY returns the current consistency semantics for data access operations on the set of file handles created by one collective open. If *flag* is true, atomic mode is enabled; if *flag* is false, nonatomic mode is enabled.

MPI_FILE_SYNC(*fh*)

INOUT	<i>fh</i>	file handle (handle)
-------	-----------	----------------------

C binding

```
int MPI_File_sync(MPI_File fh)
```

Fortran 2008 binding

```
MPI_File_sync(fh, ierror)
  TYPE(MPI_File), INTENT(IN) :: fh
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

1 Fortran binding

```
2 MPI_FILE_SYNC(FH, IERROR)
3     INTEGER FH, IERROR
```

4 Calling `MPI_FILE_SYNC` with `fh` causes all previous writes to `fh` by the calling process
5 to be transferred to the storage device. If other processes have made updates to the storage
6 device, then all such updates become visible to subsequent reads of `fh` by the calling process.
7 `MPI_FILE_SYNC` may be necessary to ensure sequential consistency in certain cases (see
8 above).

9 `MPI_FILE_SYNC` is a collective operation.

10 The user is responsible for ensuring that all nonblocking requests and split collective
11 operations on `fh` have been completed before calling `MPI_FILE_SYNC`—otherwise, the call
12 to `MPI_FILE_SYNC` is erroneous.
13

14 14.6.2 Random Access vs. Sequential Files

15 MPI distinguishes ordinary random access files from sequential stream files, such as pipes
16 and tape files. Sequential stream files must be opened with the `MPI_MODE_SEQUENTIAL`
17 flag set in the `amode`. For these files, the only permitted data access operations are shared
18 file pointer reads and writes. Filetypes and etypes with holes are erroneous. In addition, the
19 notion of file pointer is not meaningful; therefore, calls to `MPI_FILE_SEEK_SHARED` and
20 `MPI_FILE_GET_POSITION_SHARED` are erroneous, and the pointer update rules specified
21 for the data access routines do not apply. The amount of data accessed by a data access
22 operation will be the amount requested unless the end of file is reached or an error is raised.
23

24
25 *Rationale.* This implies that reading on a pipe will always wait until the requested
26 amount of data is available or until the process writing to the pipe has issued an end
27 of file. (*End of rationale.*)
28

29 Finally, for some sequential files, such as those corresponding to magnetic tapes or
30 streaming network connections, writes to the file may be destructive. In other words, a
31 write may act as a truncate (a `MPI_FILE_SET_SIZE` with `size` set to the current position)
32 followed by the write.
33

34 14.6.3 Progress

35 The *progress* rules of MPI are both a promise to users and a set of constraints on imple-
36 mentors. In cases where the progress rules restrict possible implementation choices more
37 than the interface specification alone, the progress rules take precedence.
38

39 All blocking routines must complete in finite time unless an exceptional condition (such
40 as resource exhaustion) causes an error.

41 Nonblocking data access routines inherit the following progress rule from nonblocking
42 point-to-point communication: a nonblocking write is equivalent to a nonblocking send for
43 which a receive is eventually posted, and a nonblocking read is equivalent to a nonblocking
44 receive for which a send is eventually posted.

45 Finally, an implementation is free to delay progress of collective routines until all pro-
46 cesses in the group associated with the collective call have invoked the routine. Once all
47 processes in the group have invoked the routine, the progress rule of the equivalent noncol-
48 lective routine must be followed.

14.6.4 Collective File Operations

Collective file operations are subject to the same restrictions as collective communication operations. For a complete discussion, please refer to the semantics set forth in Section 6.14.

Collective file operations are collective over a duplicate of the communicator used to open the file—this duplicate communicator is implicitly specified via the file handle argument. Different processes can pass different values for other arguments of a collective routine unless specified otherwise.

14.6.5 Nonblocking Collective File Operations

Nonblocking collective file operations are defined only for data access routines with explicit offsets and individual file pointers but not with shared file pointers.

Nonblocking collective file operations are subject to the same restrictions as blocking collective I/O operations. All processes belonging to the group of the communicator that was used to open the file must call collective I/O operations (blocking and nonblocking) in the same order. This is consistent with the ordering rules for collective operations in threaded environments. For a complete discussion, please refer to the semantics set forth in Section 6.14.

Nonblocking collective I/O operations do not match with blocking collective I/O operations. Multiple nonblocking collective I/O operations can be outstanding on a single file handle. High quality MPI implementations should be able to support a large number of *pending* nonblocking I/O operations.

All nonblocking collective I/O calls are local and return immediately, irrespective of the status of other processes. The call initiates the operation that may progress independently of any communication, computation, or I/O. The call returns a request handle, which must be passed to a completion call. Input buffers should not be modified and output buffers should not be accessed before the completion call returns. The same *progress* rules described for nonblocking collective operations apply for nonblocking collective I/O operations. For a complete discussion, please refer to the semantics set forth in Section 6.12.

14.6.6 Type Matching

The type matching rules for I/O mimic the type matching rules for communication with one exception: if `etype` is `MPI_BYTE`, then this matches any `datatype` in a data access operation. In general, the `etype` of data items written must match the `etype` used to read the items, and for each data access operation, the current `etype` must also match the type declaration of the data access buffer.

Advice to users. In most cases, use of `MPI_BYTE` as a wild card will defeat the file interoperability features of MPI. File interoperability can only perform automatic conversion between heterogeneous data representations when the exact datatypes accessed are explicitly specified. (*End of advice to users.*)

14.6.7 Miscellaneous Clarifications

Once an I/O routine completes, it is safe to free any opaque objects passed as arguments to that routine. For example, the `comm` and `info` used in an `MPI_FILE_OPEN`, or the `etype` and `filetype` used in an `MPI_FILE_SET_VIEW`, can be freed without affecting access to the

1 file. Note that for nonblocking routines and split collective operations, the operation must
2 be completed before it is safe to reuse data buffers passed as arguments.

3 As in communication, datatypes must be committed before they can be used in file
4 manipulation or data access operations. For example, the `etype` and `filetype` must be com-
5 mitted before calling `MPI_FILE_SET_VIEW`, and the `datatype` must be committed before
6 calling `MPI_FILE_READ` or `MPI_FILE_WRITE`.

8 14.6.8 MPI_Offset Type

9
10 `MPI_Offset` is an integer type of size sufficient to represent the size (in bytes) of the largest
11 file supported by MPI. Displacements and offsets are always specified as values of type
12 `MPI_Offset`.

13 In Fortran, the corresponding integer is an integer with kind parameter
14 `MPI_OFFSET_KIND`, which is defined in the `mpi_f08` module, the `mpi` module and the `mpif.h`
15 include file.

16 In Fortran 77 environments that do not support KIND parameters, `MPI_Offset` arguments
17 should be declared as an `INTEGER` of suitable size. The language interoperability implications
18 for `MPI_Offset` are similar to those for addresses (see Section 19.3).

19 14.6.9 Logical vs. Physical File Layout

20
21 MPI specifies how the data should be laid out in a virtual file structure (the view), not
22 how that file structure is to be stored on one or more disks. Specification of the physical
23 file structure was avoided because it is expected that the mapping of files to disks will be
24 system specific, and any specific control over file layout would therefore restrict program
25 portability. However, there are still cases where some information may be necessary to
26 optimize file layout. This information can be provided as *hints* specified via `info` when a file
27 is created (see Section 14.2.8).

28 14.6.10 File Size

29
30 The size of a file may be increased by writing to the file after the current end of file. The size
31 may also be changed by calling MPI *size changing* routines, such as `MPI_FILE_SET_SIZE`.
32 A call to a size changing routine does not necessarily change the file size. For example,
33 calling `MPI_FILE_PREALLOCATE` with a size less than the current size does not change the
34 size.

35
36 Consider a set of bytes that has been written to a file since the most recent call to a
37 size changing routine, or since `MPI_FILE_OPEN` if no such routine has been called. Let the
38 *high byte* be the byte in that set with the largest displacement. The file size is the larger of

- 39 • One plus the displacement of the high byte.
- 40 • The size immediately after the size changing routine, or `MPI_FILE_OPEN`, returned.

41
42 When applying consistency semantics, calls to `MPI_FILE_SET_SIZE` and
43 `MPI_FILE_PREALLOCATE` are considered writes to the file (which conflict with operations
44 that access bytes at displacements between the old and new file sizes), and
45 `MPI_FILE_GET_SIZE` is considered a read of the file (which overlaps with all accesses to the
46 file).

Advice to users. Any sequence of operations containing the collective routines MPI_FILE_SET_SIZE and MPI_FILE_PREALLOCATE is a write sequence. As such, sequential consistency in nonatomic mode is not guaranteed unless the conditions in Section 14.6.1 are satisfied. (*End of advice to users.*)

File pointer update semantics (i.e., file pointers are updated by the amount accessed) are only guaranteed if file size changes are sequentially consistent.

Advice to users. Consider the following example. Given two operations made by separate processes to a file containing 100 bytes: an MPI_FILE_READ of 10 bytes and an MPI_FILE_SET_SIZE to 0 bytes. If the user does not enforce sequential consistency between these two operations, the file pointer may be updated by the amount requested (10 bytes) even if the amount accessed is zero bytes. (*End of advice to users.*)

14.6.11 Examples

The examples in this section illustrate the application of the MPI consistency and semantics guarantees. These address

- conflicting accesses on file handles obtained from a single collective open, and
- all accesses on file handles obtained from two separate collective opens.

The simplest way to achieve consistency for conflicting accesses is to obtain sequential consistency by setting atomic mode.

Example 14.5. For the code below, process 1 will read either 0 or 10 integers. If the latter, every element of `b` will be 5. If nonatomic mode is set, the results of the read are undefined.

```

/* Process 0 */

int i, a[10];
int TRUE = 1;

for (i=0;i<10;i++)
    a[i] = 5;

MPI_File_open(MPI_COMM_WORLD, "workfile",
              MPI_MODE_RDWR | MPI_MODE_CREATE, MPI_INFO_NULL, &fh0);
MPI_File_set_view(fh0, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
MPI_File_set_atomicity(fh0, TRUE);
MPI_File_write_at(fh0, 0, a, 10, MPI_INT, &status);
/* MPI_Barrier(MPI_COMM_WORLD); */

/* Process 1 */

int b[10];
int TRUE = 1;
MPI_File_open(MPI_COMM_WORLD, "workfile",
              MPI_MODE_RDWR | MPI_MODE_CREATE, MPI_INFO_NULL, &fh1);
MPI_File_set_view(fh1, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);

```

```

1 MPI_File_set_atomicsity(fh1, TRUE);
2 /* MPI_Barrier(MPI_COMM_WORLD); */
3 MPI_File_read_at(fh1, 0, b, 10, MPI_INT, &status);

```

A user may guarantee that the write on process 0 precedes the read on process 1 by imposing temporal order with, for example, calls to MPI_BARRIER.

Advice to users. Routines other than MPI_BARRIER may be used to impose temporal order. In the example above, process 0 could use MPI_SEND to send a 0 byte message, received by process 1 using MPI_RECV. (*End of advice to users.*)

Example 14.6. Alternatively, a user can impose consistency with nonatomic mode set:

```

13 /* Process 0 */
14 int i, a[10];
15 for (i=0; i<10; i++)
16     a[i] = 5;
17
18 MPI_File_open(MPI_COMM_WORLD, "workfile",
19              MPI_MODE_RDWR | MPI_MODE_CREATE, MPI_INFO_NULL, &fh0);
20 MPI_File_set_view(fh0, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
21 MPI_File_write_at(fh0, 0, a, 10, MPI_INT, &status);
22 MPI_File_sync(fh0);
23 MPI_Barrier(MPI_COMM_WORLD);
24 MPI_File_sync(fh0);

```

```

25 /* Process 1 */
26
27 int b[10];
28 MPI_File_open(MPI_COMM_WORLD, "workfile",
29              MPI_MODE_RDWR | MPI_MODE_CREATE, MPI_INFO_NULL, &fh1);
30 MPI_File_set_view(fh1, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
31 MPI_File_sync(fh1);
32 MPI_Barrier(MPI_COMM_WORLD);
33 MPI_File_sync(fh1);
34 MPI_File_read_at(fh1, 0, b, 10, MPI_INT, &status);

```

The “sync-barrier-sync” construct is required because:

- The barrier ensures that the write on process 0 occurs before the read on process 1.
- The first sync guarantees that the data written by all processes is transferred to the storage device.
- The second sync guarantees that all data that has been transferred to the storage device is visible to all processes. (This does not affect process 0 in this example.)

Example 14.7. The following program represents an erroneous attempt to achieve consistency by eliminating the apparently superfluous second “sync” call for each process.

```

46 /* ----- THIS EXAMPLE IS ERRONEOUS ----- */
47 /* Process 0 */

```

```

1  int i, a[10];
2  for (i=0;i<10;i++)
3      a[i] = 5;
4
5  MPI_File_open(MPI_COMM_WORLD, "workfile",
6                MPI_MODE_RDWR | MPI_MODE_CREATE, MPI_INFO_NULL, &fh0);
7  MPI_File_set_view(fh0, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
8  MPI_File_write_at(fh0, 0, a, 10, MPI_INT, &status);
9  MPI_File_sync(fh0);
10 MPI_Barrier(MPI_COMM_WORLD);
11
12 /* Process 1 */
13
14 int b[10];
15 MPI_File_open(MPI_COMM_WORLD, "workfile",
16               MPI_MODE_RDWR | MPI_MODE_CREATE, MPI_INFO_NULL, &fh1);
17 MPI_File_set_view(fh1, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
18 MPI_Barrier(MPI_COMM_WORLD);
19 MPI_File_sync(fh1);
20 MPI_File_read_at(fh1, 0, b, 10, MPI_INT, &status);
21
22 /* ----- THIS EXAMPLE IS ERRONEOUS ----- */

```

The above program also violates the MPI rule against out-of-order collective operations and will deadlock for implementations in which MPI_FILE_SYNC blocks.

Advice to users. Some implementations may choose to implement MPI_FILE_SYNC as a temporally synchronizing function. When using such an implementation, the “sync-barrier-sync” construct above can be replaced by a single “sync.” The results of using such code with an implementation for which MPI_FILE_SYNC is not temporally synchronizing is undefined. (*End of advice to users.*)

Asynchronous I/O

The behavior of asynchronous I/O operations is determined by applying the rules specified above for synchronous I/O operations.

Example 14.8. The following examples all access a preexisting file “myfile.” Word 10 in myfile initially contains the integer 2. Each example writes and reads word 10. First consider the following code fragment:

```

37 int a = 4, b, TRUE=1;
38 MPI_File_open(MPI_COMM_WORLD, "myfile",
39               MPI_MODE_RDWR, MPI_INFO_NULL, &fh);
40 MPI_File_set_view(fh, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
41 /* MPI_File_set_atomics(fh, TRUE); Use this to set atomic mode. */
42 MPI_File_iread_at(fh, 10, &b, 1, MPI_INT, &reqs[1]);
43 MPI_File_iwrite_at(fh, 10, &a, 1, MPI_INT, &reqs[0]);
44 MPI_Waitall(2, reqs, statuses);

```

For asynchronous data access operations, MPI specifies that the access occurs at any time between the call to the asynchronous data access routine and the return from the corresponding request complete routine. Thus, executing either the read before the write, or the

1 write before the read is consistent with program order. If atomic mode is set, then MPI
 2 guarantees sequential consistency, and the program will read either 2 or 4 into b. If atomic
 3 mode is not set, then sequential consistency is not guaranteed and the program may read
 4 something other than 2 or 4 due to the conflicting data access.

5 Similarly, the following code fragment does not order file accesses:

```
6
7 int a = 4, b;
8 MPI_File_open(MPI_COMM_WORLD, "myfile",
9               MPI_MODE_RDWR, MPI_INFO_NULL, &fh);
10 MPI_File_set_view(fh, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
11 /* MPI_File_set_atomicity(fh, TRUE); Use this to set atomic mode. */
12 MPI_File_iread_at(fh, 10, &b, 1, MPI_INT, &reqs[1]);
13 MPI_File_iwrite_at(fh, 10, &a, 1, MPI_INT, &reqs[0]);
14 MPI_Wait(&reqs[0], &status);
15 MPI_Wait(&reqs[1], &status);
```

16 If atomic mode is set, either 2 or 4 will be read into b. Again, MPI does not guarantee
 17 sequential consistency in nonatomic mode.

18 On the other hand, the following code fragment:

```
19 int a = 4, b;
20 MPI_File_open(MPI_COMM_WORLD, "myfile",
21               MPI_MODE_RDWR, MPI_INFO_NULL, &fh);
22 MPI_File_set_view(fh, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
23 MPI_File_iwrite_at(fh, 10, &a, 1, MPI_INT, &reqs[0]);
24 MPI_Wait(&reqs[0], &status);
25 MPI_File_iread_at(fh, 10, &b, 1, MPI_INT, &reqs[1]);
26 MPI_Wait(&reqs[1], &status);
```

27 defines the same ordering as:

```
28 int a = 4, b;
29 MPI_File_open(MPI_COMM_WORLD, "myfile",
30               MPI_MODE_RDWR, MPI_INFO_NULL, &fh);
31 MPI_File_set_view(fh, 0, MPI_INT, MPI_INT, "native", MPI_INFO_NULL);
32 MPI_File_write_at(fh, 10, &a, 1, MPI_INT, &status);
33 MPI_File_read_at(fh, 10, &b, 1, MPI_INT, &status);
```

34 Since

- 35 • nonconcurrent operations on a single file handle are sequentially consistent, and
- 36 • the program fragments specify an order for the operations,

37 MPI guarantees that both program fragments will read the value 4 into b. There is no need
 38 to set atomic mode for this example.

39 Similar considerations apply to conflicting accesses of the form:

```
40 MPI_File_iwrite_all(fh, ...);
41 MPI_File_iread_all(fh, ...);
42 MPI_Waitall(...);
```

43 In addition, as mentioned in Section 14.6.5, nonblocking collective I/O operations have to
 44 be called in the same order on the file handle by all processes.

45 Similar considerations apply to conflicting accesses of the form:


```

MPI_File_write_all_begin(fh,...);
MPI_File_iread(fh,...);
MPI_Wait(fh,...);
MPI_File_write_all_end(fh,...);

```

Recall that constraints governing consistency and semantics are not relevant to the following:

```

MPI_File_write_all_begin(fh,...);
MPI_File_read_all_begin(fh,...);
MPI_File_read_all_end(fh,...);
MPI_File_write_all_end(fh,...);

```

since split collective operations on the same file handle may not overlap (see Section 14.4.5).

14.7 I/O Error Handling

By default, communication errors are fatal—`MPI_ERRORS_ARE_FATAL` is the default error handler associated with `MPI_COMM_WORLD`. I/O errors are usually less catastrophic (e.g., “file not found”) than communication errors, and common practice is to catch these errors and continue executing. For this reason, MPI provides additional error facilities for I/O.

Advice to users. MPI does not specify the state of a computation after an erroneous MPI call has occurred. A high-quality implementation will support the I/O error handling facilities, allowing users to write programs using common practice for I/O. (*End of advice to users.*)

Like communicators, each file handle has an error handler associated with it. The MPI I/O error handling routines are defined in Section 9.3.

When MPI calls a user-defined error handler resulting from an error on a particular file handle, the first two arguments passed to the file error handler are the file handle and the error code. For I/O errors that are not associated with a valid file handle (e.g., in `MPI_FILE_OPEN` or `MPI_FILE_DELETE`), the first argument passed to the error handler is `MPI_FILE_NULL`.

I/O error handling differs from communication error handling in another important aspect. By default, the predefined error handler for file handles is `MPI_ERRORS_RETURN`. The **default file error** handler has two purposes: when a new file handle is created (by `MPI_FILE_OPEN`), the error handler for the new file handle is initially set to the default file error handler, and I/O routines that have no valid file handle on which to raise an error (e.g., `MPI_FILE_OPEN` or `MPI_FILE_DELETE`) use the default file error handler. The default file error handler can be changed by specifying `MPI_FILE_NULL` as the `fh` argument to `MPI_FILE_SET_ERRHANDLER`. The current value of the default file error handler can be determined by passing `MPI_FILE_NULL` as the `fh` argument to `MPI_FILE_GET_ERRHANDLER`.

Rationale. For communication, the default error handler is inherited from `MPI_COMM_WORLD` when using the World Model. In I/O, there is no analogous “root” file handle from which default properties can be inherited. Rather than invent a new global file handle, the default file error handler is manipulated as if it were attached to `MPI_FILE_NULL`. (*End of rationale.*)

Table 14.5: I/O error classes

MPI_ERR_FILE	Invalid file handle
MPI_ERR_NOT_SAME	Collective argument not identical on all processes, or collective routines called in a different order by different processes
MPI_ERR_AMODE	Error related to the <code>amode</code> passed to <code>MPI_FILE_OPEN</code>
MPI_ERR_UNSUPPORTED_DATAREP	Unsupported <code>datarep</code> passed to <code>MPI_FILE_SET_VIEW</code>
MPI_ERR_UNSUPPORTED_OPERATION	Unsupported operation, such as seeking on a file that supports sequential access only
MPI_ERR_NO_SUCH_FILE	File does not exist
MPI_ERR_FILE_EXISTS	File exists
MPI_ERR_BAD_FILE	Invalid file name (e.g., path name too long)
MPI_ERR_ACCESS	Permission denied
MPI_ERR_NO_SPACE	Not enough space
MPI_ERR_QUOTA	Quota exceeded
MPI_ERR_READ_ONLY	Read-only file or file system
MPI_ERR_FILE_IN_USE	File operation could not be completed, as the file is currently open by some process
MPI_ERR_DUP_DATAREP	Conversion functions could not be registered because a data representation identifier that was already defined was passed to <code>MPI_REGISTER_DATAREP</code>
MPI_ERR_CONVERSION	An error occurred in a user supplied data conversion function.
MPI_ERR_IO	Other I/O error

14.8 I/O Error Classes

The implementation dependent error codes returned by the I/O routines can be converted into the error classes defined in Table 14.5.

In addition, calls to routines in this chapter may raise errors in other MPI classes, such as `MPI_ERR_TYPE`.

14.9 Examples

14.9.1 Double Buffering with Split Collective I/O

Example 14.9. This example shows how to overlap computation and output. The computation is performed by the function `compute_buffer()`.

```

/*=====
 *
 * Function:          double_buffer
 *
 * Synopsis:

```

```

*   void double_buffer(
*       MPI_File fh,                ** IN
*       MPI_Datatype buftype,      ** IN
*       int bufcount                ** IN
*   )
*
* Description:
*   Performs the steps to overlap computation with a collective write
*   by using a double-buffering technique.
*
* Parameters:
*   fh                previously opened MPI file handle
*   buftype            MPI datatype for memory layout
*                     (Assumes a compatible view has been set on fh)
*   bufcount          # buftype elements to transfer
*-----*/
/* this macro switches which buffer "x" is pointing to */
#define TOGGLE_PTR(x) (((x)==(buffer1)) ? (x=buffer2) : (x=buffer1))

void double_buffer(MPI_File fh, MPI_Datatype buftype, int bufcount)
{
    MPI_Status status;           /* status for MPI calls */
    float *buffer1, *buffer2;   /* buffers to hold results */
    float *compute_buf_ptr;     /* destination buffer */
                                /* for computing */
    float *write_buf_ptr;       /* source for writing */
    int done;                   /* determines when to quit */

    /* buffer initialization */
    buffer1 = (float *)
                malloc(bufcount*sizeof(float));
    buffer2 = (float *)
                malloc(bufcount*sizeof(float));
    compute_buf_ptr = buffer1;   /* initially point to buffer1 */
    write_buf_ptr = buffer1;     /* initially point to buffer1 */

    /* DOUBLE-BUFFER prolog:
     * compute buffer1; then initiate writing buffer1 to disk
     */
    compute_buffer(compute_buf_ptr, bufcount, &done);
    MPI_File_write_all_begin(fh, write_buf_ptr, bufcount, buftype);

    /* DOUBLE-BUFFER steady state:
     * Overlap writing old results from buffer pointed to by write_buf_ptr
     * with computing new results into buffer pointed to by compute_buf_ptr.
     *
     * There is always one write-buffer and one compute-buffer in use
     * during steady state.
     */
    while (!done) {
        TOGGLE_PTR(compute_buf_ptr);
        compute_buffer(compute_buf_ptr, bufcount, &done);
        MPI_File_write_all_end(fh, write_buf_ptr, &status);
        TOGGLE_PTR(write_buf_ptr);
    }
}

```

```

1  MPI_File_write_all_begin(fh, write_buf_ptr, bufcount, buftype);
2  }
3
4  /* DOUBLE-BUFFER epilog:
5   *   wait for final write to complete.
6   */
7  MPI_File_write_all_end(fh, write_buf_ptr, &status);
8
9  /* buffer cleanup */
10 free(buffer1);
11 free(buffer2);
12 }

```

14.9.2 Subarray Filetype Constructor




Figure 14.4: Example array file layout




Figure 14.5: Example local array filetype for process 1

Example 14.10. Assume we are writing out a 100×100 2D array of double precision floating point numbers that is distributed among 4 processes such that each process has a block of 25 columns (e.g., process 0 has columns 0–24, process 1 has columns 25–49, etc.; see Figure 14.4). To create the filetypes for each process one could use the following C program

(see Section 5.1.3):

```

1
2
3
4
5
6
7
8
9
10
11
12
13
double subarray[100][25];
MPI_Datatype filetype;
int sizes[2], subsizes[2], starts[2];
int rank;

MPI_Comm_rank(MPI_COMM_WORLD, &rank);
sizes[0]=100; sizes[1]=100;
subsizes[0]=100; subsizes[1]=25;
starts[0]=0; starts[1]=rank*subsizes[1];

MPI_Type_create_subarray(2, sizes, subsizes, starts, MPI_ORDER_C,
MPI_DOUBLE, &filetype);

```

Or, equivalently in Fortran:

Example 14.11. Writing out a 100×100 2D array of double precision floating point numbers that is distributed among 4 processes such that each process has a block of 25 columns (e.g., process 0 has columns 0–24, process 1 has columns 25–49, etc.; see Figure 14.4).

```

16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
double precision subarray(100,25)
integer filetype, rank, ierror
integer sizes(2), subsizes(2), starts(2)

call MPI_COMM_RANK(MPI_COMM_WORLD, rank, ierror)
sizes(1) = 100
sizes(2) = 100
subsizes(1) = 100
subsizes(2) = 25
starts(1) = 0
starts(2) = rank*subsizes(2)

call MPI_TYPE_CREATE_SUBARRAY(2, sizes, subsizes, starts, &
MPI_ORDER_FORTRAN, MPI_DOUBLE_PRECISION, &
filetype, ierror)

```

The generated filetype will then describe the portion of the file contained within the process's subarray with holes for the space taken by the other processes. Figure 14.5 shows the filetype created for process 1.

DRAFT

Chapter 15

Tool Support

15.1 Introduction

This chapter discusses interfaces that allow debuggers, performance analyzers, and other tools to extract information about the behavior of MPI processes. Specifically, this chapter defines both the MPI profiling interface (Section 15.2), which supports the transparent interception and inspection of MPI calls, and the MPI tool information interface (Section 15.3), which supports the inspection and manipulation of MPI control and performance variables, as well as the registration of callbacks for MPI library events. The interfaces described in this chapter are all defined in the context of an MPI process, i.e., are callable from the same code that invokes other MPI functions.

15.2 Profiling Interface

15.2.1 Requirements

To meet the requirements for the MPI profiling interface, an implementation of the MPI functions *must*

1. provide a mechanism through which all of the MPI defined functions, except those allowed as macros (See Section 2.6.4), may be accessed with a name shift. This requires, in C and Fortran, an alternate entry point name, with the prefix `PMPI_` for each MPI function in each provided language binding and language support method. For routines implemented as macros, it is still required that the `PMPI_` version be supplied and work as expected, but it is not possible to replace at link time the `MPI_` version with a user-defined version.

For Fortran, the different support methods cause several specific procedure names. Therefore, several profiling routines (with these specific procedure names) are needed for each Fortran MPI routine, as described in Section 19.1.5.

2. ensure that those MPI functions that are not replaced may still be linked into an executable image without causing name clashes.
3. document the implementation of different language bindings of the MPI interface if they are layered on top of each other, so that the profiler developer knows whether to implement the profile interface for each binding, or to economize by implementing it only for the lowest level routines.
4. where the implementation of different language bindings is done through a layered approach (e.g., the Fortran binding is a set of “wrapper” functions that call the C

1 implementation), ensure that these wrapper functions are separable from the rest of
2 the library.

3 This separability is necessary to allow a separate profiling library to be correctly
4 implemented, since (at least with Unix linker semantics) the profiling library must
5 contain these wrapper functions if it is to perform as expected. This requirement
6 allows the person who builds the profiling library to extract these functions from the
7 original MPI library and add them into the profiling library without bringing along
8 any other unnecessary code.

- 9
10 5. provide a no-op routine `MPI_PCONTROL` in the MPI library.

11 15.2.2 Discussion

12 The objective of the MPI profiling interface is to ensure that it is relatively easy for authors
13 of profiling (and other similar) tools to interface their codes to MPI implementations on
14 different machines.

15 Since MPI is a machine independent standard with many different implementations,
16 it is unreasonable to expect that the authors of profiling tools for MPI will have access to
17 the source code that implements MPI on any particular machine. It is therefore necessary
18 to provide a mechanism by which the implementors of such tools can collect whatever
19 performance information they wish *without* access to the underlying implementation.

20 We believe that having such an interface is important if MPI is to be attractive to end
21 users, since the availability of many different tools will be a significant factor in attracting
22 users to the MPI standard.

23 The profiling interface is just that, an interface. It says *nothing* about the way in which
24 it is used. There is therefore no attempt to lay down what information is collected through
25 the interface, or how the collected information is saved, filtered, or displayed.

26 While the initial impetus for the development of this interface arose from the desire to
27 permit the implementation of profiling tools, it is clear that an interface like that specified
28 may also prove useful for other purposes, such as “internetworking” multiple MPI imple-
29 mentations. Since all that is defined is an interface, there is no objection to it being used
30 wherever it is useful.

31 As the issues being addressed here are intimately tied up with the way in which ex-
32 ecutable images are built, which may differ greatly on different machines, the examples
33 given below should be treated solely as one way of implementing the objective of the MPI
34 profiling interface. The actual requirements made of an implementation are those detailed
35 in the Requirements section above, the whole of the rest of this section is only present as
36 justification and discussion of the logic for those requirements.

37 The examples below show one way in which an implementation could be constructed to
38 meet the requirements on a Unix system (there are doubtless others that would be equally
39 valid).

40 15.2.3 Logic of the Design

41 Provided that an MPI implementation meets the requirements above, it is possible for
42 the implementor of the profiling system to intercept the MPI calls that are made by the
43 user program. The profiling system implementor can then collect any required information
44 before calling the underlying MPI implementation (through its name shifted entry points)
45 to achieve the desired effects.

15.2.4 Miscellaneous Control of Profiling

There is a clear requirement for the user code to be able to control the profiler dynamically at run time. This capability is normally used for (at least) the purposes of

- Enabling and disabling profiling depending on the state of the calculation.
- Flushing trace buffers at noncritical points in the calculation.
- Adding user events to a trace file.

These requirements are met by use of `MPI_PCONTROL`.

`MPI_PCONTROL(level, ...)`

IN level Profiling level (integer)

C binding

`int MPI_Pcontrol(const int level, ...)`

Fortran 2008 binding

`MPI_Pcontrol(level)`
`INTEGER, INTENT(IN) :: level`

Fortran binding

`MPI_PCONTROL(LEVEL)`
`INTEGER LEVEL`

MPI libraries themselves make no use of this routine, and simply return immediately to the user code. However the presence of calls to this routine allows a profiling package to be explicitly called by the user.

Since MPI has no control of the implementation of the profiling code, we are unable to specify precisely the semantics that will be provided by calls to `MPI_PCONTROL`. This vagueness extends to the number of arguments to the function, and their datatypes.

However to provide some level of portability of user codes to different profiling libraries, we request the following meanings for certain values of `level`.

- `level=0` Profiling is disabled.
- `level=1` Profiling is enabled at a normal default level of detail.
- `level=2` Profile buffers are flushed, which may be a no-op in some profilers.
- All other values of `level` have profile library defined effects and additional arguments.

We also request that the default state after MPI has been initialized is for profiling to be enabled at the normal default level. (i.e., as if `MPI_PCONTROL` had just been called with the argument 1). This allows users to link with a profiling library and to obtain profile output without having to modify their source code at all.

The provision of `MPI_PCONTROL` as a no-op in the standard MPI library supports the collection of more detailed profiling information with source code that can still link against the standard MPI library.

Example 15.1. A wrapper to accumulate the total amount of data sent by the `MPI_SEND` function, along with the total elapsed time spent in the function.

```

1  static int totalBytes = 0;
2  static double totalTime = 0.0;
3
4  int MPI_Send(const void* buffer, int count, MPI_Datatype datatype,
5              int dest, int tag, MPI_Comm comm)
6  {
7      double tstart = MPI_Wtime();      /* Pass on all arguments */
8      int size;
9      int result = PMPI_Send(buffer, count, datatype, dest, tag, comm);
10
11     totalTime += MPI_Wtime() - tstart; /* Compute time */
12
13     MPI_Type_size(datatype, &size);  /* and size */
14     totalBytes += count * size;
15
16     return result;
17 }
18
19

```

15.2.5 MPI Library Implementation

If the MPI library is implemented in C on a Unix system, then there are various options, including the two presented here, for supporting the name-shift requirement. The choice between these two options depends partly on whether the linker and compiler support weak symbols.

If the compiler and linker support weak external symbols, then only a single library is required as the following example shows:

Example 15.2. Library implementation using weak symbols.

```

30 #pragma weak MPI_Example = PMPI_Example
31
32 int PMPI_Example(/* appropriate args */)
33 {
34     /* Useful content */
35 }
36

```

The effect of this `#pragma` is to define the external symbol `MPI_Example` as a weak definition. This means that the linker will not complain if there is another definition of the symbol (for instance in the profiling library); however if no other definition exists, then the linker will use the weak definition.

In the absence of weak symbols then one possible solution would be to use the C macro preprocessor as the following example shows:

Example 15.3. Library implementation using C pre-processor macros.

```

45 #ifdef PROFILELIB
46 #     ifdef __STDC__
47 #         define FUNCTION(name) P ## name
48 #     else

```

```

1 #      define FUNCTION(name) P /**/ name
2 #      endif
3 #else
4 #      define FUNCTION(name) name
5 #endif

```

Each of the user visible functions in the library would then be declared thus

```

6
7 int FUNCTION(MPI_Example)(/* appropriate args */)
8 {
9     /* Useful content */
10 }
11

```

The same source file can then be compiled to produce both versions of the library, depending on the state of the PROFILELIB macro symbol.

It is required that the standard MPI library be built in such a way that the inclusion of MPI functions can be achieved one at a time. This may mean that each external function must reside in its own compilation unit. This is necessary so that the author of the profiling library need only define those MPI functions that need to be intercepted, references to any others being fulfilled by the normal MPI library.

Example 15.4.

The following example shows a potential link step when using the profiling interface.

```

12 % cc ... -lmyprof -lpmpi -lmpi
13

```

Here libmyprof.a contains the profiler functions that intercept some of the MPI functions, libpmpi.a contains the “name shifted” MPI functions, and libmpi.a contains the normal definitions of the MPI functions.

15.2.6 Complications

Multiple Counting

Since parts of the MPI library may themselves be implemented using more basic MPI functions (e.g., a portable implementation of the collective operations implemented using point-to-point communications), there is potential for profiling functions to be called from within an MPI function that was called from a profiling function. This could lead to “double counting” of the time spent in the inner routine. Since this effect could actually be useful under some circumstances (e.g., it might allow one to answer the question “How much time is spent in the point-to-point routines when they are called from collective functions?”), we have decided not to enforce any restrictions on the author of the MPI library that would overcome this. Therefore the author of the profiling library should be aware of this problem, and guard against it. In a single-threaded world this is easily achieved through use of a static variable in the profiling code that remembers if you are already inside a profiling routine. It becomes more complex in a multithreaded environment (as does the meaning of the times recorded).

Linker Oddities

The Unix linker traditionally operates in one pass: the effect of this is that functions from libraries are only included in the image if they are needed at the time the library is scanned.

1 When combined with weak symbols, or multiple definitions of the same function, this can
2 cause odd (and unexpected) effects.

3 Consider, for instance, an implementation of MPI in which the Fortran binding is
4 achieved by using wrapper functions on top of the C implementation. The author of the
5 profile library then assumes that it is reasonable only to provide profile functions for the C
6 binding, since Fortran will eventually call these, and the cost of the wrappers is assumed
7 to be small. However, if the wrapper functions are not in the profiling library, then none
8 of the profiled entry points will be undefined when the profiling library is called. Therefore
9 none of the profiling code will be included in the image. When the standard MPI library
10 is scanned, the Fortran wrappers will be resolved, and will also pull in the base versions of
11 the MPI functions. The overall effect is that the code will link successfully, but will not be
12 profiled.

13 To overcome this we must ensure that the Fortran wrapper functions are included in
14 the profiling version of the library. We ensure that this is possible by requiring that these
15 be separable from the rest of the base MPI library. This allows them to be copied out of
16 the base library and into the profiling one using a tool such as `ar`.

18 *Fortran Support Methods*

19 The different Fortran support methods and possible options for the support of subarrays
20 (depending on whether the compiler can support `TYPE(*)`, `DIMENSION(..)` choice buffers)
21 imply different specific procedure names for the same Fortran MPI routine. The rules and
22 implications for the profiling interface are described in Section [19.1.5](#).

24 15.2.7 Multiple Levels of Interception

25 The scheme given here does not directly support the nesting of profiling functions, since it
26 provides only a single alternative name for each MPI function. Consideration was given to
27 an implementation that would allow multiple levels of call interception, however we were
28 unable to construct an implementation of this that did not have the following disadvantages

- 29 • assuming a particular implementation language, and
- 30 • imposing a run time cost even when no profiling was taking place.

31 Since one of the objectives of MPI is to permit efficient, low latency implementations, and
32 it is not the business of a standard to require a particular implementation language, we
33 decided to accept the scheme outlined above.

34 Note, however, that it is possible to use the scheme above to implement a multi-level
35 system, since the function called by the user may call many different profiling functions
36 before calling the underlying MPI function. This capability has been demonstrated in the
37 P^N MPI tool infrastructure [\[59\]](#).

42 15.3 The MPI Tool Information Interface

43 MPI implementations often use internal variables to control their behavior and performance
44 and rely on internal events for their implementation. Understanding and manipulating these
45 variables and tracking these events can provide a more efficient execution environment or
46 improve performance for many applications. This section describes the MPI tool information
47
48

interface, which provides a mechanism for MPI implementors to expose variables, each of which represents a particular property, setting, or performance measurement from within the MPI implementation, as well as expose events that can be tracked by tools. The interface is split into three parts: the first part provides information about, and supports the setting of, control variables through which the MPI implementation tunes its configuration. The second part provides access to performance variables that can provide insight into internal performance information of the MPI implementation. The third part enables tools to query available events within an MPI implementation and register callbacks for them.

To avoid restrictions on the MPI implementation, the MPI tool information interface allows the implementation to specify which control variables, performance variables, and events exist. Additionally, the user of the MPI tool information interface can obtain meta-data about each available variable or event, such as its datatype, and a textual description. The MPI tool information interface provides the necessary routines to find all variables and events that exist in a particular MPI implementation; to query their properties; to retrieve descriptions about their meaning; to access and, if appropriate, to alter their values; and (in case of events) set callbacks triggered by them.

Variables, events, and categories across connected MPI processes with equivalent names are required to have the same meaning (see the definition of “equivalent” as related to strings in Section 15.3.3). Furthermore, enumerations with equivalent names across connected MPI processes are required to have the same meaning, but are allowed to comprise different enumeration items. Enumeration items that have equivalent names across connected MPI processes in enumerations with the same meaning must also have the same meaning. In order for variables and categories to have the same meaning, routines in the tools information interface that return details for those variables and categories have requirements on what parameters must be identical. These requirements are specified in their respective sections.

Rationale. The intent of requiring the same meaning for entities with equivalent names is to enforce consistency across connected MPI processes. For example, variables describing the number of packets sent on different types of network devices should have different names to reflect their potentially different meanings. (*End of rationale.*)

The MPI tool information interface can be used independently from the MPI communication functionality. In particular, the routines of this interface can be called before MPI is initialized and after MPI is finalized. In order to support this behavior cleanly, the MPI tool information interface uses separate initialization and finalization routines. All identifiers used in the MPI tool information interface have the prefix `MPI_T_`.

On success, all MPI tool information interface routines return `MPI_SUCCESS`, otherwise they return an appropriate and unique return code indicating the reason why the call was not successfully completed. Details on return codes can be found in Section 15.3.10. However, unsuccessful calls to the MPI tool information interface are not fatal and do not impact the execution of subsequent MPI routines.

Since the MPI tool information interface primarily focuses on tools and support libraries, MPI implementations are only required to provide C bindings for functions and constants introduced in this section. Except where otherwise noted, all conventions and principles governing the C bindings of the MPI API also apply to the MPI tool information interface, which is available by including the `mpi.h` header file. All routines in this interface have local semantics.

Table 15.1: MPI tool information interface verbosity levels

MPI_T_VERBOSITY_USER_BASIC	Basic information of interest to users
MPI_T_VERBOSITY_USER_DETAIL	Detailed information of interest to users
MPI_T_VERBOSITY_USER_ALL	All remaining information of interest to users
MPI_T_VERBOSITY_TUNER_BASIC	Basic information required for tuning
MPI_T_VERBOSITY_TUNER_DETAIL	Detailed information required for tuning
MPI_T_VERBOSITY_TUNER_ALL	All remaining information required for tuning
MPI_T_VERBOSITY_MPIDEV_BASIC	Basic information for MPI implementors
MPI_T_VERBOSITY_MPIDEV_DETAIL	Detailed information for MPI implementors
MPI_T_VERBOSITY_MPIDEV_ALL	All remaining information for MPI implementors

Advice to users. The number and type of control variables, performance variables, and events can vary between MPI implementations, platforms and different builds of the same implementation on the same platform as well as between runs. Hence, any application relying on a particular variable will not be portable. Further, there is no guarantee that the number of variables and variable indices are the same across connected MPI processes.

This interface is primarily intended for performance monitoring tools, support tools, and libraries controlling the application's environment. When maximum portability is desired, application programmers should either avoid using the MPI tool information interface or avoid being dependent on the existence of a particular control or performance variable or of a particular event. (*End of advice to users.*)

15.3.1 Verbosity Levels

The MPI tool information interface provides access to internal configuration and performance information through a set of control and performance variables defined by the MPI implementation. Since some implementations may export a large number of variables, variables are classified by a verbosity level that categorizes both their intended audience (end users, performance tuners or MPI implementors) and a relative measure of level of detail (basic, detailed or all). These verbosity levels are described by a single integer. Table 15.1 lists the constants for all possible verbosity levels. The values of the constants are monotonic in the order listed in the table; i.e., `MPI_T_VERBOSITY_USER_BASIC < MPI_T_VERBOSITY_USER_DETAIL < ... < MPI_T_VERBOSITY_MPIDEV_ALL`.

15.3.2 Binding MPI Tool Information Interface Variables to MPI Objects

Each MPI tool information interface variable provides access to a particular control setting or performance property of the MPI implementation. A variable may refer to a specific MPI object such as a communicator, datatype, or one-sided communication window, or the variable may refer more generally to the MPI environment of the process. Except for the last case, the variable must be bound to exactly one MPI object before it can be used. Table 15.2 lists all MPI object types to which an MPI tool information interface variable can be bound, together with the matching constant that MPI tool information interface routines return to identify the object type. It is erroneous to bind a control variable, performance variable, or

event to a handle that would not be valid to use as an input argument to another MPI call (excluding calls to the MPI Tool Information Interface) at the same point of execution.

Table 15.2: Constants to identify associations of variables

Constant	MPI object
MPI_T_BIND_NO_OBJECT	N/A; applies globally to entire MPI process
MPI_T_BIND_MPI_COMM	MPI communicators
MPI_T_BIND_MPI_DATATYPE	MPI datatypes
MPI_T_BIND_MPI_ERRHANDLER	MPI error handlers
MPI_T_BIND_MPI_FILE	MPI file handles
MPI_T_BIND_MPI_GROUP	MPI groups
MPI_T_BIND_MPI_OP	MPI reduction operators
MPI_T_BIND_MPI_REQUEST	MPI requests
MPI_T_BIND_MPI_WIN	MPI windows for one-sided communication
MPI_T_BIND_MPI_MESSAGE	MPI message object
MPI_T_BIND_MPI_INFO	MPI info object
MPI_T_BIND_MPI_SESSION	MPI session object

Rationale. Some variables have meanings tied to a specific MPI object. Examples include the number of send or receive operations that use a particular datatype, the number of times a particular error handler has been called, or the communication protocol and “eager limit” used for a particular communicator. Creating a new MPI tool information interface variable for each MPI object would cause the number of variables to grow without bound, since they cannot be reused to avoid naming conflicts. By associating MPI tool information interface variables with a specific MPI object, the MPI implementation only must specify and maintain a single variable, which can then be applied to as many MPI objects of the respective type as created during the program’s execution. (*End of rationale.*)

15.3.3 Convention for Returning Strings

Several MPI tool information interface functions return one or more strings. These functions have two arguments for each string to be returned: an OUT parameter that identifies a pointer to the buffer in which the string will be returned, and an INOUT parameter to pass the length of the buffer. The user is responsible for the memory allocation of the buffer and must pass the size of the buffer (n) as the length argument. Let n be the length value specified to the function. On return, the function writes at most $n - 1$ of the string’s characters into the buffer, followed by a null terminator. If the returned string’s length is greater than or equal to n , the string will be truncated to $n - 1$ characters. In this case, the length of the string plus one (for the terminating null character) is returned in the length argument. If the user passes the null pointer as the buffer argument or passes 0 as the length argument, the function does not return the string and only returns the length of the string plus one in the length argument. If the user passes the null pointer as the length argument, the buffer argument is ignored and nothing is returned.

MPI implementations behave as if they have an internal character array that is copied to the output character array supplied by the user. Such output strings are only defined

to be equivalent if their notional source-internal character arrays are identical (up to and including the null terminator), even if the output string is truncated due to a small input length parameter n .

15.3.4 Initialization and Finalization

The MPI tool information interface requires a separate set of initialization and finalization routines.

`MPI_T_INIT_THREAD(required, provided)`

IN	required	desired level of thread support (integer)
OUT	provided	provided level of thread support (integer)

C binding

`int MPI_T_init_thread(int required, int *provided)`

All programs or tools that use the MPI tool information interface must initialize the MPI tool information interface in the processes that will use the interface before calling any other of its routines. A user can initialize the MPI tool information interface by calling `MPI_T_INIT_THREAD`, which can be called multiple times. In addition, this routine initializes the thread environment for all routines in the MPI tool information interface. Calling this routine when the MPI tool information interface is already initialized has no effect beyond increasing the reference count of how often the interface has been initialized. The argument `required` is used to specify the desired level of thread support. The possible values and their semantics are identical to the ones that can be used with `MPI_INIT_THREAD` listed in Section 11.6. The call returns in `provided` information about the actual level of thread support that will be provided by the MPI implementation for calls to MPI tool information interface routines. It can be one of the four values listed in Section 11.6.

The MPI specification does not require all MPI processes to exist before MPI is initialized. If the MPI tool information interface is used before initialization of MPI, the user is responsible for ensuring that the MPI tool information interface is initialized on all processes it is used in. Processes created by the MPI implementation during initialization inherit the status of the MPI tool information interface (whether it is initialized or not as well as all active sessions and handles) from the process from which they are created.

Processes created at runtime as a result of calls to MPI's dynamic process management require their own initialization before they can use the MPI tool information interface.

Advice to users. If `MPI_T_INIT_THREAD` is called before `MPI_INIT_THREAD`, the requested and provided thread level for `MPI_T_INIT_THREAD` may influence the behavior and return value of `MPI_INIT_THREAD`. The same is true for the reverse order. Likewise, when using the Sessions Model (Section 11.3), the requested and provided thread level for `MPI_T_INIT_THREAD` may influence the behavior and return values of `MPI_SESSION_INIT` (see Section 11.3), with the same being true for the reverse order. (*End of advice to users.*)

Advice to implementors. MPI implementations should strive to make as many control or performance variables available before MPI initialization (instead of adding them

during initialization) to allow tools the most flexibility. In particular, control variables should be available before MPI initialization if their value cannot be changed after MPI initialization. (*End of advice to implementors.*)

MPI_T_FINALIZE()

C binding

int MPI_T_finalize(void)

This routine finalizes the use of the MPI tool information interface and may be called as often as the corresponding MPI_T_INIT_THREAD routine up to the current point of execution. Calling it more times returns a corresponding return code. As long as the number of calls to MPI_T_FINALIZE is smaller than the number of calls to MPI_T_INIT_THREAD up to the current point of execution, the MPI tool information interface remains initialized and calls to its routines are permissible. Further, additional calls to MPI_T_INIT_THREAD after one or more calls to MPI_T_FINALIZE are permissible.

Once MPI_T_FINALIZE is called the same number of times as the routine MPI_T_INIT_THREAD up to the current point of execution, the MPI tool information interface is no longer initialized. The user can reinitialize the interface by a subsequent call to MPI_T_INIT_THREAD.

At the end of the program execution, unless MPI_ABORT is called, an application must have called MPI_T_INIT_THREAD and MPI_T_FINALIZE an equal number of times.

15.3.5 Datatype System

All variables managed through the MPI tool information interface represent their values through typed buffers of a given length and type using an MPI datatype (similar to regular send/receive buffers). Since the initialization of the MPI tool information interface is separate from the initialization of MPI, MPI tool information interface routines can be called before MPI initialization. Consequently, these routines can also use MPI datatypes before MPI initialization. Therefore, within the context of the MPI tool information interface, it is permissible to use a subset of MPI datatypes as specified below before MPI initialization.

Table 15.3: MPI datatypes that can be used by the MPI tool information interface

MPI_INT
MPI_INT32_T
MPI_INT64_T
MPI_UNSIGNED
MPI_UNSIGNED_LONG
MPI_UNSIGNED_LONG_LONG
MPI_UINT32_T
MPI_UINT64_T
MPI_COUNT
MPI_CHAR
MPI_DOUBLE

Rationale. The MPI tool information interface relies mainly on unsigned datatypes for integer values since most variables are expected to represent counters or resource sizes. MPI_INT is provided for additional flexibility and is expected to be used mainly for control variables and enumeration types (see below).

Providing all basic datatypes, in particular providing all signed and unsigned variants of integer types, would lead to a larger number of types, which tools need to interpret. This would cause unnecessary complexity in the implementation of tools based on the MPI tool information interface. (*End of rationale.*)

The MPI tool information interface only relies on a subset of the basic MPI datatypes and does not use any derived MPI datatypes. Table 15.3 lists all MPI datatypes that can be returned by the MPI tool information interface to represent its variables.

The use of the datatype MPI_CHAR in the MPI tool information interface implies a null-terminated character array, i.e., a string in the C language. If a variable has type MPI_CHAR, the value of the count parameter returned by MPI_T_CVAR_HANDLE_ALLOC and MPI_T_PVAR_HANDLE_ALLOC must be large enough to include any valid value, including its terminating null character. The contents of returned MPI_CHAR arrays are only defined from index 0 through the location of the first null character.

Rationale. The MPI tool information interface requires a significantly simpler type system than MPI itself. Therefore, only its required subset must be present before MPI initialization and MPI implementations do not need to initialize the complete MPI datatype system. (*End of rationale.*)

For variables of type MPI_INT, an MPI implementation can provide additional information by associating names with a fixed number of values. We refer to this information in the following as an enumeration. In this case, the respective calls that provide additional metadata for each control or performance variable, i.e., MPI_T_CVAR_GET_INFO (Section 15.3.6), MPI_T_PVAR_GET_INFO (Section 15.3.7), and MPI_T_EVENT_GET_INFO (Section 15.3.8), return a handle of type MPI_T_enum that can be passed to the following functions to extract additional information. Thus, the MPI implementation can describe variables with a fixed set of values that each represents a particular state. Each enumeration type can have N different values, with a fixed N that can be queried using MPI_T_ENUM_GET_INFO.

MPI_T_ENUM_GET_INFO(enumtype, num, name, name_len)

IN	enumtype	enumeration to be queried (handle)
OUT	num	number of discrete values represented by this enumeration (integer)
OUT	name	buffer to return the string containing the name of the enumeration item (string)
INOUT	name_len	length of the string and/or buffer for name (integer)

C binding

```
int MPI_T_enum_get_info(MPI_T_enum enumtype, int *num, char *name,
                       int *name_len)
```

If `enumtype` is a valid enumeration, this routine returns the number of items represented by this enumeration type as well as its name. N must be greater than 0, i.e., the enumeration must represent at least one value.

The arguments `name` and `name_len` are used to return the name of the enumeration as described in Section 15.3.3.

The routine is required to return a name of at least length one. This name must be unique with respect to all other names for enumerations that the MPI implementation uses.

Names associated with individual values in each enumeration `enumtype` can be queried using `MPI_T_ENUM_GET_ITEM`.

`MPI_T_ENUM_GET_ITEM(enumtype, index, value, name, name_len)`

IN	<code>enumtype</code>	enumeration to be queried (handle)
IN	<code>index</code>	number of the value to be queried in this enumeration (integer)
OUT	<code>value</code>	variable value (integer)
OUT	<code>name</code>	buffer to return the string containing the name of the enumeration item (string)
INOUT	<code>name_len</code>	length of the string and/or buffer for <code>name</code> (integer)

C binding

```
int MPI_T_enum_get_item(MPI_T_enum enumtype, int index, int *value, char *name,
                       int *name_len)
```

The arguments `name` and `name_len` are used to return the name of the enumeration item as described in Section 15.3.3.

If completed successfully, the routine returns the name/value pair that describes the enumeration at the specified index. The call is further required to return a name of at least length one. This name must be unique with respect to all other names of items for the same enumeration.

15.3.6 Control Variables

The routines described in this section of the MPI tool information interface specification focus on the ability to list, query, and possibly set control variables exposed by the MPI implementation. These variables can typically be used by the user to fine tune properties and configuration settings of the MPI implementation. On many systems, such variables can be set using environment variables, although other configuration mechanisms may be available, such as configuration files or central configuration registries. A typical example that is available in several existing MPI implementations is the ability to specify an “eager limit,” i.e., an upper bound on the size of messages sent or received using an eager protocol.

Control Variable Query Functions

An MPI implementation exports a set of N control variables through the MPI tool information interface. If N is zero, then the MPI implementation does not export any control variables, otherwise the provided control variables are indexed from 0 to $N - 1$. This index number is used in subsequent calls to identify the individual variables.

1 An MPI implementation is allowed to increase the number of control variables during
 2 the execution of an MPI application when new variables become available through dynamic
 3 loading. However, MPI implementations are not allowed to change the index of a control
 4 variable or to delete a variable once it has been added to the set. When a variable becomes
 5 inactive, e.g., through dynamic unloading, accessing its value should return a corresponding
 6 return code.

7
 8 *Advice to users.* While the MPI tool information interface guarantees that indices or
 9 variable properties do not change during a particular run of an MPI program, it does
 10 not provide a similar guarantee between runs. (*End of advice to users.*)

11 The following function can be used to query the number of control variables, `num_cvar`:

12
 13
 14 `MPI_T_CVAR_GET_NUM(num_cvar)`

15 OUT `num_cvar` returns number of control variables (integer)

17 C binding

18 `int MPI_T_cvar_get_num(int *num_cvar)`

19 The function `MPI_T_CVAR_GET_INFO` provides access to additional information for
 20 each variable.
 21

22
 23
 24 `MPI_T_CVAR_GET_INFO(cvar_index, name, name_len, verbosity, datatype, enumtype, desc,`
 25 `desc_len, bind, scope)`

26 IN `cvar_index` index of the control variable to be queried, value
 27 between 0 and `num_cvar - 1` (integer)

28 OUT `name` buffer to return the string containing the name of the
 29 control variable (string)

30 INOUT `name_len` length of the string and/or buffer for `name` (integer)

31 OUT `verbosity` verbosity level of this variable (integer)

32 OUT `datatype` MPI datatype of the information stored in the
 33 control variable (handle)

34 OUT `enumtype` optional descriptor for enumeration information
 35 (handle)

36 OUT `desc` buffer to return the string containing a description of
 37 the control variable (string)

38 INOUT `desc_len` length of the string and/or buffer for `desc` (integer)

39 OUT `bind` type of MPI object to which this variable must be
 40 bound (integer)

41 OUT `scope` scope of when changes to this variable are possible
 42 (integer)

43 C binding

44 `int MPI_T_cvar_get_info(int cvar_index, char *name, int *name_len,`

```
int *verbosity, MPI_Datatype *datatype, MPI_T_enum *enumtype,
char *desc, int *desc_len, int *bind, int *scope)
```

After a successful call to `MPI_T_CVAR_GET_INFO` for a particular variable, subsequent calls to this routine that query information about the same variable must return the same information. An MPI implementation is not allowed to alter any of the returned values.

If any OUT parameter to `MPI_T_CVAR_GET_INFO` is a NULL pointer, the implementation will ignore the parameter and not return a value for the parameter.

The arguments `name` and `name_len` are used to return the name of the control variable as described in Section 15.3.3.

If completed successfully, the routine is required to return a name of at least length one. The name must be unique with respect to all other names for control variables used by the MPI implementation.

The argument `verbosity` returns the verbosity level of the variable (see Section 15.3.1).

The argument `datatype` returns the MPI datatype that is used to represent the control variable.

If the variable is of type `MPI_INT`, MPI can optionally specify an enumeration for the values represented by this variable and return it in `enumtype`. In this case, MPI returns an enumeration identifier, which can then be used to gather more information as described in Section 15.3.5. Otherwise, `enumtype` is set to `MPI_T_ENUM_NULL`. If the datatype is not `MPI_INT` or the argument `enumtype` is the null pointer, no enumeration type is returned.

The arguments `desc` and `desc_len` are used to return a description of the control variable as described in Section 15.3.3.

Returning a description is optional. If an MPI implementation does not return a description, the first character for `desc` must be set to the null character and `desc_len` must be set to one at the return of this call.

The parameter `bind` returns the type of the MPI object to which the variable must be bound or the value `MPI_T_BIND_NO_OBJECT` (see Section 15.3.2).

The scope of a variable determines whether changing a variable's value is either local to the MPI process or must be done by the user across multiple connected MPI processes. The latter is further split into variables that require changes in a group of MPI processes and those that require collective changes among all connected MPI processes. Both cases can require variables on all participating MPI processes either to be set to consistent (but potentially different) values or to equal values. The description provided with the variable must contain an explanation about the requirements and/or restrictions for setting the particular variable.

On successful return from `MPI_T_CVAR_GET_INFO`, the argument `scope` will be set to one of the constants listed in Table 15.4.

If the name of a control variable is equivalent across connected MPI processes, the following OUT parameters must be identical: `verbosity`, `datatype`, `enumtype`, `bind`, and `scope`. The returned description must be equivalent.

Advice to users. The `scope` of a variable only indicates if a variable might be changeable; it is not a guarantee that it can be changed at any time. (*End of advice to users.*)

Table 15.4: Scopes for control variables

Scope Constant	Description
MPI_T_SCOPE_CONSTANT	read-only, value is constant
MPI_T_SCOPE_READONLY	read-only, cannot be written, but can change
MPI_T_SCOPE_LOCAL	may be writeable, writing only affects the calling MPI process
MPI_T_SCOPE_GROUP	may be writeable, must be set to consistent values across a group of connected MPI processes
MPI_T_SCOPE_GROUP_EQ	may be writeable, must be set to the same value across a group of connected MPI processes
MPI_T_SCOPE_ALL	may be writeable, must be set to consistent values across all connected MPI processes
MPI_T_SCOPE_ALL_EQ	may be writeable, must be set to the same value across all connected MPI processes

```
MPI_T_CVAR_GET_INDEX(name, cvar_index)
```

IN	name	name of the control variable (string)
OUT	cvar_index	index of the control variable (integer)

C binding

```
int MPI_T_cvar_get_index(const char *name, int *cvar_index)
```

MPI_T_CVAR_GET_INDEX is a function for retrieving the index of a control variable given a known variable name. The name parameter is provided by the caller, and cvar_index is returned by the MPI implementation. The name parameter is a string terminated with a null character.

This routine returns MPI_SUCCESS on success and returns MPI_T_ERR_INVALID_NAME if name does not match the name of any control variable provided by the implementation at the time of the call.

Rationale. This routine is provided to enable fast retrieval of control variables by a tool, assuming it knows the name of the variable for which it is looking. The number of variables exposed by the implementation can change over time, so it is not possible for the tool to simply iterate over the list of variables once at initialization. Although using MPI implementation specific variable names is not portable across MPI implementations, tool developers may choose to take this route for lower overhead at runtime because the tool will not have to iterate over the entire set of variables to find a specific one. (*End of rationale.*)

Example 15.5. Querying and printing the names of all available control variables.

```
#include <stdio.h>
#include <stdlib.h>
#include <mpi.h>

int main(int argc, char *argv[])
```

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

```

{
    int i, err, num, namelen, bind, verbose, scope;
    int threadssupport;
    char name[100];

    MPI_Datatype datatype;

    err = MPI_T_init_thread(MPI_THREAD_SINGLE, &threadssupport);
    if (err != MPI_SUCCESS)
        return err;

    err = MPI_T_cvar_get_num(&num);
    if (err != MPI_SUCCESS)
        return err;

    for (i = 0; i < num; i++) {
        namelen = 100;
        err = MPI_T_cvar_get_info(i, name, &namelen,
                                &verbose, &datatype, NULL,
                                NULL, NULL, /*no description */
                                &bind, &scope);
        if ((err != MPI_SUCCESS) && (err != MPI_T_ERR_INVALID_INDEX))
            return err;

        printf("Var %i: %s\n", i, name);
    }

    err = MPI_T_finalize();
    if (err != MPI_SUCCESS)
        return 1;
    else
        return 0;
}

```

Handle Allocation and Deallocation

Before reading or writing the value of a variable, a user must first allocate a handle of type `MPI_T_cvar_handle` for the variable by binding it to an MPI object (see also Section 15.3.2).

Rationale. Handles used in the MPI tool information interface are distinct from handles used in the remaining parts of the MPI standard because they must be usable before MPI is initialized and after MPI is finalized. Further, accessing handles, in particular for performance variables, can be time critical and having a separate handle space enables optimizations. (*End of rationale.*)

`MPI_T_CVAR_HANDLE_ALLOC(cvar_index, obj_handle, handle, count)`

IN `cvar_index` index of control variable for which handle is to be allocated (index)

1	IN	obj_handle	reference to a handle of the MPI object to which this
2			variable is supposed to be bound (pointer)
3	OUT	handle	allocated handle (handle)
4			
5	OUT	count	number of elements used to represent this variable
6			(integer)

C binding

```

9 int MPI_T_cvar_handle_alloc(int cvar_index, void *obj_handle,
10 MPI_T_cvar_handle *handle, int *count)

```

11 This routine binds the control variable specified by the argument `index` to an MPI object.
12 The object is passed in the argument `obj_handle` as an address to a local variable that stores
13 the object's handle. The argument `obj_handle` is ignored if the `MPI_T_CVAR_GET_INFO`
14 call for this control variable returned `MPI_T_BIND_NO_OBJECT` in the argument `bind`. The
15 handle allocated to reference the variable is returned in the argument `handle`. Upon success-
16 ful return, `count` contains the number of elements (of the datatype returned by a previous
17 `MPI_T_CVAR_GET_INFO` call) used to represent this variable.

19 *Advice to users.* The `count` can be different based on the MPI object to which the
20 control variable was bound. For example, variables bound to communicators could
21 have a count that matches the size of the communicator.

22 It is not portable to pass references to predefined MPI object handles, such as
23 `MPI_COMM_WORLD` to this routine, since their implementation depends on the MPI
24 library. Instead, such object handles should be stored in a local variable and the
25 address of this local variable should be passed into `MPI_T_CVAR_HANDLE_ALLOC`.
26 (*End of advice to users.*)

28 The value of `cvar_index` should be in the range from 0 to `num_cvar - 1`, where `num_cvar`
29 is the number of available control variables as determined from a prior call to
30 `MPI_T_CVAR_GET_NUM`. The type of the MPI object it references must be consistent with
31 the type returned in the `bind` argument in a prior call to `MPI_T_CVAR_GET_INFO`.

```

34 MPI_T_CVAR_HANDLE_FREE(handle)

```

35	INOUT	handle	handle to be freed (handle)
----	-------	--------	-----------------------------

C binding

```

39 int MPI_T_cvar_handle_free(MPI_T_cvar_handle *handle)

```

40 When a handle is no longer needed, a user of the MPI tool information interface should
41 call `MPI_T_CVAR_HANDLE_FREE` to free the handle and the associated resources in the
42 MPI implementation. On a successful return, MPI sets the handle to
43 `MPI_T_CVAR_HANDLE_NULL`.

Control Variable Access Functions

MPI_T_CVAR_READ(handle, buf)

IN	handle	handle to the control variable to be read (handle)
OUT	buf	initial address of storage location for variable value (choice)

C binding

int MPI_T_cvar_read(MPI_T_cvar_handle handle, void *buf)

This routine queries the value of a control variable identified by the argument `handle` and stores the result in the buffer identified by the parameter `buf`. The user must ensure that the buffer is of the appropriate size to hold the entire value of the control variable (based on the returned datatype and count from prior corresponding calls to `MPI_T_CVAR_GET_INFO` and `MPI_T_CVAR_HANDLE_ALLOC`, respectively).

MPI_T_CVAR_WRITE(handle, buf)

INOUT	handle	handle to the control variable to be written (handle)
IN	buf	initial address of storage location for variable value (choice)

C binding

int MPI_T_cvar_write(MPI_T_cvar_handle handle, const void *buf)

This routine sets the value of the control variable identified by the argument `handle` to the data stored in the buffer identified by the parameter `buf`. The user must ensure that the buffer is of the appropriate size to hold the entire value of the control variable (based on the returned datatype and count from prior corresponding calls to `MPI_T_CVAR_GET_INFO` and `MPI_T_CVAR_HANDLE_ALLOC`, respectively).

If the variable has a global scope (as returned by a prior corresponding `MPI_T_CVAR_GET_INFO` call), any write call to this variable must be issued by the user in all connected (as defined in Section 11.10.4) MPI processes. If the variable has group scope, any write call to this variable must be issued by the user in all MPI processes in the group, which must be described by the MPI implementation in the description by the `MPI_T_CVAR_GET_INFO`.

In both cases, the user must ensure that the writes in all participating MPI processes are consistent. If the scope is either `MPI_T_SCOPE_ALL_EQ` or `MPI_T_SCOPE_GROUP_EQ` this means that the variable in all connected MPI processes or MPI processes of the group, respectively, must be set to the same value.

If it is not possible to change the variable at the time the call is made, the function returns either `MPI_T_ERR_CVAR_SET_NOT_NOW`, if there may be a later time at which the variable could be set, or `MPI_T_ERR_CVAR_SET_NEVER`, if the variable cannot be set for the remainder of the application's execution.

Example 15.6. Reading the value of a control variable.

```

1  int getValue_int_comm(int index, MPI_Comm comm, int *val)
2  {
3      int err, count;
4      MPI_T_cvar_handle handle;
5
6      /* This example assumes that the variable index */
7      /* can be bound to a communicator */
8
9
10     err = MPI_T_cvar_handle_alloc(index, &comm, &handle, &count);
11     if (err != MPI_SUCCESS)
12         return err;
13
14     /* The following assumes that the variable is */
15     /* represented by a single integer */
16
17     err = MPI_T_cvar_read(handle, val);
18     if (err != MPI_SUCCESS)
19         return err;
20
21     err = MPI_T_cvar_handle_free(&handle);
22
23     return err;
24 }

```

15.3.7 Performance Variables

The following section focuses on the ability to list and to query performance variables provided by the MPI implementation. Performance variables provide insight into MPI implementation-specific internals and can represent information such as the state of the MPI implementation (e.g., waiting blocked, receiving, not active), aggregated timing data for submodules, or queue sizes and lengths.

Rationale. The interface for performance variables is separate from the interface for control variables, since performance variables have different requirements and parameters. By keeping them separate, the interface provides cleaner semantics and allows for more performance optimization opportunities. (*End of rationale.*)

Some performance variables and classes refer to **events**. In general, such events describe state transitions within software or hardware related to the performance of an MPI application. The events offered through the callback-driven event-notification interface described in Section 15.3.8 also refer to such state transitions; however, the set of state transitions referred to by performance variables and events as described in Section 15.3.8 may not be identical.

Performance Variable Classes

Each performance variable is associated with a class that describes its basic semantics, possible datatypes, basic behavior, its starting value, whether it can overflow, and when and how an MPI implementation can change the variable's value. The starting value is the value that is assigned to the variable the first time that it is used or whenever it is reset.

Advice to users. If a performance variable belongs to a class that can overflow, it is up to the user to protect against this overflow, e.g., by frequently reading and resetting the variable value. (*End of advice to users.*)

Advice to implementors. MPI implementations should use large enough datatypes for each performance variable to avoid overflows under normal circumstances. (*End of advice to implementors.*)

The classes are defined by the following constants:

MPI_T_PVAR_CLASS_STATE: A performance variable in this class represents a set of discrete states. Variables of this class are represented by MPI_INT and can be set by the MPI implementation at any time. Variables of this type should be described further using an enumeration, as discussed in Section 15.3.5. The starting value is the current state of the implementation at the time that the starting value is set. MPI implementations must ensure that variables of this class cannot overflow.

MPI_T_PVAR_CLASS_LEVEL: A performance variable in this class represents a value that describes the utilization level of a resource. The value of a variable of this class can change at any time to match the current utilization level of the resource. Values returned from variables in this class are nonnegative and represented by one of the following datatypes: MPI_UNSIGNED, MPI_UNSIGNED_LONG, MPI_UNSIGNED_LONG_LONG, MPI_DOUBLE. The starting value is the current utilization level of the resource at the time that the starting value is set. MPI implementations must ensure that variables of this class cannot overflow.

MPI_T_PVAR_CLASS_SIZE: A performance variable in this class represents a value that is the size of a resource. Values returned from variables in this class are nonnegative and represented by one of the following datatypes: MPI_UNSIGNED, MPI_UNSIGNED_LONG, MPI_UNSIGNED_LONG_LONG, MPI_DOUBLE. The starting value is the current size of the resource at the time that the starting value is set. MPI implementations must ensure that variables of this class cannot overflow.

MPI_T_PVAR_CLASS_PERCENTAGE: The value of a performance variable in this class represents the percentage utilization of a finite resource. The value of a variable of this class can change at any time to match the current utilization level of the resource. It will be returned as an MPI_DOUBLE datatype. The value must always be between 0.0 (resource not used at all) and 1.0 (resource completely used). The starting value is the current percentage utilization level of the resource at the time that the starting value is set. MPI implementations must ensure that variables of this class cannot overflow.

MPI_T_PVAR_CLASS_HIGHWATERMARK: A performance variable in this class represents a value that describes the maximum observed utilization of a resource. The value of a variable of this class is nonnegative and grows monotonically from the initialization or reset of the variable. It can be represented by one of the following datatypes: MPI_UNSIGNED, MPI_UNSIGNED_LONG, MPI_UNSIGNED_LONG_LONG, MPI_DOUBLE. The starting value is the current utilization level of the resource at the time that the variable is started or reset. MPI implementations must ensure that variables of this class cannot overflow.

1 **MPI_T_PVAR_CLASS_LOWWATERMARK:** A performance variable in this class represents
2 a value that describes the minimum observed utilization of a resource. The value of a
3 variable of this class is nonnegative and decreases monotonically from the initialization
4 or reset of the variable. It can be represented by one of the following datatypes:
5 MPI_UNSIGNED, MPI_UNSIGNED_LONG, MPI_UNSIGNED_LONG_LONG,
6 MPI_DOUBLE. The starting value is the current utilization level of the resource at the
7 time that the variable is started or reset. MPI implementations must ensure that
8 variables of this class cannot overflow.

9
10 **MPI_T_PVAR_CLASS_COUNTER:** A performance variable in this class counts the number
11 of occurrences of a specific event (e.g., the number of memory allocations within an
12 MPI library). The value of a variable of this class increases monotonically from the
13 initialization or reset of the performance variable by one for each specific event that
14 is observed. Values must be nonnegative and represented by one of the following
15 datatypes: MPI_UNSIGNED, MPI_UNSIGNED_LONG, MPI_UNSIGNED_LONG_LONG. The
16 starting value for variables of this class is 0. Variables of this class can overflow.

17 **MPI_T_PVAR_CLASS_AGGREGATE:** The value of a performance variable in this class is an
18 an aggregated value that represents a sum of arguments processed during a specific
19 event (e.g., the amount of memory allocated by all memory allocations). This class is
20 similar to the counter class, but instead of counting individual events, the value can
21 be incremented by arbitrary amounts. The value of a variable of this class increases
22 monotonically from the initialization or reset of the performance variable. It must
23 be nonnegative and represented by one of the following datatypes: MPI_UNSIGNED,
24 MPI_UNSIGNED_LONG, MPI_UNSIGNED_LONG_LONG, MPI_DOUBLE. The starting value
25 for variables of this class is 0. Variables of this class can overflow.

26
27 **MPI_T_PVAR_CLASS_TIMER:** The value of a performance variable in this class represents
28 the aggregated time that the MPI implementation spends executing a particular event,
29 type of event, or section of the MPI library. This class has the same basic semantics
30 as MPI_T_PVAR_CLASS_AGGREGATE, but explicitly records a timing value. The value
31 of a variable of this class increases monotonically from the initialization or reset of the
32 performance variable. It must be nonnegative and represented by one of the following
33 datatypes: MPI_UNSIGNED, MPI_UNSIGNED_LONG, MPI_UNSIGNED_LONG_LONG,
34 MPI_DOUBLE. The starting value for variables of this class is 0. If the type
35 MPI_DOUBLE is used, the units that represent time in this datatype must match the
36 units used by MPI_WTIME. Otherwise, the time units should be documented, e.g.,
37 in the description returned by MPI_T_PVAR_GET_INFO. Variables of this class can
38 overflow.

39
40 **MPI_T_PVAR_CLASS_GENERIC:** This class can be used to describe a variable that does
41 not fit into any of the other classes. For variables in this class, the starting value is
42 variable-specific and implementation-defined.

43 *Performance Variable Query Functions*

44
45 An MPI implementation exports a set of N performance variables through the MPI tool
46 information interface. If N is zero, then the MPI implementation does not export any
47 performance variables; otherwise the provided performance variables are indexed from 0 to
48 $N - 1$. This index number is used in subsequent calls to identify the individual variables.

An MPI implementation is allowed to increase the number of performance variables during the execution of an MPI application when new variables become available through dynamic loading. However, MPI implementations are not allowed to change the index of a performance variable or to delete a variable once it has been added to the set. When a variable becomes inactive, e.g., through dynamic unloading, accessing its value should return a corresponding return code.

The following function can be used to query the number of performance variables, `num_pvar`:

`MPI_T_PVAR_GET_NUM(num_pvar)`

OUT `num_pvar` returns number of performance variables (integer)

C binding

`int MPI_T_pvar_get_num(int *num_pvar)`

The function `MPI_T_PVAR_GET_INFO` provides access to additional information for each variable.

`MPI_T_PVAR_GET_INFO(pvar_index, name, name_len, verbosity, var_class, datatype, enumtype, desc, desc_len, bind, readonly, continuous, atomic)`

IN	<code>pvar_index</code>	index of the performance variable to be queried between 0 and <code>num_pvar - 1</code> (integer)
OUT	<code>name</code>	buffer to return the string containing the name of the performance variable (string)
INOUT	<code>name_len</code>	length of the string and/or buffer for <code>name</code> (integer)
OUT	<code>verbosity</code>	verbosity level of this variable (integer)
OUT	<code>var_class</code>	class of performance variable (integer)
OUT	<code>datatype</code>	MPI datatype of the information stored in the performance variable (handle)
OUT	<code>enumtype</code>	optional descriptor for enumeration information (handle)
OUT	<code>desc</code>	buffer to return the string containing a description of the performance variable (string)
INOUT	<code>desc_len</code>	length of the string and/or buffer for <code>desc</code> (integer)
OUT	<code>bind</code>	type of MPI object to which this variable must be bound (integer)
OUT	<code>readonly</code>	flag indicating whether the variable can be written/reset (integer)
OUT	<code>continuous</code>	flag indicating whether the variable can be started and stopped or is continuously active (integer)
OUT	<code>atomic</code>	flag indicating whether the variable can be atomically read and reset (integer)

C binding

```

1 int MPI_T_pvar_get_info(int pvar_index, char *name, int *name_len,
2     int *verbosity, int *var_class, MPI_Datatype *datatype,
3     MPI_T_enum *enumtype, char *desc, int *desc_len, int *bind,
4     int *readonly, int *continuous, int *atomic)
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

After a successful call to `MPI_T_PVAR_GET_INFO` for a particular variable, subsequent calls to this routine that query information about the same variable must return the same information. An MPI implementation is not allowed to alter any of the returned values.

If any OUT parameter to `MPI_T_PVAR_GET_INFO` is a NULL pointer, the implementation will ignore the parameter and not return a value for the parameter.

The arguments `name` and `name_len` are used to return the name of the performance variable as described in Section 15.3.3. If completed successfully, the routine is required to return a name of at least length one.

The argument `verbosity` returns the verbosity level of the variable (see Section 15.3.1).

The class of the performance variable is returned in the parameter `var_class`. The class must be one of the constants defined in Section 15.3.7.

The combination of the name and the class of the performance variable must be unique with respect to all other names for performance variables used by the MPI implementation.

Advice to implementors. Groups of variables that belong closely together, but have different classes, can have the same name. This choice is useful, e.g., to refer to multiple variables that describe a single resource (like the level, the total size, as well as high- and low-water marks). (*End of advice to implementors.*)

The argument `datatype` returns the MPI datatype that is used to represent the performance variable.

If the variable is of type `MPI_INT`, MPI can optionally specify an enumeration for the values represented by this variable and return it in `enumtype`. In this case, MPI returns an enumeration identifier, which can then be used to gather more information as described in Section 15.3.5. Otherwise, `enumtype` is set to `MPI_T_ENUM_NULL`. If the datatype is not `MPI_INT` or the argument `enumtype` is the null pointer, no enumeration type is returned.

Returning a description is optional. If an MPI implementation does not return a description, the first character for `desc` must be set to the null character and `desc_len` must be set to one at the return from this function.

The parameter `bind` returns the type of the MPI object to which the variable must be bound or the value `MPI_T_BIND_NO_OBJECT` (see Section 15.3.2).

Upon return, the argument `readonly` is set to zero if the variable can be written or reset by the user. It is set to one if the variable can only be read.

Upon return, the argument `continuous` is set to zero if the variable can be started and stopped by the user, i.e., it is possible for the user to control if and when the value of a variable is updated. It is set to one if the variable is always active and cannot be controlled by the user.

Upon return, the argument `atomic` is set to zero if the variable cannot be read and reset atomically. Only variables for which the call sets `atomic` to one can be used in a call to `MPI_T_PVAR_READRESET`.

If a performance variable has an equivalent name and has the same class across connected MPI processes, the following OUT parameters must be identical: `verbosity`, `varclass`,

datatype, enumtype, bind, readonly, continuous, and atomic. The returned description must be equivalent.

MPI_T_PVAR_GET_INDEX(name, var_class, pvar_index)

IN	name	the name of the performance variable (string)
IN	var_class	the class of the performance variable (integer)
OUT	pvar_index	the index of the performance variable (integer)

C binding

int MPI_T_pvar_get_index(const char *name, int var_class, int *pvar_index)

MPI_T_PVAR_GET_INDEX is a function for retrieving the index of a performance variable given a known variable name and class. The `name` and `var_class` parameters are provided by the caller, and `pvar_index` is returned by the MPI implementation. The `name` parameter is a string terminated with a null character.

This routine returns `MPI_SUCCESS` on success and returns `MPI_T_ERR_INVALID_NAME` if `name` does not match the name of any performance variable of the specified `var_class` provided by the implementation at the time of the call.

Rationale. This routine is provided to enable fast retrieval of performance variables by a tool, assuming it knows the name of the variable for which it is looking. The number of variables exposed by the implementation can change over time, so it is not possible for the tool to simply iterate over the list of variables once at initialization. Although using MPI implementation specific variable names is not portable across MPI implementations, tool developers may choose to take this route for lower overhead at runtime because the tool will not have to iterate over the entire set of variables to find a specific one. (*End of rationale.*)

Performance Experiment Sessions

Within a single program, multiple components can use the MPI tool information interface. To avoid collisions with respect to accesses to performance variables, users of the MPI tool information interface must first create a performance experiment session. Subsequent calls that access performance variables can then be made within the context of this performance experiment session. Starting, stopping, reading, writing, or resetting a variable in one performance experiment session shall not influence whether a variable is started, stopped, read, written, or reset in another performance experiment session.

MPI_T_PVAR_SESSION_CREATE(pe_session)

OUT	pe_session	identifier of performance experiment session (handle)
-----	------------	---

C binding

int MPI_T_pvar_session_create(MPI_T_pvar_session *pe_session)

This call creates a new performance experiment session for accessing performance variables and returns a handle for this performance experiment session in the argument `pe_session` of type `MPI_T_pvar_session`.

1 MPI_T_PVAR_SESSION_FREE(pe_session)

2 INOUT pe_session identifier of performance experiment session (handle)

4 **C binding**

5 int MPI_T_pvar_session_free(MPI_T_pvar_session *pe_session)

7 This call frees an existing performance experiment session. Calls to the MPI tool
8 information interface can no longer be made within the context of a performance experiment
9 session after it is freed. On a successful return, MPI sets the performance experiment session
10 identifier to MPI_T_PVAR_SESSION_NULL.

11 *Handle Allocation and Deallocation*

12 Before using a performance variable, a user must first allocate a handle of type
13 MPI_T_pvar_handle for the variable by binding it to an MPI object (see also Section 15.3.2).

14 MPI_T_PVAR_HANDLE_ALLOC(pe_session, pvar_index, obj_handle, handle, count)

15 INOUT pe_session identifier of performance experiment session (handle)

16 IN pvar_index index of performance variable for which handle is to
17 be allocated (integer)

18 IN obj_handle reference to a handle of the MPI object to which this
19 variable is supposed to be bound (pointer)

20 OUT handle allocated handle (handle)

21 OUT count number of elements used to represent this variable
22 (integer)

23 **C binding**

24 int MPI_T_pvar_handle_alloc(MPI_T_pvar_session pe_session, int pvar_index,
25 void *obj_handle, MPI_T_pvar_handle *handle, int *count)

26 This routine binds the performance variable specified by the argument index to an MPI
27 object in the performance experiment session identified by the parameter pe_session. The
28 object is passed in the argument obj_handle as an address to a local variable that stores
29 the object's handle. The argument obj_handle is ignored if the MPI_T_PVAR_GET_INFO
30 call for this performance variable returned MPI_T_BIND_NO_OBJECT in the argument bind.
31 The handle allocated to reference the variable is returned in the argument handle. Upon
32 successful return, count contains the number of elements (of the datatype returned by a
33 previous MPI_T_PVAR_GET_INFO call) used to represent this variable.

34 *Advice to users.* The count can be different based on the MPI object to which the
35 performance variable was bound. For example, variables bound to communicators
36 could have a count that matches the size of the communicator.

37 It is not portable to pass references to predefined MPI object handles, such as
38 MPI_COMM_WORLD, to this routine, since their implementation depends on the MPI
39 library. Instead, such an object handle should be stored in a local variable and the
40 address of this local variable should be passed into MPI_T_PVAR_HANDLE_ALLOC.
41 (*End of advice to users.*)

The value of index should be in the range from 0 to `num_pvar - 1`, where `num_pvar` is the number of available performance variables as determined from a prior call to `MPI_T_PVAR_GET_NUM`. The type of the MPI object it references must be consistent with the type returned in the `bind` argument in a prior call to `MPI_T_PVAR_GET_INFO`.

For all routines in the rest of this section that take both `handle` and `pe_session` as IN or INOUT arguments, if the `handle` argument passed in is not associated with the `pe_session` argument, `MPI_T_ERR_INVALID_HANDLE` is returned.

`MPI_T_PVAR_HANDLE_FREE(pe_session, handle)`

INOUT	<code>pe_session</code>	identifier of performance experiment session (handle)
INOUT	<code>handle</code>	handle to be freed (handle)

C binding

```
int MPI_T_pvar_handle_free(MPI_T_pvar_session pe_session,
                          MPI_T_pvar_handle *handle)
```

When a handle is no longer needed, a user of the MPI tool information interface should call `MPI_T_PVAR_HANDLE_FREE` to free the handle in the performance experiment session identified by the parameter `pe_session` and the associated resources in the MPI implementation. On a successful return, MPI sets the handle to `MPI_T_PVAR_HANDLE_NULL`.

Starting and Stopping of Performance Variables

Performance variables that have the continuous flag set during the query procedure are continuously updated once a handle has been allocated. Such variables may be queried at any time, but they cannot be started or stopped by the user. All other variables are in a stopped state after their handle has been allocated; their values are not updated until they have been started by the user.

`MPI_T_PVAR_START(pe_session, handle)`

IN	<code>pe_session</code>	identifier of performance experiment session (handle)
INOUT	<code>handle</code>	handle of a performance variable (handle)

C binding

```
int MPI_T_pvar_start(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle)
```

This functions starts the performance variable with the handle identified by the parameter `handle` in the performance experiment session identified by the parameter `pe_session`.

If the constant `MPI_T_PVAR_ALL_HANDLES` is passed in `handle`, the MPI implementation attempts to start all variables within the performance experiment session identified by the parameter `pe_session` for which handles have been allocated. In this case, the routine returns `MPI_SUCCESS` if all variables are started successfully (even if there are no noncontinuous variables to be started), otherwise `MPI_T_ERR_PVAR_NO_STARTSTOP` is returned. Continuous variables and variables that are already started are ignored when `MPI_T_PVAR_ALL_HANDLES` is specified.

```

1 MPI_T_PVAR_STOP(pe_session, handle)
2     IN      pe_session      identifier of performance experiment session (handle)
3
4     INOUT   handle          handle of a performance variable (handle)
5

```

C binding

```

7 int MPI_T_pvar_stop(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle)
8

```

This function stops the performance variable with the handle identified by the parameter `handle` in the performance experiment session identified by the parameter `pe_session`.

If the constant `MPI_T_PVAR_ALL_HANDLES` is passed in `handle`, the MPI implementation attempts to stop all variables within the performance experiment session identified by the parameter `pe_session` for which handles have been allocated. In this case, the routine returns `MPI_SUCCESS` if all variables are stopped successfully (even if there are no noncontinuous variables to be stopped), otherwise `MPI_T_ERR_PVAR_NO_STARTSTOP` is returned. Continuous variables and variables that are already stopped are ignored when `MPI_T_PVAR_ALL_HANDLES` is specified.

Performance Variable Access Functions

```

21 MPI_T_PVAR_READ(pe_session, handle, buf)
22
23     IN      pe_session      identifier of performance experiment session (handle)
24
25     IN      handle          handle of a performance variable (handle)
26
27     OUT     buf             initial address of storage location for variable value
28                             (choice)
29

```

C binding

```

30 int MPI_T_pvar_read(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle,
31                   void *buf)
32

```

The `MPI_T_PVAR_READ` call queries the value of the performance variable with the handle `handle` in the performance experiment session identified by the parameter `pe_session` and stores the result in the buffer identified by the parameter `buf`. The user is responsible to ensure that the buffer is of the appropriate size to hold the entire value of the performance variable (based on the datatype and count returned by the corresponding previous calls to `MPI_T_PVAR_GET_INFO` and `MPI_T_PVAR_HANDLE_ALLOC`, respectively).

The constant `MPI_T_PVAR_ALL_HANDLES` cannot be used as an argument for the function `MPI_T_PVAR_READ`.

```

41
42 MPI_T_PVAR_WRITE(pe_session, handle, buf)
43
44     IN      pe_session      identifier of performance experiment session (handle)
45
46     INOUT   handle          handle of a performance variable (handle)
47
48     IN      buf             initial address of storage location for variable value
49                             (choice)
50

```

C binding

```
int MPI_T_pvar_write(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle,
                    const void *buf)
```

The `MPI_T_PVAR_WRITE` call attempts to write the value of the performance variable with the handle identified by the parameter `handle` in the performance experiment session identified by the parameter `pe_session`. The value to be written is passed in the buffer identified by the parameter `buf`. The user must ensure that the buffer is of the appropriate size to hold the entire value of the performance variable (based on the datatype and count returned by the corresponding previous calls to `MPI_T_PVAR_GET_INFO` and `MPI_T_PVAR_HANDLE_ALLOC`, respectively).

If it is not possible to change the variable, the function returns `MPI_T_ERR_PVAR_NO_WRITE`.

The constant `MPI_T_PVAR_ALL_HANDLES` cannot be used as an argument for the function `MPI_T_PVAR_WRITE`.

```
MPI_T_PVAR_RESET(pe_session, handle)
```

IN	<code>pe_session</code>	identifier of performance experiment session (handle)
INOUT	<code>handle</code>	handle of a performance variable (handle)

C binding

```
int MPI_T_pvar_reset(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle)
```

The `MPI_T_PVAR_RESET` call sets the performance variable with the handle identified by the parameter `handle` to its starting value specified in Section 15.3.7. If it is not possible to change the variable, the function returns `MPI_T_ERR_PVAR_NO_WRITE`.

If the constant `MPI_T_PVAR_ALL_HANDLES` is passed in `handle`, the MPI implementation attempts to reset all variables within the performance experiment session identified by the parameter `pe_session` for which handles have been allocated. In this case, the routine returns `MPI_SUCCESS` if all variables are reset successfully (even if there are no valid handles or all are read-only), otherwise `MPI_T_ERR_PVAR_NO_WRITE` is returned. Read-only variables are ignored when `MPI_T_PVAR_ALL_HANDLES` is specified.

```
MPI_T_PVAR_READRESET(pe_session, handle, buf)
```

IN	<code>pe_session</code>	identifier of performance experiment session (handle)
INOUT	<code>handle</code>	handle of a performance variable (handle)
OUT	<code>buf</code>	initial address of storage location for variable value (choice)

C binding

```
int MPI_T_pvar_readreset(MPI_T_pvar_session pe_session,
                        MPI_T_pvar_handle handle, void *buf)
```

This call atomically combines the functionality of `MPI_T_PVAR_READ` and `MPI_T_PVAR_RESET` with the same semantics as if these two calls were called separately.

1 If the variable cannot be read and reset atomically, this routine returns
2 MPI_T_ERR_PVAR_NO_ATOMIC.

3 The constant MPI_T_PVAR_ALL_HANDLES cannot be used as an argument for the func-
4 tion MPI_T_PVAR_READRESET.

5
6 *Advice to implementors.* Sampling-based tools rely on the ability to call the MPI tool
7 information interface, in particular routines to start, stop, read, write, and reset per-
8 formance variables, from any program context, including asynchronous contexts such
9 as signal handlers. MPI implementations should strive, if possible in their particular
10 environment, to enable these usage scenarios for all or a subset of the routines men-
11 tioned above. If implementing only a subset, the read, write, and reset routines are
12 typically the most critical for sampling based tools. An MPI implementation should
13 clearly document any restrictions on the program contexts in which the MPI tool infor-
14 mation interface can be used. Restrictions might include guaranteeing usage outside
15 of all signals or outside a specific set of signals. Any restrictions could be documented,
16 for example, through the description returned by MPI_T_PVAR_GET_INFO. (*End of*
17 *advice to implementors.*)

18
19 *Rationale.* All routines to read, to write or to reset performance variables require
20 the performance experiment session argument. This requirement keeps the interface
21 consistent and allows the use of MPI_T_PVAR_ALL_HANDLES where appropriate. Fur-
22 ther, this opens up additional performance optimizations for the implementation of
23 handles. (*End of rationale.*)

24
25 **Example 15.7.** Detecting Receives with long unexpected message queues.

26 The following example shows a sample tool to identify receive operations that occur during
27 times with long message queues. This examples assumes that the MPI implementation
28 exports a variable with the name “MPI_T_UMQ_LENGTH” to represent the current length of the
29 unexpected message queue. The tool is implemented as a PMPI tool using the MPI profiling
30 interface.

31 The tool consists of three parts: (1) the initialization (by intercepting the call to MPI_INIT),
32 (2) the test for long unexpected message queues (by intercepting calls to MPI_RECV), and
33 (3) the clean-up phase (by intercepting the call to MPI_FINALIZE). To capture all receives,
34 the example would have to be extended to have similar wrappers for all receive operations.

35
36 **Part 1—Initialization:** During initialization, the tool searches for the variable and, once
37 the right index is found, allocates a performance experiment session and a handle for the
38 variable with the found index, and starts the performance variable.

```
39 #include <stdio.h>
40 #include <stdlib.h>
41 #include <string.h>
42 #include <assert.h>
43 #include <mpi.h>
44
45 /* Global variables for the tool */
46 static MPI_T_pvar_session pe_session;
47 static MPI_T_pvar_handle handle;
48
49 int MPI_Init(int *argc, char ***argv )
```

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
{
    int err, num, i, index, namelen, verbosity;
    int var_class, bind, threadsup;
    int readonly, continuous, atomic, count;
    char name[18];

    MPI_Comm comm;
    MPI_Datatype datatype;
    MPI_T_enum enumtype;

    err = PMPI_Init(argc, argv);
    if (err != MPI_SUCCESS)
        return err;

    err = PMPI_T_init_thread(MPI_THREAD_SINGLE, &threadsup);
    if (err != MPI_SUCCESS)
        return err;

    err = PMPI_T_pvar_get_num(&num);
    if (err != MPI_SUCCESS)
        return err;

    index = -1;
    i = 0;
    while ((i < num) && (index < 0) && (err == MPI_SUCCESS)) {
        /* Pass a buffer that is at least one character longer than */
        /* the name of the variable being searched for to avoid */
        /* finding variables that have a name that has a prefix */
        /* equal to the name of the variable being searched. */
        namelen = 18;
        err = PMPI_T_pvar_get_info(i, name, &namelen, &verbosity,
                                   &var_class, &datatype, &enumtype,
                                   NULL, NULL, &bind, &readonly,
                                   &continuous, &atomic);
        if (strcmp(name, "MPI_T_UMQ_LENGTH") == 0) index = i;

        i++;
    }
    if (err != MPI_SUCCESS)
        return err;

    /* this could be handled in a more flexible way for a generic tool */
    assert(index >= 0);
    assert(var_class == MPI_T_PVAR_CLASS_LEVEL);
    assert(datatype == MPI_INT);
    assert(bind == MPI_T_BIND_MPI_COMM);

    /* Create a session */
    err = PMPI_T_pvar_session_create(&pe_session);
    if (err != MPI_SUCCESS)
        return err;

    /* Get a handle and bind to MPI_COMM_WORLD */
    comm = MPI_COMM_WORLD;
    err = PMPI_T_pvar_handle_alloc(pe_session, index, &comm, &handle,
                                   &count);
    if (err != MPI_SUCCESS)

```

```

1      return err;
2
3      /* this could be handled in a more flexible way for a generic tool */
4      assert(count == 1);
5
6      /* Start variable */
7      err = PMPI_T_pvar_start(pe_session, handle);
8      if (err != MPI_SUCCESS)
9          return err;
10
11     return MPI_SUCCESS;
12 }

```

Part 2—Testing the Queue Lengths During Receives: During every receive operation, the tool reads the unexpected queue length through the matching performance variable and compares it against a predefined threshold.

```

13 #define THRESHOLD 5
14
15 int MPI_Recv(void *buf, int count, MPI_Datatype datatype, int source,
16             int tag, MPI_Comm comm, MPI_Status *status)
17 {
18     int value, err;
19
20     if (comm == MPI_COMM_WORLD) {
21         err = PMPI_T_pvar_read(pe_session, handle, &value);
22         if ((err == MPI_SUCCESS) && (value > THRESHOLD)) {
23             /* tool identified receive called with long UMQ */
24             /* execute tool functionality, */
25             /* e.g., gather and print call stack */
26         }
27     }
28
29     return PMPI_Recv(buf, count, datatype, source, tag, comm, status);
30 }

```

Part 3—Termination: In the wrapper for MPI_FINALIZE, the MPI tool information interface is finalized.

```

31 int MPI_Finalize(void)
32 {
33     int err;
34
35     err = PMPI_T_pvar_handle_free(pe_session, &handle);
36     err = PMPI_T_pvar_session_free(&pe_session);
37     err = PMPI_T_finalize();
38
39     return PMPI_Finalize();
40 }

```

15.3.8 Events

During the execution of an MPI application, the MPI implementation can raise *events* of a specific type to inform the user of a state change in the implementation. **Event types**

describe specific state changes within the MPI implementation. In comparison to aggregate performance variables, events provide per-instance information on such state changes. The MPI implementation is said to **raise an event** when it invokes a callback function previously registered for the corresponding event type by the user. Each callback invocation for a specific event instance has a timestamp associated with it, which can be queried by the user, describing the time when the event was observed by the implementation. This decouples the observation of the state change from the communication of this information to the user. A timestamp in this context is a count of clock ticks elapsed since some time in the past and represented as a variable of type `MPI_Count`.

Event Sources

As a means to manage multiple state changes to be observed concurrently by different parts of the software and hardware system, the event interface of the MPI Tool Information Interface uses the concept of *sources*. A source in this context is a concept describing the logical entity raising the event. A source may or may not directly represent a concrete part of the software or hardware system. This concept is used primarily to describe partial ordering of events across different components where total ordering cannot necessarily be determined or is too costly to enforce.

The following function can be used to query the number of event sources, `num_sources`:

```
MPI_T_SOURCE_GET_NUM(num_sources)
```

OUT	<code>num_sources</code>	returns number of event sources (integer)
-----	--------------------------	---

C binding

```
int MPI_T_source_get_num(int *num_sources)
```

The number of available event sources can be queried with a call to `MPI_T_SOURCE_GET_NUM`. An MPI implementation is allowed to increase the number of sources during the execution of an MPI process. However, MPI implementations are not allowed to change the index of an event source or to delete an event source once it has been made visible to the user (e.g., if new event sources become available via dynamic loading of additional components in the MPI implementation).

```
MPI_T_SOURCE_GET_INFO(source_index, name, name_len, desc, desc_len, ordering,
                      ticks_per_second, max_ticks, info)
```

IN	<code>source_index</code>	index of the source to be queried between 0 and <code>num_sources - 1</code> (integer)
OUT	<code>name</code>	buffer to return the string containing the name of the source (string)
INOUT	<code>name_len</code>	length of the string and/or buffer for <code>name</code> (integer)
OUT	<code>desc</code>	buffer to return the string containing the description of the source (string)
INOUT	<code>desc_len</code>	length of the string and/or buffer for <code>desc</code> (integer)

1	OUT	ordering	flag indicating chronological ordering guarantees given by the source (integer)
2			
3	OUT	ticks_per_second	the number of ticks per second for the timer of this source (integer)
4			
5	OUT	max_ticks	the maximum count of ticks reported by this source before overflow occurs (integer)
6			
7	OUT	info	optional info object (handle)
8			
9			

10 C binding

```
11 int MPI_T_source_get_info(int source_index, char *name, int *name_len,
12     char *desc, int *desc_len, MPI_T_source_order *ordering,
13     MPI_Count *ticks_per_second, MPI_Count *max_ticks,
14     MPI_Info *info)
15
```

16 A call to `MPI_T_SOURCE_GET_INFO` returns additional information on the source identified by the `source_index` argument.

17 The arguments `name` and `name_len` are used to return the name of the source as described in Section 15.3.3.

18 The arguments `desc` and `desc_len` are used to return the description of the source as described in Section 15.3.3.

19 The `ordering` argument returns whether event callbacks of this source will be invoked in chronological order, i.e., the timestamps reported by `MPI_T_EVENT_GET_TIMESTAMP` of subsequent events of the same source are monotonically increasing. The value of `ordering` can be `MPI_T_SOURCE_ORDERED` or `MPI_T_SOURCE_UNORDERED`.

20 The `ticks_per_seconds` argument returns the number of ticks elapsed in one second for the timer used for the specific source.

21 The `max_ticks` argument returns the largest number of ticks reported by this source as a timestamp before the value overflows.

22 *Advice to users.* As the size of `MPI_Count` is defined in relation to the types `MPI_Aint` and `MPI_Offset`, the effective size of `MPI_Count` may lead to overflows of the timestamp values reported. Users can use the argument `max_ticks` to mitigate resulting problems. *(End of advice to users.)*

23 MPI can optionally return an info object containing the default hints set for this source. If the argument to `info` provided by the user is the `NULL` pointer, this argument is ignored, otherwise an MPI implementation is required to return all hints that are supported by the implementation for this source and have default values specified; any user-supplied hints that were not ignored by the implementation; and any additional hints that were set by the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no key/value pair. The user is responsible for freeing `info` via `MPI_INFO_FREE`.

Table 15.5: Hierarchy of safety requirement levels for event callback routines

Safety Requirement
MPI_T_CB_REQUIRE_NONE
MPI_T_CB_REQUIRE_MPI_RESTRICTED
MPI_T_CB_REQUIRE_THREAD_SAFE
MPI_T_CB_REQUIRE_ASYNC_SIGNAL_SAFE

MPI_T_SOURCE_GET_TIMESTAMP(source_index, timestamp)

IN	source_index	index of the source (integer)
OUT	timestamp	current timestamp from specified source (integer)

C binding

```
int MPI_T_source_get_timestamp(int source_index, MPI_Count *timestamp)
```

To enable proper query of a reference timestamp for a specific source, a user can obtain a current timestamp using `MPI_T_SOURCE_GET_TIMESTAMP`. The argument `source_index` identifies the index of the source to query. The call returns `MPI_SUCCESS` and a current timestamp in the argument `timestamp` if the source supports ad-hoc generation of timestamps. The call returns `MPI_T_ERR_INVALID_INDEX` if the index does not identify a valid source. The call returns `MPI_T_ERR_NOT_SUPPORTED` if the source does not support the ad-hoc generation of timestamps.

Callback Safety Requirements

The actions a user is allowed to perform inside a callback function may vary with its execution context. As the user has no control over the execution context of specific callback function invocations, MPI provides a way to communicate this information using callback safety levels.

Table 15.5 provides the hierarchy of callback safety requirements levels within user-defined callback functions. The MPI implementation provides the safety requirement as an argument to the callback when it is invoked.

The level of `MPI_T_CB_REQUIRE_NONE` is the lowest level and does not impose any restrictions on the callback function.

The level of `MPI_T_CB_REQUIRE_MPI_RESTRICTED` restricts the set of MPI functions that can be called from inside the callback to all functions with the prefix `MPI_T` as well as `MPI_WTICK` and `MPI_WTIME`.

Advice to users. While some MPI functions are safe to be called inside a callback function used in the MPI tool information interface—which may in some implementations be issued from asynchronous contexts such as signal handlers—this does not imply that those MPI functions are generally safe to be called in asynchronous contexts such as signal handlers. (*End of advice to users.*)

The level of `MPI_T_CB_REQUIRE_THREAD_SAFE` includes all the limitations of `MPI_T_CB_REQUIRE_MPI_RESTRICTED` and additionally requires the callback to be reentrant

Table 15.6: List of MPI functions that when called from within a callback function may not return MPI_T_ERR_NOT_ACCESSIBLE

```

MPI_T_EVENT_COPY
MPI_T_EVENT_GET_SOURCE
MPI_T_EVENT_GET_TIMESTAMP
MPI_T_EVENT_READ
MPI_T_PVAR_READ
MPI_T_PVAR_READRESET
MPI_T_PVAR_RESET
MPI_T_PVAR_START
MPI_T_PVAR_STOP
MPI_T_PVAR_WRITE
MPI_T_SOURCE_GET_TIMESTAMP

```

and thread-safe. This means the callback must allow its execution to be interrupted by or happen concurrently with any other callback including itself.

The level of MPI_T_CB_REQUIRE_ASYNC_SIGNAL_SAFE includes all the limitations of MPI_T_CB_REQUIRE_THREAD_SAFE and additionally requires the callback to meet the safety requirements needed to support invocations from asynchronous contexts, such as signal handlers.

Advice to users. It is always safe to assume the highest restrictions for a callback invocation (i.e., MPI_T_CB_REQUIRE_ASYNC_SIGNAL_SAFE). By evaluating the specific requirements at runtime, a tool may obtain more freedom of action within the callback. (*End of advice to users.*)

Advice to implementors. A high-quality implementation will strive to set callback safety requirements to the most permissive level for a given callback invocation. (*End of advice to implementors.*)

All functions with the prefix MPI_T, except those listed in Table 15.6, may return the return code MPI_T_ERR_NOT_ACCESSIBLE to indicate that the user may not access this function at this time. The functions (and their respective PMPI versions) listed in Table 15.6 are exceptions to this rule and shall not return MPI_T_ERR_NOT_ACCESSIBLE.

Rationale. A call may be implemented in a way that is not safe for all execution contexts of a callback function, e.g., inside a signal handler. An MPI implementation therefore needs a way to communicate its inability to perform a certain action due to the execution context of a callback invocation. (*End of rationale.*)

Advice to implementors. A high-quality implementation shall not return MPI_T_ERR_NOT_ACCESSIBLE except where absolutely necessary. (*End of advice to implementors.*)

Advice to users. Users intercepting calls into the MPI tool information interface using the PMPI interface must ensure that the safety requirements for the calling context are met. This means that users may have to implement the wrapper with the highest safety level used by the MPI implementation. (*End of advice to users.*)

Event Type Query Functions

An MPI implementation exports a set of N event types through the MPI tool information interface. If N is zero, then the MPI implementation does not export any event types; otherwise, the provided event types are indexed from 0 to $N - 1$. This index number is used in subsequent calls to identify a specific event type.

An MPI implementation is allowed to increase the number of event types during the execution of an MPI process. However, MPI implementations are not allowed to change the index of an event type or to delete an event type once it has been made visible to the user (e.g., if new event types become available via dynamic loading of additional components in the MPI implementation).

The following function can be used to query the number of event types, `num_events`:

`MPI_T_EVENT_GET_NUM(num_events)`

OUT `num_events` returns number of event types (integer)

C binding

`int MPI_T_event_get_num(int *num_events)`

The function `MPI_T_EVENT_GET_INFO` provides access to additional information about a specific event type.

`MPI_T_EVENT_GET_INFO(event_index, name, name_len, verbosity, array_of_datatypes, array_of_displacements, num_elements, enumtype, info, desc, desc_len, bind)`

IN	<code>event_index</code>	index of the event type to be queried between 0 and <code>num_events - 1</code> (integer)
OUT	<code>name</code>	buffer to return the string containing the name of the event type (string)
INOUT	<code>name_len</code>	length of the string and/or buffer for <code>name</code> (integer)
OUT	<code>verbosity</code>	verbosity level of this event type (integer)
OUT	<code>array_of_datatypes</code>	array of MPI basic datatypes used to encode the event data (array of handles)
OUT	<code>array_of_displacements</code>	array of byte displacements of the elements in the event buffer (array of non-negative integers)
INOUT	<code>num_elements</code>	length of <code>array_of_datatypes</code> and <code>array_of_displacements</code> arrays (non-negative integer)
OUT	<code>enumtype</code>	optional descriptor for enumeration information (handle)
OUT	<code>info</code>	optional info object (handle)
OUT	<code>desc</code>	buffer to return the string containing a description of the event type (string)
INOUT	<code>desc_len</code>	length of the string and/or buffer for <code>desc</code> (integer)

1 OUT bind type of MPI object to which an event of this type
2 must be bound (integer)

4 C binding

```
5 int MPI_T_event_get_info(int event_index, char *name, int *name_len,
6                         int *verbosity, MPI_Datatype array_of_datatypes[],
7                         MPI_Aint array_of_displacements[], int *num_elements,
8                         MPI_T_enum *enumtype, MPI_Info *info, char *desc, int *desc_len,
9                         int *bind)
```

11 After a successful call to `MPI_T_EVENT_GET_INFO` for a particular event type, subsequent calls to this routine that query information about the same event type must return the same information. If any `INOUT` or `OUT` argument to `MPI_T_EVENT_GET_INFO` is a `NULL` pointer, the implementation will ignore the argument and not return a value for the specific argument.

16 The arguments `name` and `name_len` are used to return the name of the event type as described in Section 15.3.3. If completed successfully, the routine is required to return a name of at least length one. The name of the event type must be unique with respect to all other names for event types used by the MPI implementation.

20 The argument `verbosity` returns the verbosity level of the event type (see Section 15.3.1).

21 The argument `array_of_datatypes` returns an array of MPI datatype handles that describe the elements returned for an instance of the event type with index `event_index`. The event data can either be queried element by element with `MPI_T_EVENT_READ` or copied into a contiguous event buffer with `MPI_T_EVENT_COPY`. For the latter case, the argument `array_of_displacements` returns an array of byte displacements in the event buffer in ascending order starting with zero.

27 The user is responsible for the memory allocation for the `array_of_datatypes` and `array_of_displacements` arrays. The number of elements in each array is supplied by the user in `num_elements`. If the number of elements used by the event type is larger than the value of `num_elements` provided by the user, the number of datatype handles and displacements returned in the corresponding arrays is truncated to the value of `num_elements` passed in by the user. If the user passes the `NULL` pointer for `array_of_datatypes` or `array_of_displacements`, the respective arguments are ignored. Unless the user passes the `NULL` pointer for `num_elements`, the function returns the number of elements required for this event type. If the user passes the `NULL` pointer for `num_elements`, the arguments `num_elements`, `array_of_datatypes`, and `array_of_displacements` are ignored.

37 MPI can optionally return an enumeration identifier in the `enumtype` argument, describing the individual elements in the `array_of_datatypes` argument. Otherwise, `enumtype` is set to `MPI_T_ENUM_NULL`. If the argument to `enumtype` provided by the user is the `NULL` pointer, no enumeration type is returned.

41 MPI can optionally return an info object containing the default hints set for a registration handle for this event type. If the argument to `info` provided by the user is the `NULL` pointer, this argument is ignored, otherwise an MPI implementation is required to return all hints that are supported by the implementation for a registration handle for this event type and have default values specified; any user-supplied hints that were not ignored by the implementation; and any additional hints that were set by the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no key/value pair. The user is responsible for freeing `info` via `MPI_INFO_FREE`.

The arguments `desc` and `desc_len` are used to return the description of the event type as described in Section 15.3.3. Returning a description is optional. If an MPI implementation does not return a description, the first character for `desc` must be set to the null character and `desc_len` must be set to one at the return from this function.

The parameter `bind` returns the type of the MPI object to which the event type must be bound or the value `MPI_T_BIND_NO_OBJECT` (see Section 15.3.2).

If an event type has an equivalent name across connected MPI processes, the following OUT parameters must be identical: `verbosity`, `array_of_datatypes`, `num_elements`, `enumtype`, and `bind`. The returned description must be equivalent. As the argument `array_of_displacements` is process dependent, it may differ across connected MPI processes.

This routine returns `MPI_SUCCESS` on success and returns `MPI_T_ERR_INVALID_INDEX` if `event_index` does not match a valid event type index provided by the implementation at the time of the call.

`MPI_T_EVENT_GET_INDEX(name, event_index)`

IN	<code>name</code>	name of the event type (string)
OUT	<code>event_index</code>	index of the event type (integer)

C binding

`int MPI_T_event_get_index(const char *name, int *event_index)`

`MPI_T_EVENT_GET_INDEX` returns the index of an event type identified by a known event type name. The name parameter is provided by the caller, and `event_index` is returned by the MPI implementation. The name parameter is a string terminated with a null character.

This routine returns `MPI_SUCCESS` on success and returns `MPI_T_ERR_INVALID_NAME` if `name` does not match the name of any event type provided by the implementation at the time of the call.

Rationale. This routine is provided to enable fast retrieval of an event index by a tool, assuming it knows the name of the event type for which it is looking. The number of event types exposed by the implementation can change over time, so it is not possible for the tool to simply iterate over the list of event types once at initialization. Although using MPI implementation specific event type names is not portable across MPI implementations, tool developers may choose to take this route for lower overhead at runtime because the tool will not have to iterate over the entire set of event types to find a specific one. (*End of rationale.*)

Handle Allocation and Deallocation

Before the MPI implementation calls a callback function on the occurrence of a specific event, the user needs to register a callback function to be called for that event type and obtain a handle of type `MPI_T_event_registration`.

```

1 MPI_T_EVENT_HANDLE_ALLOC(event_index, obj_handle, info, event_registration)
2   IN      event_index      index of event type for which the registration handle
3                               is to be allocated (integer)
4
5   IN      obj_handle       reference to a handle of the MPI object to which this
6                               event is supposed to be bound (pointer)
7
8   IN      info             info object (handle)
9
10  OUT     event_registration event registration (handle)

```

C binding

```

11 int MPI_T_event_handle_alloc(int event_index, void *obj_handle, MPI_Info info,
12                             MPI_T_event_registration *event_registration)
13

```

MPI_T_EVENT_HANDLE_ALLOC creates a **registration handle** for the event type identified by `event_index`. Furthermore, if required by the event type, the registration handle is bound to the object referred to by the argument `obj_handle`. The argument `obj_handle` is ignored if the MPI_T_EVENT_GET_INFO call for this event type returned MPI_T_BIND_NO_OBJECT in the argument `bind`. The user can pass hints for the handle allocation to the MPI implementation via the `info` argument. The allocated event-registration handle is returned in the argument `event_registration`.

```

24 MPI_T_EVENT_HANDLE_SET_INFO(event_registration, info)

```

```

25   INOUT  event_registration  event registration (handle)
26   IN     info                info object (handle)

```

C binding

```

28 int MPI_T_event_handle_set_info(MPI_T_event_registration event_registration,
29                               MPI_Info info)
30

```

MPI_T_EVENT_HANDLE_SET_INFO updates the hints of the event-registration handle associated with `event_registration` using the hints provided in `info`. A call to this procedure has no effect on previously set or defaulted hints that are not specified by `info`. It also has no effect on previously set or defaulted hints that are specified by `info`, but are ignored by the MPI implementation in this call to MPI_T_EVENT_HANDLE_SET_INFO.

Advice to users. Some info items that an implementation can use when it creates an event-registration handle cannot easily be changed once the registration handle is created. Thus, an implementation may ignore hints issued in this call that it would have accepted in a handle allocation call. An implementation may also be unable to update certain info hints in a call to MPI_T_EVENT_HANDLE_SET_INFO. MPI_T_EVENT_HANDLE_GET_INFO can be used to determine whether info changes were ignored by the implementation. (*End of advice to users.*)

MPI_T_EVENT_HANDLE_GET_INFO(event_registration, info_used)				1
IN	event_registration	event registration (handle)		2
OUT	info_used	info object (handle)		3
				4
				5

C binding

```
int MPI_T_event_handle_get_info(MPI_T_event_registration event_registration,
                               MPI_Info *info_used)
```

MPI_T_EVENT_HANDLE_GET_INFO returns a new info object containing the hints of the event-registration handle associated with `event_registration`. The current setting of all hints related to this registration handle is returned in `info_used`. An MPI implementation is required to return all hints that are supported by the implementation and have default values specified; any user-supplied hints that were not ignored by the implementation; and any additional hints that were set by the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no key/value pairs. The user is responsible for freeing `info_used` via `MPI_INFO_FREE`.

MPI_T_EVENT_REGISTER_CALLBACK(event_registration, cb_safety, info, user_data,				19
event_cb_function)				20

INOUT	event_registration	event registration (handle)		21
IN	cb_safety	maximum callback safety level (integer)		22
IN	info	info object (handle)		23
IN	user_data	pointer to a user-controlled buffer		24
IN	event_cb_function	pointer to user-defined callback function (function)		25
				26
				27
				28

C binding

```
int MPI_T_event_register_callback(MPI_T_event_registration event_registration,
                                 MPI_T_cb_safety cb_safety, MPI_Info info, void *user_data,
                                 MPI_T_event_cb_function event_cb_function)
```

MPI_T_EVENT_REGISTER_CALLBACK associates a user-defined function pointed to by `event_cb_function` with an allocated event-registration handle. The maximum callback safety level supported by the callback function is passed in the argument `cb_safety`. The safety levels are defined in Table 15.5. A user can register multiple callback functions for a given event-registration handle, potentially specifying one for each callback safety level. Registering a callback function for a specific callback safety level overwrites any previously-registered callback function pointer and info object associated with the event registration for the specific callback safety level. If `event_cb_function` is the NULL pointer, an existing association of a callback function for that callback safety level is removed.

When an event is triggered, the implementation will select from all registered callbacks the callback with the lowest safety level valid in the context in which the callback is invoked. In situations where the required callback safety level exceeds the highest level for which a callback function is registered for a given registration handle, the event instance is dropped.

At callback invocation time, the implementation passes the pointer to a user-defined memory region specified during callback registration with the argument `user_data`.

The user can pass hints for the registration of the specified callback function to the MPI implementation via the info argument.

Advice to users. As event instances can be raised as soon as the registration handle is associated with the first callback function, the callback function with the highest callback safety guarantees should be registered before any further registrations for lower callback safety guarantees, to avoid dropped events due to insufficient callback safety guarantees. (*End of advice to users.*)

The callback function passed to `MPI_T_EVENT_REGISTER_CALLBACK` in the argument `event_cb_function` needs to have the following type:

```
typedef void MPI_T_event_cb_function(MPI_T_event_instance event_instance,
                                     MPI_T_event_registration event_registration,
                                     MPI_T_cb_safety cb_safety, void *user_data);
```

The argument `event_instance` corresponds to a handle for the opaque event-instance object of type `MPI_T_event_instance`. This handle is only valid inside the corresponding invocation of the function to which it is passed. The argument `event_registration` corresponds to the event-registration handle returned by `MPI_T_EVENT_HANDLE_ALLOC` for the user function to the same event type and bound object combination. The handle can be used to identify the specific event registration information, such as event type and bound object, or even to deallocate the handle from within the callback invocation. The argument `cb_safety` describes the safety requirements the callback function must fulfill in the current invocation. The argument `user_data` is the pointer to user-allocated memory that was passed to the MPI implementation during callback registration.

```
MPI_T_EVENT_CALLBACK_SET_INFO(event_registration, cb_safety, info)
```

INOUT	<code>event_registration</code>	event registration (handle)
IN	<code>cb_safety</code>	callback safety level (integer)
IN	<code>info</code>	info object (handle)

C binding

```
int MPI_T_event_callback_set_info(MPI_T_event_registration event_registration,
                                  MPI_T_cb_safety cb_safety, MPI_Info info)
```

`MPI_T_EVENT_CALLBACK_SET_INFO` updates the hints of the callback function registered for the callback safety level specified by `cb_safety` of the event-registration handle associated with `event_registration` using the hints provided in `info`. A call to this procedure has no effect on previously set or defaulted hints that are not specified by `info`. It also has no effect on previously set or defaulted hints that are specified by `info`, but are ignored by the MPI implementation in this call to `MPI_T_EVENT_CALLBACK_SET_INFO`.

```
MPI_T_EVENT_CALLBACK_GET_INFO(event_registration, cb_safety, info_used)
```

IN	<code>event_registration</code>	event registration (handle)
IN	<code>cb_safety</code>	callback safety level (integer)
OUT	<code>info_used</code>	info object (handle)

C binding

```
int MPI_T_event_callback_get_info(MPI_T_event_registration event_registration,
                                MPI_T_cb_safety cb_safety, MPI_Info *info_used)
```

`MPI_T_EVENT_CALLBACK_GET_INFO` returns a new info object containing the hints of the callback function registered for the callback safety level specified by `cb_safety` of the event-registration handle associated with `event_registration`. The current set of all hints related to this callback safety level of the event-registration handle is returned in `info_used`. An MPI implementation is required to return all hints that are supported by the implementation and have default values specified, any user-supplied hints that were not ignored by the implementation, and any additional hints that were set by the implementation. If no such hints exist, a handle to a newly created info object is returned that contains no key/value pairs. The user is responsible for freeing `info_used` via `MPI_INFO_FREE`.

To stop the MPI implementation from raising events for a specific registration, a user needs to free the corresponding event-registration handle.

```
MPI_T_EVENT_HANDLE_FREE(event_registration, user_data, free_cb_function)
```

INOUT	<code>event_registration</code>	event registration (handle)
IN	<code>user_data</code>	pointer to a user-controlled buffer
IN	<code>free_cb_function</code>	pointer to user-defined callback function (function)

C binding

```
int MPI_T_event_handle_free(MPI_T_event_registration event_registration,
                            void *user_data, MPI_T_event_free_cb_function free_cb_function)
```

`MPI_T_EVENT_HANDLE_FREE` returns `MPI_SUCCESS` when deallocation of the handle was initiated successfully and returns `MPI_T_ERR_INVALID_HANDLE` if `event_registration` does not match a valid allocated event-registration handle at the time of the call. The callback function `free_cb_function` is called by the MPI implementation, when it is able to guarantee that no further event instances for the corresponding event-registration handle will be raised. If the pointer to `free_cb_function` is the `NULL` pointer, no user function is invoked after successful deallocation of the event registration handle. The pointer to user-controlled memory provided in the `user_data` argument will be passed to the function provided in the `free_cb_function` on invocation.

Advice to users. A free-callback function associated with a registration handle should always be prepared to postpone any pending actions, should the provided callback safety requirements exceed those required by the pending actions. (*End of advice to users.*)

The callback function passed to `MPI_T_EVENT_HANDLE_FREE` in the argument `free_cb_function` needs to have the following type:

```
typedef void MPI_T_event_free_cb_function(
    MPI_T_event_registration event_registration,
    MPI_T_cb_safety cb_safety, void *user_data);
```

1 Handling Dropped Events

2 Events may occur at times when the MPI implementation cannot invoke the user function
 3 corresponding to a matching event handle. An implementation is allowed to buffer such
 4 events and delay the callback invocation. If an event occurs at times when the corresponding
 5 callback function cannot be called and the corresponding data cannot be buffered, or no
 6 callback function meeting the required callback safety level is registered, the event data may
 7 be dropped. To discover such data loss, the user can set a handler function for a specific
 8 event-registration handle.
 9

10
 11 MPI_T_EVENT_SET_DROPPED_HANDLER(event_registration, dropped_cb_function)

12 INOUT event_registration valid event registration (handle)
 13 IN dropped_cb_function pointer to user-defined callback function (function)
 14
 15

16 C binding

17 int MPI_T_event_set_dropped_handler(
 18 MPI_T_event_registration event_registration,
 19 MPI_T_event_dropped_cb_function dropped_cb_function)
 20

21 MPI_T_EVENT_SET_DROPPED_HANDLER registers the function
 22 `dropped_cb_function` to be called by the MPI implementation when event information is
 23 dropped for the registration handle specified in `event_registration`. Subsequent calls to
 24 MPI_T_EVENT_SET_DROPPED_HANDLER with the same registration handle will replace
 25 previously-registered callback functions for that registration handle. If the pointer to
 26 `dropped_cb_function` is the NULL pointer, no data loss is recorded or reported until a new
 27 valid callback function is registered.

28 *Advice to users.* The invocation of the dropped handler callback function may not
 29 necessarily occur close to the time the event was actually lost. (*End of advice to*
 30 *users.*)
 31

32 The callback function passed to MPI_T_EVENT_SET_DROPPED_HANDLER in the
 33 argument `dropped_cb_function` needs to have the following type:

34 typedef void MPI_T_event_dropped_cb_function(MPI_Count count,
 35 MPI_T_event_registration event_registration, int source_index,
 36 MPI_T_cb_safety cb_safety, void *user_data);
 37

38 The argument `event_registration` corresponds to the event registration handle to which
 39 the dropped data corresponds. The argument `count` provides a best effort estimation of
 40 the number of invocations to a registered event callback corresponding to
 41 `event_registration` that were not executed since the registration of the dropped-callback
 42 handler or the last invocation of a registered dropped-callback handler. If the number
 43 of dropped events observed by the implementation exceeds the limit of
 44 `count`, an implementation shall set `count` to the maximum possible value for the type of
 45 `count`. The `source_index` provides the index of the source that dropped the corresponding
 46 event information. The argument `cb_safety` describes the safety requirements the callback
 47 function must fulfill in the current invocation. The possible values for `cb_safety` are described
 48 in Table 15.5. The argument `user_data` is the pointer to user-allocated memory that was

passed to the MPI implementation during callback registration. If no event callback is registered for safety requirement levels that an implementation uses to invoke the dropped handler callback function for a specific event, the corresponding dropped handler callback function will not be invoked.

Advice to users. A callback function for dropped events associated with a registration handle should always be prepared to postpone any pending actions, should the provided callback safety requirements exceed those required by the pending actions. (*End of advice to users.*)

Advice to implementors. A high-quality implementation should strive to find a good balance between timely notification, completeness of information, and the freedom of action for a tool when invoking the callback function for dropped events associated with a registration handle. (*End of advice to implementors.*)

If dropped event notifications have been observed for a specific source since the last event notification of that source, the corresponding dropped handler callback function must be called before other events are raised for that source. This means in a sequence of five events E1 to E5 from the same source, where E3 and E4 were dropped, any handler function set through `MPI_T_EVENT_SET_DROPPED_HANDLER` for event-registration handles associated with E3 or E4 must be called before E5 is raised.

Reading Event Data

In event callbacks, the parameter `event_instance` provides access to the per-instance event data, i.e., the data encoded by the specific event type for this instance. The user can obtain event data as well as event meta data, such as a time stamp and the source, by providing this handle to the respective query functions. The event-instance handle is invalid beyond the scope of the current invocation of the callback function to which it is provided.

The callback function argument `event_registration` identifies the registration handle that was used to register the callback function.

The callback function argument `cb_safety` indicates the requirements for the specific callback invocation. The value is one of the safety requirements levels described in Table 15.5. The argument `user_data` passes the pointer provided by the user during callback registration back to the function call.

Advice to users. Depending on the registered event and usage of MPI by the application, a callback function may be invoked with high frequency. Users should therefore strive to minimize the amount of work done inside callback functions. Furthermore, the time spent in a callback function may influence the capability of an implementation to buffer events; long execution times may lead to an increased number of dropped events. (*End of advice to users.*)

MPI provides the following function calls to access data of a specific event instance and its corresponding meta data (such as its time and source).

1 MPI_T_EVENT_READ(event_instance, element_index, buffer)

2	IN	event_instance	event-instance handle provided to the callback
3			function (handle)
4			
5	IN	element_index	index into the array of datatypes of the item to be
6			queried (integer)
7	OUT	buffer	pointer to a memory location to store the item data
8			(choice)
9			

10 C binding

11 int MPI_T_event_read(MPI_T_event_instance event_instance, int element_index,
12 void *buffer)

13
14 MPI_T_EVENT_READ allows users to copy one element of the event data to a user-
15 specified buffer at a time.

16 The event_instance argument identifies the event instance to query. It is erroneous
17 to provide any other event-instance handle to the call than the one passed by the MPI
18 implementation to the callback function in which the data is read. The buffer argument
19 must point to a memory location the MPI implementation can copy the element of the event
20 data to identified by element_index.

21
22 MPI_T_EVENT_COPY(event_instance, buffer)

23	IN	event_instance	event instance provided to the callback function
24			(handle)
25			
26	OUT	buffer	user-allocated buffer for event data (choice)
27			

28 C binding

29 int MPI_T_event_copy(MPI_T_event_instance event_instance, void *buffer)

30
31 MPI_T_EVENT_COPY copies the event data as a whole into the user-provided buffer.
32 The user must assure that the buffer is of at least the size of the extent of the event
33 type, which can be computed from the type and displacement information returned by
34 the corresponding call to MPI_T_EVENT_GET_INFO. The data may include padding bytes
35 between individual elements of the event data in the buffer. A user can reconstruct the
36 location and size of the data contained in the buffer through the information returned by
37 MPI_T_EVENT_GET_INFO.

38
39 *Advice to implementors.* An implementation should strive to use an appropriately
40 compact representation when copying event instance data to a user buffer via
41 MPI_T_EVENT_COPY to reduce the amount of memory required for the user buffer.
42 (*End of advice to implementors.*)

43 Reading Event Meta Data

44
45 Additional to the specific event data encoded by each event type, supplemental information
46 available across all event types can be queried.

MPI_T_EVENT_GET_TIMESTAMP(event_instance, event_timestamp)

IN	event_instance	event instance provided to the callback function (handle)
OUT	event_timestamp	timestamp the event was observed (integer)

C binding

```
int MPI_T_event_get_timestamp(MPI_T_event_instance event_instance,
                             MPI_Count *event_timestamp)
```

MPI_T_EVENT_GET_TIMESTAMP returns the timestamp of when the event was initially observed by the implementation. The `event_instance` argument identifies the event instance to query. It is erroneous to provide any other handle to the call than the one passed by the MPI implementation to the callback function in which the timestamp is read.

Advice to users. An MPI implementation may postpone the call to the user's callback function. In this case, the call to MPI_T_EVENT_GET_TIMESTAMP may yield a timestamp in the past that is closer to the time the event was initially observed, as opposed to a timestamp captured during callback function invocation. (*End of advice to users.*)

Advice to implementors. A high-quality implementation will return a timestamp as close as possible to the earliest time the event was observed by the MPI implementation. (*End of advice to implementors.*)

An event may be raised from different components acting as event sources in the MPI implementation. A source in this context is an abstract concept that helps to define partial ordering of raised events, as each source provides its own ordering guarantees. A source describes the entity that raises the event, rather than the origin of the data.

To identify the source of an event instance, the user can query the index of the source within the corresponding event callback function invocation.

Advice to implementors. An excessive number of event sources may negatively impact performance of a tool due to per-source overhead in event handling. (*End of advice to implementors.*)

MPI_T_EVENT_GET_SOURCE(event_instance, source_index)

IN	event_instance	event instance provided to the callback function (handle)
OUT	source_index	index identifying the source (integer)

C binding

```
int MPI_T_event_get_source(MPI_T_event_instance event_instance,
                           int *source_index)
```

The `event_instance` argument identifies the event instance to query. It is erroneous to provide any other event-instance handle to the call than the one passed by the MPI implementation to the callback function in which the source is queried.

1 The `source_index` argument returns the index of the source of the event instance. It
 2 can be used to query more information on the source using `MPI_T_SOURCE_GET_INFO`.

3
 4 *Rationale.* Event callback function invocations are associated with a source to enable
 5 chronological processing of events on the tool side, when required, while retaining low
 6 overhead on the side of the MPI implementation. (*End of rationale.*)

8 15.3.9 Variable Categorization

9
 10 MPI implementations can optionally group performance and control variables into categories
 11 to express logical relationships between various variables. For example, an MPI implemen-
 12 tation could group all control and performance variables that refer to message transfers in
 13 the MPI implementation and thereby distinguish them from variables that refer to local
 14 resources such as memory allocations or other interactions with the operating system.

15 Categories can also contain other categories to form a hierarchical grouping. Categories
 16 can never include themselves, either directly or transitively within other included categories.
 17 Expanding on the example above, this allows MPI to refine the grouping of variables referring
 18 to message transfers into variables to control and to monitor message queues, message
 19 matching activities and communication protocols. Each of these groups of variables would
 20 be represented by a separate category and these categories would then be listed in a single
 21 category representing variables for message transfers.

22 The category information may be queried in a fashion similar to the mechanism for
 23 querying variable information. The MPI implementation exports a set of N categories via
 24 the MPI tool information interface. If $N = 0$, then the MPI implementation does not export
 25 any categories, otherwise the provided categories are indexed from 0 to $N - 1$. This index
 26 number is used in subsequent calls to functions of the MPI tool information interface to
 27 identify the individual categories.

28 An MPI implementation is permitted to increase the number of categories during the
 29 execution of an MPI program when new categories become available through dynamic load-
 30 ing. However, MPI implementations are not allowed to change the index of a category or
 31 delete it once it has been added to the set.

32 Similarly, MPI implementations are allowed to add variables to categories, but they
 33 are not allowed to remove variables from categories or change the order in which they are
 34 returned.

35 *Category Query Functions*

36
 37 The following function can be used to query the number of categories, `num_cat`.

38
 39
 40 `MPI_T_CATEGORY_GET_NUM(num_cat)`

41 OUT `num_cat` current number of categories (integer)

42 **C binding**

43
 44 `int MPI_T_category_get_num(int *num_cat)`

45
 46 Individual category information can then be queried by calling the following function:

MPI_T_CATEGORY_GET_INFO(cat_index, name, name_len, desc, desc_len, num_cvars,			1
num_pvars, num_categories)			2
IN	cat_index	index of the category to be queried (integer)	3
OUT	name	buffer to return the string containing the name of the	4
		category (string)	5
INOUT	name_len	length of the string and/or buffer for name (integer)	6
OUT	desc	buffer to return the string containing the description	7
		of the category (string)	8
INOUT	desc_len	length of the string and/or buffer for desc (integer)	9
OUT	num_cvars	number of control variables in the category (integer)	10
OUT	num_pvars	number of performance variables in the category	11
		(integer)	12
OUT	num_categories	number of categories contained in the category	13
		(integer)	14

C binding

```
int MPI_T_category_get_info(int cat_index, char *name, int *name_len,
                           char *desc, int *desc_len, int *num_cvars, int *num_pvars,
                           int *num_categories)
```

The arguments `name` and `name_len` are used to return the name of the category as described in Section 15.3.3.

The routine is required to return a name of at least length one. This name must be unique with respect to all other names for categories used by the MPI implementation.

If any OUT parameter to `MPI_T_CATEGORY_GET_INFO` is the NULL pointer, the implementation will ignore the parameter and not return a value for the parameter.

The arguments `desc` and `desc_len` are used to return the description of the category as described in Section 15.3.3.

Returning a description is optional. If an MPI implementation decides not to return a description, the first character for `desc` must be set to the null character and `desc_len` must be set to one at the return of this call.

The function returns the number of control variables, performance variables and other categories contained in the queried category in the arguments `num_cvars`, `num_pvars`, and `num_categories`, respectively.

If the name of a category is equivalent across connected MPI processes, then the returned description must be equivalent.

```
MPI_T_CATEGORY_GET_NUM_EVENTS(cat_index, num_events)
```

IN	cat_index	index of the category to be queried (integer)	41
OUT	num_events	number of event types in the category (integer)	42

C binding

```
int MPI_T_category_get_num_events(int cat_index, int *num_events)
```

1 MPI_T_CATEGORY_GET_NUM_EVENTS returns the number of event types contained
2 in the queried category.

3
4
5 MPI_T_CATEGORY_GET_INDEX(name, cat_index)

6 IN name the name of the category (string)
7 OUT cat_index the index of the category (integer)
8

9 C binding

10 int MPI_T_category_get_index(const char *name, int *cat_index)

11
12 MPI_T_CATEGORY_GET_INDEX is a function for retrieving the index of a category
13 given a known category name. The name parameter is provided by the caller, and cat_index
14 is returned by the MPI implementation. The name parameter is a string terminated with a
15 null character.

16 This routine returns MPI_SUCCESS on success and returns MPI_T_ERR_INVALID_NAME
17 if name does not match the name of any category provided by the implementation at the
18 time of the call.

19
20 *Rationale.* This routine is provided to enable fast retrieval of a category index
21 by a tool, assuming it knows the name of the category for which it is looking. The
22 number of categories exposed by the implementation can change over time, so it is not
23 possible for the tool to simply iterate over the list of categories once at initialization.
24 Although using MPI implementation specific category names is not portable across
25 MPI implementations, tool developers may choose to take this route for lower overhead
26 at runtime because the tool will not have to iterate over the entire set of categories
27 to find a specific one. (*End of rationale.*)
28

29 Category Member Query Functions

30
31
32 MPI_T_CATEGORY_GET_CVARS(cat_index, len, indices)

33 IN cat_index index of the category to be queried, in the range
34 from 0 to num_cat - 1 (integer)
35
36 IN len the length of the indices array (integer)
37 OUT indices an integer array of size len, indicating control
38 variable indices (array of integers)
39

40 C binding

41 int MPI_T_category_get_cvars(int cat_index, int len, int indices[])

42
43 MPI_T_CATEGORY_GET_CVARS can be used to query which control variables are
44 contained in a particular category. A category contains zero or more control variables.
45
46
47
48

MPI_T_CATEGORY_GET_PVARS(cat_index, len, indices)			1
IN	cat_index	index of the category to be queried, in the range from 0 to num_cat - 1 (integer)	2
			3
			4
IN	len	the length of the indices array (integer)	5
OUT	indices	an integer array of size len, indicating performance variable indices (array of integers)	6
			7
			8

C binding

```
int MPI_T_category_get_pvars(int cat_index, int len, int indices[])
```

MPI_T_CATEGORY_GET_PVARS can be used to query which performance variables are contained in a particular category. A category contains zero or more performance variables.

MPI_T_CATEGORY_GET_EVENTS(cat_index, len, indices)			16
IN	cat_index	index of the category to be queried, in the range from 0 to num_cat - 1 (integer)	17
			18
			19
IN	len	the length of the indices array (integer)	20
OUT	indices	an integer array of size len, indicating event type indices (array of integers)	21
			22
			23

C binding

```
int MPI_T_category_get_events(int cat_index, int len, int indices[])
```

MPI_T_CATEGORY_GET_EVENTS can be used to query which event types are contained in a particular category. A category contains zero or more event types.

MPI_T_CATEGORY_GET_CATEGORIES(cat_index, len, indices)			30
IN	cat_index	index of the category to be queried, in the range from 0 to num_cat - 1 (integer)	31
			32
			33
IN	len	the length of the indices array (integer)	34
OUT	indices	an integer array of size len, indicating category indices (array of integers)	35
			36
			37

C binding

```
int MPI_T_category_get_categories(int cat_index, int len, int indices[])
```

MPI_T_CATEGORY_GET_CATEGORIES can be used to query which other categories are contained in a particular category. A category contains zero or more other categories.

As mentioned above, MPI implementations can grow the number of categories as well as the number of variables or other categories within a category. In order to allow users of the MPI tool information interface to check quickly whether new categories have been added or new variables or categories have been added to a category, MPI maintains an update number that is monotonically increasing during the execution and is returned by the following function:

```

1 MPI_T_CATEGORY_CHANGED(update_number)
2     OUT     update_number           update number (integer)
3
4

```

C binding

```

5 int MPI_T_category_changed(int *update_number)
6
7

```

8 If two calls to this routine return the same update number, it is guaranteed that the
9 category information has not changed between the two calls. If the update number retrieved
10 from the second call is higher, then some categories have been added or expanded. If the
11 number of changes to categories exceeds the limit of `update_number`, an implementation
12 shall set `update_number` to the maximum possible value for the type of `update_number`.

13 The index values returned in indices by `MPI_T_CATEGORY_GET_CVARS`,
14 `MPI_T_CATEGORY_GET_PVARS`, `MPI_T_CATEGORY_GET_EVENTS`, and
15 `MPI_T_CATEGORY_GET_CATEGORIES` can be used as input to
16 `MPI_T_CVAR_GET_INFO`, `MPI_T_PVAR_GET_INFO`, `MPI_T_EVENT_GET_INFO`, and
17 `MPI_T_CATEGORY_GET_INFO`, respectively.

18 The user is responsible for allocating the arrays passed into the functions
19 `MPI_T_CATEGORY_GET_CVARS`, `MPI_T_CATEGORY_GET_PVARS`,
20 `MPI_T_CATEGORY_GET_EVENTS`, and `MPI_T_CATEGORY_GET_CATEGORIES`. Starting
21 from array index 0, each function writes up to `len` elements into the array. If the category
22 contains more than `len` elements, the function returns an arbitrary subset of size `len`. Oth-
23 erwise, the entire set of elements is returned in the beginning entries of the array, and any
24 remaining array entries are not modified.

15.3.10 Return Codes for the MPI Tool Information Interface

25 All procedures defined as part of the MPI tool information interface return an integer *return*
26 *code* (see Table 15.7) to indicate whether the function was completed successfully or was
27 aborted. For the former case, the value `MPI_SUCCESS` is returned. In the latter case, the
28 return code indicates the reason for not completing the routine. Regardless of whether the
29 return code is `MPI_SUCCESS` or indicates that the procedure abnormally terminated, the
30 MPI process continues normal execution and does not invoke any MPI error handler. The
31 MPI implementation is not required to check all user-provided parameters; if a user passes
32 invalid parameter values to any routine, the behavior of the implementation is undefined.

33 All return codes with the prefix `MPI_T_ERR_` must be unique values and cannot overlap
34 with any error codes or error classes returned by the MPI implementation. They must also
35 satisfy

$$0 = \text{MPI_SUCCESS} < \text{MPI_T_ERR_XXX} \leq \text{MPI_ERR_LASTCODE}.$$

15.3.11 Profiling Interface

36 All requirements for the profiling interface, as described in Section 15.2, also apply to
37 the MPI tool information interface. All rules, guidelines, and recommendations from Sec-
38 tion 15.2 apply equally to procedures defined as part of the MPI tool information interface.

Table 15.7: Return codes used in procedures of the MPI tool information interface.

Return Code	Description
Return Codes for All Procedures in the MPI Tool Information Interface	
MPI_SUCCESS	Call completed successfully
MPI_T_ERR_INVALID	Invalid or bad parameter value(s)
MPI_T_ERR_MEMORY	Out of memory
MPI_T_ERR_NOT_INITIALIZED	Interface not initialized
MPI_T_ERR_CANNOT_INIT	Interface not in the state to be initialized
MPI_T_ERR_NOT_ACCESSIBLE	Requested functionality not accessible
Return Codes for Datatype Procedures: MPI_T_ENUM_XXX	
MPI_T_ERR_INVALID_INDEX	The enumeration index is invalid
Return Codes for Variable, Category, and Event Query Procedures: MPI_T_XXX_GET_XXX	
MPI_T_ERR_INVALID_INDEX	The variable or category index is invalid
MPI_T_ERR_INVALID_NAME	The variable or category name is invalid
Return Codes for Handle Procedures: MPI_T_XXX_{ALLOC FREE}	
MPI_T_ERR_INVALID_INDEX	The variable index is invalid
MPI_T_ERR_INVALID_HANDLE	The handle is invalid
MPI_T_ERR_OUT_OF_HANDLES	No more handles available
Return Codes for Performance Experiment Session Procedures: MPI_T_PVAR_SESSION_XXX	
MPI_T_ERR_OUT_OF_SESSIONS	No more sessions available
MPI_T_ERR_INVALID_SESSION	Session argument is not a valid session
Return Codes for Control Variable Access Procedures: MPI_T_CVAR_{READ WRITE}	
MPI_T_ERR_CVAR_SET_NOT_NOW	Variable cannot be set at this moment
MPI_T_ERR_CVAR_SET_NEVER	Variable cannot be set until end of execution
MPI_T_ERR_INVALID_HANDLE	The handle is invalid
Return Codes for Performance Variable Access and Control Procedures:	
MPI_T_PVAR_{START STOP READ WRITE RESET READREST}	
MPI_T_ERR_INVALID_HANDLE	The handle is invalid
MPI_T_ERR_INVALID_SESSION	Performance experiment session argument is invalid
MPI_T_ERR_PVAR_NO_STARTSTOP	Variable cannot be started or stopped (for MPI_T_PVAR_START and MPI_T_PVAR_STOP)
MPI_T_ERR_PVAR_NO_WRITE	Variable cannot be written or reset (for MPI_T_PVAR_WRITE and MPI_T_PVAR_RESET)
MPI_T_ERR_PVAR_NO_ATOMIC	Variable cannot be read and written atomically (for MPI_T_PVAR_READRESET)
Return Codes for Source Procedures: MPI_T_SOURCE_XXX	
MPI_T_ERR_INVALID_INDEX	The source index is invalid
MPI_T_ERR_NOT_SUPPORTED	Requested functionality not supported
Return Codes for Category Procedures: MPI_T_CATEGORY_XXX	
MPI_T_ERR_INVALID_INDEX	The category index is invalid

DRAFT

Chapter 16

Deprecated Interfaces

16.1 Deprecated since MPI-2.0

- The following function is deprecated and is superseded by `MPI_COMM_CREATE_KEYVAL` in MPI-2.0. The language independent definition of the deprecated function is the same as that of the new function, except for the function name and a different behavior in the C/Fortran language interoperability, see Section 19.3.7. The language bindings are modified.

`MPI_KEYVAL_CREATE(copy_fn, delete_fn, keyval, extra_state)`

IN	<code>copy_fn</code>	Copy callback function for keyval
IN	<code>delete_fn</code>	Delete callback function for keyval
OUT	<code>keyval</code>	key value for future access (integer)
IN	<code>extra_state</code>	Extra state for callback functions

C binding

```
int MPI_Keyval_create(MPI_Copy_function *copy_fn,  
MPI_Delete_function *delete_fn, int *keyval, void *extra_state)
```

For this routine, an interface within the `mpi_f08` module was never defined.

Fortran binding

```
MPI_KEYVAL_CREATE(COPY_FN, DELETE_FN, KEYVAL, EXTRA_STATE, IERROR)  
EXTERNAL COPY_FN, DELETE_FN  
INTEGER KEYVAL, EXTRA_STATE, IERROR
```

The `copy_fn` function is invoked when a communicator is duplicated by `MPI_COMM_DUP`. `copy_fn` should be of type `MPI_Copy_function`, which is defined as follows:

```
typedef int MPI_Copy_function(MPI_Comm oldcomm, int keyval,  
void *extra_state, void *attribute_val_in, void *attribute_val_out,  
int *flag);
```

A Fortran declaration for such a function is as follows:

For this routine, an interface within the `mpi_f08` module was never defined.

```

1  SUBROUTINE COPY_FUNCTION(OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
2  ATTRIBUTE_VAL_OUT, FLAG, IERR)
3      INTEGER OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
4      ATTRIBUTE_VAL_OUT, IERR
5      LOGICAL FLAG
6

```

copy_fn may be specified as MPI_NULL_COPY_FN or MPI_DUP_FN from either C or Fortran; MPI_NULL_COPY_FN is a function that does nothing other than return flag = 0 and MPI_SUCCESS. MPI_DUP_FN is a simple-minded copy function that sets flag = 1, returns the value of attribute_val_in in attribute_val_out, and returns MPI_SUCCESS. Note that MPI_NULL_COPY_FN and MPI_DUP_FN are also deprecated.

Analogous to copy_fn is a callback deletion function, defined as follows. The delete_fn function is invoked when a communicator is deleted by MPI_COMM_FREE or when a call is made explicitly to MPI_ATTR_DELETE. delete_fn should be of type MPI_Delete_function, which is defined as follows:

```

17  typedef int MPI_Delete_function(MPI_Comm comm, int keyval,
18  void *attribute_val, void *extra_state);
19

```

A Fortran declaration for such a function is as follows:

For this routine, an interface within the mpi_f08 module was never defined.

```

23  SUBROUTINE DELETE_FUNCTION(COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR)
24  INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR
25

```

delete_fn may be specified as MPI_NULL_DELETE_FN from either C or Fortran; MPI_NULL_DELETE_FN is a function that does nothing other than return MPI_SUCCESS. Note that MPI_NULL_DELETE_FN is also deprecated.

- The following function is deprecated and is superseded by MPI_COMM_FREE_KEYVAL in MPI-2.0. The language independent definition of the deprecated function is the same as the new function, except for the function name. The language bindings are modified.

```

36  MPI_KEYVAL_FREE(keyval)
37

```

```

38  INOUT  keyval                                Frees the integer key value (integer)
39

```

C binding

```

41  int MPI_Keyval_free(int *keyval)
42

```

For this routine, an interface within the mpi_f08 module was never defined.

Fortran binding

```

46  MPI_KEYVAL_FREE(KEYVAL, IERROR)
47  INTEGER KEYVAL, IERROR
48

```

- The following function is deprecated and is superseded by `MPI_COMM_SET_ATTR` in MPI-2.0. The language independent definition of the deprecated function is the same as the new function, except for the function name. The language bindings are modified.

`MPI_ATTR_PUT(comm, keyval, attribute_val)`

INOUT	<code>comm</code>	communicator to which attribute will be attached (handle)
IN	<code>keyval</code>	key value, as returned by <code>MPI_KEYVAL_CREATE</code> (integer)
IN	<code>attribute_val</code>	attribute value

C binding

```
int MPI_Attr_put(MPI_Comm comm, int keyval, void *attribute_val)
```

For this routine, an interface within the `mpi_f08` module was never defined.

Fortran binding

```
MPI_ATTR_PUT(COMM, KEYVAL, ATTRIBUTE_VAL, IERROR)
      INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR
```

- The following function is deprecated and is superseded by `MPI_COMM_GET_ATTR` in MPI-2.0. The language independent definition of the deprecated function is the same as the new function, except for the function name. The language bindings are modified.

`MPI_ATTR_GET(comm, keyval, attribute_val, flag)`

IN	<code>comm</code>	communicator to which attribute is attached (handle)
IN	<code>keyval</code>	key value (integer)
OUT	<code>attribute_val</code>	attribute value, unless <code>flag = false</code>
OUT	<code>flag</code>	true if an attribute value was extracted; false if no attribute is associated with the key

C binding

```
int MPI_Attr_get(MPI_Comm comm, int keyval, void *attribute_val, int *flag)
```

For this routine, an interface within the `mpi_f08` module was never defined.

Fortran binding

```
MPI_ATTR_GET(COMM, KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)
```

1 INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR
 2
 3 LOGICAL FLAG

- 4
 5 • The following function is deprecated and is superseded by
 6 MPI_COMM_DELETE_ATTR in MPI-2.0. The language independent definition of the
 7 deprecated function is the same as the new function, except for the function name.
 8 The language bindings are modified.
 9

10
 11 MPI_ATTR_DELETE(comm, keyval)

12
 13 INOUT comm communicator to which attribute is attached
 14 (handle)
 15 IN keyval The key value of the deleted attribute (integer)
 16

17 **C binding**

18 int MPI_Attr_delete(MPI_Comm comm, int keyval)

19 For this routine, an interface within the `mpi_f08` module was never defined.
 20

21 **Fortran binding**

22 MPI_ATTR_DELETE(COMM, KEYVAL, IERROR)
 23 INTEGER COMM, KEYVAL, IERROR
 24

25 16.2 Deprecated since MPI-2.2

- 26
 27
 28 • The entire set of C++ language bindings was deprecated as of MPI-2.2 and removed
 29 in MPI-3.0. See Chapter 17, [Removed Interfaces](#) for more information.
 30
 31 • The following function typedefs have been deprecated and are superseded by new
 32 names. Other than the typedef names, the function signatures are exactly the same;
 33 the names were updated to match conventions of other function typedef names.
 34

35 Deprecated Name	36 New Name
37 MPI_Comm_errhandler_fn	38 MPI_Comm_errhandler_function
39 MPI_File_errhandler_fn	40 MPI_File_errhandler_function
41 MPI_Win_errhandler_fn	42 MPI_Win_errhandler_function

43 16.3 Deprecated since MPI-4.0

- 44 • Cancelling a send request by calling MPI_CANCEL has been deprecated and may be
 45 removed in a future version of the MPI specification.
 46
 47
 48

- The following function is deprecated and is superseded by the new MPI_INFO_GET_STRING call in MPI-4.0.

MPI_INFO_GET(info, key, valuelen, value, flag)

IN	info	info object (handle)
IN	key	key (string)
IN	valuelen	length of value associated with key (integer)
OUT	value	value (string)
OUT	flag	true if key defined, false if not (logical)

C binding

```
int MPI_Info_get(MPI_Info info, const char *key, int valuelen,
char *value, int *flag)
```

Fortran 2008 binding

```
MPI_Info_get(info, key, valuelen, value, flag, ierror)
  TYPE(MPI_Info), INTENT(IN) :: info
  CHARACTER(LEN=*), INTENT(IN) :: key
  INTEGER, INTENT(IN) :: valuelen
  CHARACTER(LEN=valuelen), INTENT(OUT) :: value
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_INFO_GET(INFO, KEY, VALUELEN, VALUE, FLAG, IERROR)
  INTEGER INFO, VALUELEN, IERROR
  CHARACTER*(*) KEY, VALUE
  LOGICAL FLAG
```

This function retrieves the value associated with key in a previous call to MPI_INFO_SET. If such a key exists, it sets **flag** to **true** and returns the value in **value**, otherwise it sets **flag** to **false** and leaves **value** unchanged. **valuelen** is the number of characters available in **value**. If it is less than the actual size of the value, the value is truncated. In C, **valuelen** should be one less than the amount of allocated space to allow for the null terminator.

If **key** is larger than MPI_MAX_INFO_KEY, the call is erroneous.

The function MPI_INFO_GET is allowed to be called at any time, following the description for MPI functionality that is always available in Section 11.4.1.

- The following function is deprecated and is superseded by the new MPI_INFO_GET_STRING call in MPI-4.0.

MPI_INFO_GET_VALUELEN(info, key, valuelen, flag)

IN	info	info object (handle)
IN	key	key (string)
OUT	valuelen	length of value associated with key (integer)
OUT	flag	true if key defined, false if not (logical)

C binding

```
int MPI_Info_get_valuelen(MPI_Info info, const char *key, int *valuelen,
int *flag)
```

Fortran 2008 binding

```
MPI_Info_get_valuelen(info, key, valuelen, flag, ierror)
  TYPE(MPI_Info), INTENT(IN) :: info
  CHARACTER(LEN=*), INTENT(IN) :: key
  INTEGER, INTENT(OUT) :: valuelen
  LOGICAL, INTENT(OUT) :: flag
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_INFO_GET_VALUELEN(INFO, KEY, VALUELEN, FLAG, IERROR)
  INTEGER INFO, VALUELEN, IERROR
  CHARACTER*(*) KEY
  LOGICAL FLAG
```

Retrieves the length of the value associated with key. If key is defined, valuelen is set to the length of its associated value and flag is set to true. If key is not defined, valuelen is not touched and flag is set to false. The length returned in C does not include the end-of-string character.

If key is larger than MPI_MAX_INFO_KEY, the call is erroneous.

The function MPI_INFO_GET_VALUELEN is allowed to be called at any time, following the description for MPI functionality that is always available in Section 11.4.1.

- The following return code has been deprecated and is superseded by a new name in MPI-4.0.

<u>Deprecated Name</u>	<u>Replacement Name</u>
MPI_T_ERR_INVALID_ITEM	MPI_T_ERR_INVALID_INDEX

- The following Fortran subroutines are deprecated because the Fortran language `storage_size()` and `c_sizeof()` intrinsic functions provide similar functionality. Note that while `MPI_SIZEOF` and `c_sizeof()` return the size in bytes, `storage_size()` provides the size in bits.

`MPI_SIZEOF(x, size)`

IN	x	a Fortran variable of numeric intrinsic type (choice)
OUT	size	size of machine representation of that type (integer)

Fortran 2008 binding

```
MPI_Sizeof(x, size, ierror)
  TYPE(*), DIMENSION(..) :: x
  INTEGER, INTENT(OUT) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_SIZEOF(X, SIZE, IERROR)
  <type> X
  INTEGER SIZE, IERROR
```

This function returns the size in bytes of the machine representation of the given variable. It is a generic Fortran routine and has a Fortran binding only.

Advice to users. This function is similar to the C `sizeof` operator but behaves slightly differently. If given an array argument, it returns the size of the base element, not the size of the whole array. (*End of advice to users.*)

Rationale. This function is not available in other languages because it would not be useful. (*End of rationale.*)

16.4 Deprecated since MPI-4.1

- The use of the `mpif.h` include file has been deprecated. Information supporting the transition to `USE mpi` or `USE mpi_f08` is provided in Section [19.1.4](#).
- The predefined attribute key `MPI_HOST` for `MPI_COMM_WORLD` when using the World Model is deprecated.

MPI_HOST: Host process rank, if such exists, `MPI_PROC_NULL`, otherwise.

Host Rank

The value returned for `MPI_HOST` gets the rank of the *HOST* process in the group associated with communicator `MPI_COMM_WORLD`, if there is such. `MPI_PROC_NULL` is returned if there is no host. MPI does not specify what it means for a process to be a *HOST*, nor does it requires that a *HOST* exists.

The attribute `MPI_HOST` has the same value on all processes of `MPI_COMM_WORLD`.

Environmental inquiry keys

C type: `const int` (or unnamed enum)

Fortran type: `INTEGER`

`MPI_HOST`

- All `MPI_XXX_X` procedures have been deprecated and may be removed in a future version of the MPI specification. In the case of their C binding and their Fortran binding through the `mpi_f08` module, they are superseded by the large count and large byte displacement bindings of their counterpart in the form of `MPI_XXX`.

`MPI_TYPE_SIZE_X(datatype, size)`

IN	<code>datatype</code>	<code>datatype</code> to get information on (handle)
OUT	<code>size</code>	<code>datatype</code> size (integer)

C binding

`int MPI_Type_size_x(MPI_Datatype datatype, MPI_Count *size)`

Fortran 2008 binding

`MPI_Type_size_x(datatype, size, ierror)`
`TYPE(MPI_Datatype), INTENT(IN) :: datatype`
`INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size`
`INTEGER, OPTIONAL, INTENT(OUT) :: ierror`

Fortran binding

`MPI_TYPE_SIZE_X(DATATYPE, SIZE, IERROR)`
`INTEGER DATATYPE, IERROR`
`INTEGER(KIND=MPI_COUNT_KIND) SIZE`

The description of `MPI_TYPE_SIZE` is applicable to this deprecated `MPI_TYPE_SIZE_X` accordingly, see Section [5.1.5](#).

MPI_TYPE_GET_EXTENT_X(datatype, lb, extent)			1
			2
IN	datatype	datatype to get information on (handle)	3
OUT	lb	lower bound of datatype (integer)	4
OUT	extent	extent of datatype (integer)	5

C binding

```
int MPI_Type_get_extent_x(MPI_Datatype datatype, MPI_Count *lb,
MPI_Count *extent)
```

Fortran 2008 binding

```
MPI_Type_get_extent_x(datatype, lb, extent, ierror)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: lb, extent
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_GET_EXTENT_X(DATATYPE, LB, EXTENT, IERROR)
  INTEGER DATATYPE, IERROR
  INTEGER(KIND=MPI_COUNT_KIND) LB, EXTENT
```

The description of MPI_TYPE_GET_EXTENT is applicable to this deprecated MPI_TYPE_GET_EXTENT_X accordingly, see Section 5.1.7.

MPI_TYPE_GET_TRUE_EXTENT_X(datatype, true_lb, true_extent)			28
			29
IN	datatype	datatype to get information on (handle)	30
OUT	true_lb	true lower bound of datatype (integer)	31
OUT	true_extent	true extent of datatype (integer)	32

C binding

```
int MPI_Type_get_true_extent_x(MPI_Datatype datatype, MPI_Count *true_lb,
MPI_Count *true_extent)
```

Fortran 2008 binding

```
MPI_Type_get_true_extent_x(datatype, true_lb, true_extent, ierror)
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: true_lb, true_extent
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_TYPE_GET_TRUE_EXTENT_X(DATATYPE, TRUE_LB, TRUE_EXTENT, IERROR)
  INTEGER DATATYPE, IERROR
```

1 INTEGER(KIND=MPI_COUNT_KIND) TRUE_LB, TRUE_EXTENT

2

3 The description of MPI_TYPE_GET_TRUE_EXTENT is applicable to this deprecated
4 MPI_TYPE_GET_TRUE_EXTENT_X accordingly, see Section 5.1.8.

5

6

7 MPI_GET_ELEMENTS_X(status, datatype, count)

8

9 IN status return status of receive operation (status)

10

11 IN datatype datatype used by receive operation (handle)

12

13 OUT count number of received basic elements (integer)

14

15 **C binding**

16

17 int MPI_Get_elements_x(const MPI_Status *status, MPI_Datatype datatype,
18 MPI_Count *count)

19

20 **Fortran 2008 binding**

21

22 MPI_Get_elements_x(status, datatype, count, ierror)

23

24 TYPE(MPI_Status), INTENT(IN) :: status

25

26 TYPE(MPI_Datatype), INTENT(IN) :: datatype

27

28 INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: count

29

30 INTEGER, OPTIONAL, INTENT(OUT) :: ierror

31

32 **Fortran binding**

33

34 MPI_GET_ELEMENTS_X(STATUS, DATATYPE, COUNT, IERROR)

35

36 INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, IERROR

37

38 INTEGER(KIND=MPI_COUNT_KIND) COUNT

39

40 The description of MPI_GET_ELEMENTS is applicable to this deprecated
41 MPI_GET_ELEMENTS_X accordingly, see Section 5.1.11.

42

43

44 MPI_STATUS_SET_ELEMENTS_X(status, datatype, count)

45

46 INOUT status status with which to associate count (status)

47

48 IN datatype datatype associated with count (handle)

49

50 IN count number of elements to associate with status
51 (integer)

52

53

54 **C binding**

55

56 int MPI_Status_set_elements_x(MPI_Status *status, MPI_Datatype datatype,
57 MPI_Count count)

58

59 **Fortran 2008 binding**

60

61 MPI_Status_set_elements_x(status, datatype, count, ierror)

```
TYPE(MPI_Status), INTENT(INOUT) :: status
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

Fortran binding

```
MPI_STATUS_SET_ELEMENTS_X(STATUS, DATATYPE, COUNT, IERROR)
  INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, IERROR
  INTEGER(KIND=MPI_COUNT_KIND) COUNT
```

The description of MPI_STATUS_SET_ELEMENTS is applicable to this deprecated MPI_STATUS_SET_ELEMENTS_X accordingly, see Section [13.3](#).

DRAFT

DRAFT

Chapter 17

Removed Interfaces

17.1 Removed MPI-1 Bindings

17.1.1 Overview

The following MPI-1 bindings were deprecated as of MPI-2 and were removed in MPI-3. They may be provided by an implementation for backwards compatibility, but are not required. Removal of these bindings affects all language-specific definitions thereof. Only the language-neutral bindings are listed when possible.

17.1.2 Removed MPI-1 Functions

Table 17.1 shows the removed MPI-1 functions and their replacements.

Table 17.1: Removed MPI-1 functions and their replacements

Removed	MPI-2 Replacement
MPI_ADDRESS	MPI_GET_ADDRESS
MPI_ERRHANDLER_CREATE	MPI_COMM_CREATE_ERRHANDLER
MPI_ERRHANDLER_GET	MPI_COMM_GET_ERRHANDLER
MPI_ERRHANDLER_SET	MPI_COMM_SET_ERRHANDLER
MPI_TYPE_EXTENT	MPI_TYPE_GET_EXTENT
MPI_TYPE_HINDEXED	MPI_TYPE_CREATE_HINDEXED
MPI_TYPE_HVECTOR	MPI_TYPE_CREATE_HVECTOR
MPI_TYPE_LB	MPI_TYPE_GET_EXTENT
MPI_TYPE_STRUCT	MPI_TYPE_CREATE_STRUCT
MPI_TYPE_UB	MPI_TYPE_GET_EXTENT

17.1.3 Removed MPI-1 Datatypes

Table 17.2 shows the removed MPI-1 datatypes and their replacements.

17.1.4 Removed MPI-1 Constants

Table 17.3 shows the removed MPI-1 constants. There are no replacements.

17.1.5 Removed MPI-1 Callback Prototypes

Table 17.4 shows the removed MPI-1 callback prototypes and their replacements.

Table 17.2: Removed MPI-1 datatypes. The indicated routine may be used for changing the lower and upper bound respectively.

Removed	MPI-2 Replacement
MPI_LB	MPI_TYPE_CREATE_RESIZED
MPI_UB	MPI_TYPE_CREATE_RESIZED

Table 17.3: Removed MPI-1 constants

Removed MPI-1 Constants
C type: <code>const int</code> (or unnamed enum)
Fortran type: <code>INTEGER</code>
MPI_COMBINER_HINDEXED_INTEGER
MPI_COMBINER_HVECTOR_INTEGER
MPI_COMBINER_STRUCT_INTEGER

Table 17.4: Removed MPI-1 callback prototypes and their replacements

Removed	MPI-2 Replacement
<code>MPI_Handler_function</code>	<code>MPI_Comm_errhandler_function</code>

17.2 C++ Bindings

The C++ bindings were deprecated as of MPI-2.2. The C++ bindings were removed in MPI-3.0. The namespace is still reserved, however, and bindings may only be provided by an implementation as described in the MPI-2.2 standard.

Chapter 18

Semantic Changes and Warnings

This chapter lists semantic changes that have been introduced into the MPI standard as well as warnings that could potentially impact program behavior. In addition to those listed here, Chapter 17 also lists changes and backward incompatibilities caused by removing interfaces. Unlike Chapter 17, the changes in this chapter did not go through a deprecation process.

18.1 Semantic Changes

This section describes semantics that have changed in a way that would potentially cause an MPI program to behave differently when using this version of the MPI standard without changing the program's code.

18.1.1 Semantic Changes Starting in MPI-4.0

- `MPI_COMM_DUP` and `MPI_COMM_IDUP` no longer propagate info hints from the input communicator to the output communicator. This behavior can be achieved using `MPI_COMM_DUP_WITH_INFO` and `MPI_COMM_IDUP_WITH_INFO`.
- The default communicator where errors are raised when not involving a communicator, window, or file was changed from `MPI_COMM_WORLD` to `MPI_COMM_SELF`.

18.2 Additional Warnings

This section describes additional changes that could potentially cause a program that relies on the semantics described in a previous version of the MPI standard to behave differently than with this version of MPI. The changes in this section are limited in scope and unlikely to impact most programs.

18.2.1 Warnings Starting in MPI-4.1

Implementations are no longer allowed to implement `MPI_WTICK`, `PMPI_WTICK`, `MPI_WTIME`, and `PMPI_WTIME` as well as handle conversion functions as macros (Sections 9.6 and 19.3.4). This should not impact applications but may require changes in some implementations.

18.2.2 Warnings Starting in MPI-4.0

The limit for length of MPI identifiers was removed. Prior to MPI-4.0, MPI identifiers were limited to 30 characters (31 with the profiling interface). This limitation was initially introduced to avoid exceeding the limit on some compilation systems.

Rationale. For Fortran, this limit was already relaxed for the Fortran specific function names, see Section 19.1.5, and the Fortran language specification 2003 requires support for a minimum of 63 characters for internal and external identifiers. Starting with the ISO/IEC 9899:1999 C programming language standard, support for a minimum of 63 characters is required for internal identifiers, but only 31 characters are required to be significant for external identifiers. At the time of the release of MPI-4.0, most or nearly all compilers allow external identifiers longer than 31 characters. Therefore, the restriction is removed. (*End of rationale.*)

Advice to users. This affects users only if they store MPI identifiers into fixed sized strings. (*End of advice to users.*)

DRAFT

Chapter 19

Language Bindings

19.1 Support for Fortran

19.1.1 Overview

The Fortran MPI language bindings have been designed to be compatible with the Fortran 90 standard with additional features from Fortran 2018 [48]. In previous versions of this document, references were made to Fortran 2003 and Fortran 2008 [46] with TS 29113 [47]; where appropriate, the specific features of Fortran 2018 that MPI requires will be noted explicitly.

Rationale. Fortran 90 contains numerous features designed to make it a more “modern” language than Fortran 77. It seems natural that MPI should be able to take advantage of these new features with a set of bindings tailored to Fortran 90. In Fortran 2008 with TS 29113 and later Fortran 2018, the major new language features used are the ASYNCHRONOUS attribute to protect nonblocking MPI operations, and assumed-type and assumed-rank dummy arguments for choice buffer arguments. Further requirements for compiler support are listed in Section 19.1.7. (*End of rationale.*)

MPI defines three methods of Fortran support:

1. **USE mpi_f08:** This method is described in Section 19.1.2. It requires compile-time argument checking with unique MPI handle types and provides techniques to fully solve the optimization problems with nonblocking calls. This is the only Fortran support method that is consistent with the Fortran standard (Fortran 2008 with TS 29113 and later Fortran 2018). This method is highly recommended for all MPI applications.
2. **USE mpi:** This method is described in Section 19.1.3 and requires compile-time argument checking. Handles are defined as INTEGER. This Fortran support method is inconsistent with the Fortran standard, and its use is therefore not recommended. It exists only for backwards compatibility.
3. **INCLUDE 'mpif.h':** This method is described in Section 19.1.4. The use of the include file mpif.h has been strongly discouraged starting with MPI-3.0 and deprecated with MPI-4.1, because this method neither guarantees compile-time argument checking nor provides sufficient techniques to solve the optimization problems with nonblocking calls, and is therefore inconsistent with the Fortran standard. It exists only for backwards compatibility with legacy MPI applications.

MPI implementations providing a Fortran interface must provide one or both of the following:

- 1 • The USE `mpi_f08` Fortran support method.
- 2
- 3 • The USE `mpi` and INCLUDE `'mpif.h'` Fortran support methods.

4 Section 19.1.6 describes restrictions if the compiler does not support all the needed features.

5 Application subroutines and functions may use either one of the modules or the (depre-
6 cated) `mpif.h` include file. An implementation may require the use of one of the modules
7 to prevent type mismatch errors.

8
9 *Advice to users.* Users are advised to utilize one of the MPI modules even if `mpif.h`
10 enforces type checking on a particular system. Using a module provides several poten-
11 tial advantages over using an include file; the `mpi_f08` module offers the most robust
12 and complete Fortran support. (*End of advice to users.*)

13
14 In a single application, it must be possible to link together routines that USE `mpi_f08`,
15 USE `mpi`, and INCLUDE `'mpif.h'`.

16 The LOGICAL constant `MPI_SUBARRAYS_SUPPORTED` is set to `.TRUE.` if all buffer choice
17 arguments are defined in explicit interfaces with assumed-type and assumed-rank [48]; oth-
18 erwise it is set to `.FALSE.` The LOGICAL constant `MPI_ASYNC_PROTECTS_NONBLOCKING`
19 is set to `.TRUE.` if the `ASYNCHRONOUS` attribute was added to the choice buffer arguments of
20 all nonblocking interfaces **and** the underlying Fortran compiler supports the `ASYNCHRONOUS`
21 attribute for MPI communication (as part of TS 29113, which has been superceded by For-
22 tran 2018), otherwise it is set to `.FALSE.` These constants exist for each Fortran support
23 method, but not in the C header file. The values may be different for each Fortran support
24 method. All other constants and the integer values of handles must be the same for each
25 Fortran support method.

26 Sections 19.1.2 through 19.1.4 define the Fortran support methods. The Fortran in-
27 terfaces of each MPI routine are shorthands. Section 19.1.5 defines the corresponding full
28 interface specification together with the specific procedure names and implications for the
29 profiling interface. Section 19.1.6 describes the implementation of the MPI routines for dif-
30 ferent versions of the Fortran standard. Section 19.1.7 summarizes major requirements for
31 MPI implementations with Fortran support. Section 19.1.8 and Section 19.1.9 describe ad-
32 ditional functionality that is part of the Fortran support. `MPI_F_SYNC_REG` is needed
33 for one of the methods to prevent register optimization problems. A set of functions
34 provides additional support for Fortran intrinsic numeric types, including parameterized
35 types: `MPI_TYPE_MATCH_SIZE`, `MPI_TYPE_CREATE_F90_INTEGER`,
36 `MPI_TYPE_CREATE_F90_REAL` and `MPI_TYPE_CREATE_F90_COMPLEX`. In the context
37 of MPI, parameterized types are Fortran intrinsic types that are specified using `KIND` type
38 parameters. Sections 19.1.10 through 19.1.19 give an overview and details on known prob-
39 lems when using Fortran together with MPI; Section 19.1.20 compares the Fortran problems
40 with those in C.

41 19.1.2 Fortran Support Through the `mpi_f08` Module

42
43 An MPI implementation providing a Fortran interface must provide a module named `mpi_f08`
44 that can be used in a Fortran program. Section 19.1.6 describes restrictions if the compiler
45 does not support all the needed features. Within all MPI function specifications, the first
46 of the set of two Fortran routine interface specifications is provided by this module. This
47 module must:
48

- Define all named MPI constants. 1
- Declare MPI functions that return a value. 2
- Provide explicit interfaces according to the Fortran routine interface specifications. This module therefore guarantees compile-time argument checking for all arguments that are not `TYPE(*)`, with the following exception: 3

Only one Fortran interface is defined for functions that are deprecated as of MPI-3.0. This interface must be provided as an explicit interface according to the rules defined for the `mpi` module, see Section 19.1.3. 4

Advice to users. It is strongly recommended that developers substitute calls to deprecated routines when upgrading from the (deprecated) `mpif.h` or the `mpi` module to the `mpi_f08` module. (*End of advice to users.*) 5

- Define the derived type `MPI_Status`, and define all MPI handles with uniquely named handle types (instead of `INTEGER` handles, as in the `mpi` module). This is reflected in the first Fortran binding in each MPI function definition throughout this document (except for the deprecated routines). 6
- Overload the operators `.EQ.` and `.NE.` to allow the comparison of these MPI handles with `.EQ.`, `.NE.`, `==` and `/=`. 7
- Use the `ASYNCHRONOUS` attribute to protect the buffers of nonblocking operations, and set the `LOGICAL` constant `MPI_ASYNC_PROTECTS_NONBLOCKING` to `.TRUE.` if the underlying Fortran compiler supports the `ASYNCHRONOUS` attribute for MPI communication (as part of TS 29113). See Section 19.1.6 for older compiler versions. 8
- Set the `LOGICAL` constant `MPI_SUBARRAYS_SUPPORTED` to `.TRUE.` and declare choice buffers using the Fortran 2018 features `assumed-type` and `assumed-rank`, i.e., `TYPE(*)`, `DIMENSION(..)` in all nonblocking, split collective and persistent communication routines, if the underlying Fortran compiler supports it. With this, noncontiguous sub-arrays can be used as buffers in nonblocking routines. 9

Rationale. In all blocking routines, i.e., if the choice-buffer is not declared as `ASYNCHRONOUS`, the Fortran 2018 feature is not needed for the support of noncontiguous buffers because the compiler can pass the buffer by in-and-out-copy through a contiguous scratch array. (*End of rationale.*) 10

- Set the `MPI_SUBARRAYS_SUPPORTED` constant to `.FALSE.` and declare choice buffers with a compiler-dependent mechanism that overrides type checking if the underlying Fortran compiler does not support the Fortran 2018 `assumed-type` and `assumed-rank` notation. In this case, the use of noncontiguous sub-arrays as buffers in nonblocking calls may be invalid. See Section 19.1.6 for details. 11
- Declare each argument with an `INTENT` of `IN`, `OUT`, or `INOUT` as defined in this standard. 12

Rationale. For these definitions in the `mpi_f08` bindings, in most cases, `INTENT(IN)` is used if the C interface uses call-by-value. For all buffer arguments and for `OUT` and `INOUT` dummy arguments that allow one of the nonordinary Fortran constants (see `MPI_BOTTOM`, etc. in Section 2.5.4) as input, an `INTENT` is not specified. (*End of rationale.*) 13

Advice to users. If a dummy argument is declared with `INTENT(OUT)`, then the Fortran standard stipulates that the actual argument becomes undefined upon invocation of the MPI routine, i.e., it may be overwritten by some other values, e.g. zeros; according to [46], 12.5.2.4 Ordinary dummy variables, Paragraph 17: “If a dummy argument has `INTENT(OUT)`, the actual argument becomes undefined at the time the association is established, except [...]”. For example, if the dummy argument is an assumed-size array and the actual argument is a strided array, the call may be implemented with copy-in and copy-out of the argument. In the case of `INTENT(OUT)` the copy-in may be suppressed by the optimization and the routine starts execution using an array of undefined values. If the routine stores fewer elements into the dummy argument than is provided in the actual argument, then the remaining locations are overwritten with these undefined values. See also both advices to implementors in Section 19.1.3. (*End of advice to users.*)

- Declare all `ierror` output arguments as `OPTIONAL`, except for user-defined callback functions (e.g., of type `MPI_Comm_copy_attr_function` or `COMM_COPY_ATTR_FUNCTION`) and predefined callbacks (e.g., `MPI_COMM_NULL_COPY_FN`).

Rationale. For user-defined callback functions (e.g., of type `MPI_Comm_copy_attr_function` or `COMM_COPY_ATTR_FUNCTION`) and their predefined callbacks (e.g., `MPI_COMM_NULL_COPY_FN`), the `ierror` argument is not optional. The MPI library must always call these routines with an actual `ierror` argument. Therefore, these user-defined functions need not check whether the MPI library calls these routines with or without an actual `ierror` output argument. (*End of rationale.*)

The MPI Fortran bindings in the `mpi_f08` module are designed based on the Fortran 2008 standard [46] together with the Technical Specification “TS 29113 Further Interoperability with C” [47] of the ISO/IEC JTC1/SC22/WG5 (Fortran) working group, which is now integrated in Fortran 2018 standard [48].

Rationale. The features in TS 29113 on further interoperability with C were decided on by ISO/IEC JTC1/SC22/WG5 and designed by PL22.3 (formerly J3) to support a higher level of integration between Fortran-specific features and C than was provided in the Fortran 2008 standard; part of this design is based on requirements from the MPI Forum to support MPI-3.0. These features became part of Fortran 2018 [48], so references to TS 29113 are obsolete, except insofar as to specify a particular feature set from Fortran 2018 or minimal requirements to a compiler.

Fortran 2018 contains the following language features that are needed for the MPI bindings in the `mpi_f08` module: assumed-type and assumed-rank. It is important that any possible actual argument can be used for such dummy arguments, e.g., scalars, arrays, assumed-shape arrays, assumed-size arrays, allocatable arrays, and with any element type, e.g., `REAL`, `CHARACTER*5`, `CHARACTER*(*)`, sequence derived types, or `BIND(C)` derived types. Especially for backward compatibility reasons, it is important that any possible actual argument in an implicit interface implementation of a choice buffer dummy argument (e.g., with the deprecated `mpif.h` without argument-checking) can be used in an implementation with assumed-type and assumed-rank argument in an explicit interface (e.g., with the `mpi_f08` module).

A further feature useful for MPI is the extension of the semantics of the `ASYNCHRONOUS` attribute: In F2003 and F2008, this attribute could be used only to protect buffers of Fortran asynchronous I/O. With TS 29113 and now Fortran 2018, this attribute also covers asynchronous communication occurring within library routines written in C.

The MPI Forum hereby wishes to acknowledge this important effort by the Fortran PL22.3 and WG5 committee. (*End of rationale.*)

19.1.3 Fortran Support Through the `mpi` Module

An MPI implementation providing a Fortran interface must provide a module named `mpi` that can be used in a Fortran program. Within all MPI function specifications, the second of the set of two Fortran routine interface specifications is provided by this module. This module must:

- Define all named MPI constants.
- Declare MPI functions that return a value.
- Provide explicit interfaces according to the Fortran routine interface specifications. This module therefore guarantees compile-time argument checking and allows positional and keyword-based argument lists. If an implementation is paired with a compiler that either does not support `TYPE(*)`, `DIMENSION(..)` from Fortran 2018, or is otherwise unable to ignore the types of choice buffers, then the implementation must provide explicit interfaces only for MPI routines with no choice buffer arguments. See Section 19.1.6 for more details.
- Define all MPI handles as type `INTEGER`.
- Define the derived type `MPI_Status` and all named handle types that are used in the `mpi_f08` module. For these named handle types, overload the operators `.EQ.` and `.NE.` to allow handle comparison via the `.EQ.`, `.NE.`, `==` and `/=` operators.

Rationale. They are needed only when the application converts old-style `INTEGER` handles into new-style handles with a named type. (*End of rationale.*)

- A high quality MPI implementation may enhance the interface by using the `ASYNCHRONOUS` attribute in the same way as in the `mpi_f08` module if it is supported by the underlying compiler.
- Set the LOGICAL constant `MPI_ASYNC_PROTECTS_NONBLOCKING` to `.TRUE.` if the `ASYNCHRONOUS` attribute is used in all nonblocking interfaces **and** the underlying Fortran compiler supports the `ASYNCHRONOUS` attribute for MPI communication (as part of Fortran 2018), otherwise to `.FALSE..`

Advice to users. For an MPI implementation that fully supports nonblocking calls with the `ASYNCHRONOUS` attribute for choice buffers, an existing MPI-2.2 application may fail to compile even if it compiled and executed with expected results with an MPI-2.2 implementation. One reason may be that the application uses “contiguous” but not “simply contiguous” `ASYNCHRONOUS` arrays as actual arguments for choice buffers of nonblocking routines, e.g., by using subscript triplets with stride one or specifying

(1:n) for a whole dimension instead of using (:). This should be fixed to fulfill the Fortran constraints for ASYNCHRONOUS dummy arguments. This is not considered a violation of backward compatibility because existing applications can not use the ASYNCHRONOUS attribute to protect nonblocking calls. Another reason may be that the application does not conform either to the MPI standard or to the Fortran standard, typically because the program forces the compiler to perform copy-in/out for a choice buffer argument in a nonblocking MPI call. This is also not a violation of backward compatibility because the application itself is nonconforming. See Section 19.1.12 for more details. (*End of advice to users.*)

- A high quality MPI implementation may enhance the interface by using TYPE(*), DIMENSION(..) choice buffer dummy arguments instead of using nonstandardized extensions such as !\$PRAGMA IGNORE_TKR or a set of overloaded functions as described by M. Hennecke in [33], if the compiler supports this Fortran 2018 language feature. See Section 19.1.6 for further details.
- Set the LOGICAL constant MPI_SUBARRAYS_SUPPORTED to .TRUE. if all choice buffer arguments in all nonblocking, split collective and persistent communication routines are declared with TYPE(*), DIMENSION(..), otherwise set it to .FALSE.. When MPI_SUBARRAYS_SUPPORTED is defined as .TRUE., noncontiguous sub-arrays can be used as buffers in nonblocking routines.
- Set the MPI_SUBARRAYS_SUPPORTED constant to .FALSE. and declare choice buffers with a compiler-dependent mechanism that overrides type checking if the underlying Fortran compiler does not support the Fortran 2018 assumed-type and assumed-rank features. In this case, the use of noncontiguous sub-arrays in nonblocking calls may be disallowed. See Section 19.1.6 for details.

An MPI implementation may provide other features in the mpi module that enhance the usability of MPI while maintaining adherence to the standard. For example, it may provide INTENT information in these interface blocks.

Advice to implementors. The appropriate INTENT may be different from what is given in the MPI language-neutral bindings. Implementations must choose INTENT so that the function adheres to the MPI standard, e.g., by defining the INTENT as provided in the mpi_f08 bindings. (*End of advice to implementors.*)

Rationale. The intent given by the MPI generic interface is not precisely defined and does not in all cases correspond to the correct Fortran INTENT. For instance, receiving into a buffer specified by a datatype with absolute addresses may require associating MPI_BOTTOM with a dummy OUT argument. Moreover, “constants” such as MPI_BOTTOM and MPI_STATUS_IGNORE are not constants as defined by Fortran, but “special addresses” used in a nonstandard way. Finally, the MPI-1 generic intent was changed in several places in MPI-2. For instance, MPI_IN_PLACE changes the intent of an OUT argument to be INOUT. (*End of rationale.*)

Advice to implementors. The Fortran 2008 standard illustrates in its Note 5.17 that “INTENT(OUT) means that the value of the argument after invoking the procedure is entirely the result of executing that procedure. If an argument should retain

its value rather than being redefined, `INTENT(INOUT)` should be used rather than `INTENT(OUT)`, even if there is no explicit reference to the value of the dummy argument. Furthermore, `INTENT(INOUT)` is not equivalent to omitting the `INTENT` attribute, because `INTENT(INOUT)` always requires that the associated actual argument is definable.” Applications that include the (deprecated) `mpif.h` may not expect that `INTENT(OUT)` is used. In particular, output array arguments are expected to keep their content as long as the MPI routine does not modify them. To keep this behavior, it is recommended that implementations not use `INTENT(OUT)` in the `mpi` module and the (deprecated) `mpif.h` include file, even though `INTENT(OUT)` is specified in an interface description of the `mpi_f08` module. (*End of advice to implementors.*)

19.1.4 Fortran Support Through the `mpif.h` Include File

The use of the `mpif.h` include file has been deprecated in MPI-4.1.

An MPI implementation providing a Fortran interface must provide an include file named `mpif.h` that can be used in a Fortran program. Within all MPI function specifications, the second of the set of two Fortran routine interface specifications is supported by this include file. This include file must:

- Define all named MPI constants.
- Declare MPI functions that return a value.
- Define all handles as `INTEGER`.
- Be valid and equivalent for both fixed and free source form.

For each MPI routine, an implementation can choose to use an implicit or explicit interface for the second Fortran binding (in deprecated routines, the first one may be omitted).

- Set the `LOGICAL` constants `MPI_SUBARRAYS_SUPPORTED` and `MPI_ASYNC_PROTECTS_NONBLOCKING` according to the same rules as for the `mpi` module. In the case of implicit interfaces for choice buffer or nonblocking routines, the constants must be set to `.FALSE..`

Advice to users. Instead of using `mpif.h`, the use of the `mpi_f08` or `mpi` module is strongly encouraged for the following reasons:

- Most `mpif.h` implementations do not include compile-time argument checking.
- Therefore, many bugs in MPI applications remain undetected at compile-time, such as:
 - Missing `ierror` as last argument in most Fortran bindings.
 - Declaration of a `status` as an `INTEGER` variable instead of an `INTEGER` array with size `MPI_STATUS_SIZE`.
 - Incorrect argument positions; e.g., interchanging the `count` and `datatype` arguments.
 - Passing incorrect MPI handles; e.g., passing a `datatype` instead of a communicator.

- The migration from `mpif.h` to the `mpi` module should be relatively straightforward (i.e., substituting `INCLUDE 'mpif.h'` after an `implicit` statement by `use mpi` before that `implicit` statement) as long as the application syntax is correct.
- Migrating portable and correctly written applications to the `mpi` module is not expected to be difficult. No compile or runtime problems should occur because an `mpif.h` include file was always allowed to provide explicit Fortran interfaces.

(End of advice to users.)

19.1.5 Interface Specifications, Procedure Names, and the Profiling Interface

The Fortran interface specification of each MPI routine specifies the routine name that must be called by the application program, and the names and types of the dummy arguments together with additional attributes. The Fortran standard allows a given Fortran interface to be implemented with several methods, e.g., within or outside of a module, with or without `BIND(C)`, or the buffers with or without Fortran 2018 (as successor of Fortran 2008 with TS 29113). Such implementation decisions imply different binary interfaces and different specific procedure names. The requirements for several implementation schemes together with the rules for the specific procedure names and its implications for the profiling interface are specified within this section, but not the implementation details.

Rationale. When this section was originally introduced in MPI-3.0, the major goals for the three Fortran support methods were:

- Portable implementation of the wrappers from the MPI Fortran interfaces to the MPI routines in C.
- Binary backward compatible implementation path when switching `MPI_SUBARRAYS_SUPPORTED` from `.FALSE.` to `.TRUE.`.
- The Fortran PMPI interface need not be backward compatible, but a method must be included that a tools layer can use to examine the MPI library about the specific procedure names and interfaces used.
- No performance drawbacks.
- Consistency between all three Fortran support methods.
- Consistent with Fortran 2018.

The design expected that all dummy arguments in the MPI Fortran interfaces are interoperable with C according to Fortran 2018. This expectation was not fulfilled. The `LOGICAL` arguments are not interoperable with C, mainly because the internal representations for `.FALSE.` and `.TRUE.` are compiler dependent. The provided interface was mainly based on `BIND(C)` interfaces and therefore inconsistent with Fortran. To be consistent with Fortran, the `BIND(C)` had to be removed from the callback procedure interfaces and the predefined callbacks, e.g., `MPI_COMM_DUP_FN`. Non-`BIND(C)` procedures are also not interoperable with C, and therefore the `BIND(C)` had to be removed from all routines with `PROCEDURE` arguments, e.g., from `MPI_OP_CREATE`.

Therefore, this section was rewritten as an erratum to MPI-3.0. *(End of rationale.)*

A Fortran call to an MPI routine shall result in a call to a procedure with one of the specific procedure names and calling conventions, as described in Table 19.1. Case is not significant in the names.

Table 19.1: Specific Fortran procedure names and related calling conventions. `MPI_ISEND` is used as an example. For routines without choice buffers, only 1A and 2A apply.

No.	Specific procedure name	Calling convention
1A	<code>MPI_Isend_f08</code>	Fortran interface and arguments, as in Annex A.4, except that in routines with a choice buffer dummy argument, this dummy argument is implemented with nonstandard extensions like <code>!\$PRAGMA IGNORE_TKR</code> , which provides a call-by-reference argument without type, kind, and dimension checking.
1B	<code>MPI_Isend_f08ts</code>	Fortran interface and arguments, as in Annex A.4, but only for routines with one or more choice buffer dummy arguments; these dummy arguments are implemented with <code>TYPE(*)</code> , <code>DIMENSION(..)</code> .
2A	<code>MPI_ISEND</code>	Fortran interface and arguments, as in Annex A.5, except that in routines with a choice buffer dummy argument, this dummy argument is implemented with nonstandard extensions like <code>!\$PRAGMA IGNORE_TKR</code> , which provides a call-by-reference argument without type, kind, and dimension checking.
2B	<code>MPI_ISEND_FTS</code>	Fortran interface and arguments, as in Annex A.5, but only for routines with one or more choice buffer dummy arguments; these dummy arguments are implemented with <code>TYPE(*)</code> , <code>DIMENSION(..)</code> . In the (deprecated) <code>mpif.h</code> only, the postfix “_FTS” for <code>MPI_NEIGHBOR_ALLGATHERV_INIT</code> , <code>MPI_NEIGHBOR_ALLTOALLV_INIT</code> , and <code>MPI_NEIGHBOR_ALLTOALLW_INIT</code> is shortened to “_F”.

1 Note that for the deprecated routines in Section 16.1, which are reported only in An-
 2 nex A.5, scheme 2A is utilized in the `mpi` module and (deprecated) `mpif.h`, and also in the
 3 `mpi_f08` module.

4 To set `MPI_SUBARRAYS_SUPPORTED` to `.TRUE.` within a Fortran support method, it
 5 is required that all nonblocking and split-collective routines with buffer arguments are im-
 6 plemented according to 1B and 2B, i.e., with `MPI_Xxxx_f08ts` in the `mpi_f08` module, and
 7 with `MPI_XXXX_FTS` in the `mpi` module and the (deprecated) `mpif.h` include file.

8 The `mpi` and `mpi_f08` modules and the (deprecated) `mpif.h` include file will each cor-
 9 respond to exactly one implementation scheme from Table 19.1. However, the MPI library
 10 may contain multiple implementation schemes from Table 19.1.

11 *Advice to implementors.* This may be desirable for backwards binary compatibility
 12 in the scope of a single MPI implementation, for example. (*End of advice to imple-*
 13 *mentors.*)

14 *Rationale.* After a compiler provides the facilities `TYPE(*)`, `DIMENSION(..)` from
 15 Fortran 2018, it is possible to change the bindings within a Fortran support method
 16 to support subarrays without recompiling the complete application provided that
 17 the previous interfaces with their specific procedure names are still included in the
 18 library. Of course, only recompiled routines can benefit from the added facilities.
 19 There is no binary compatibility conflict because each interface uses its own spe-
 20 cific procedure names and all interfaces use the same constants (except the value of
 21 `MPI_SUBARRAYS_SUPPORTED` and `MPI_ASYNC_PROTECTS_NONBLOCKING`) and type
 22 definitions. After a compiler also ensures that buffer arguments of nonblocking MPI
 23 operations can be protected through the `ASYNCHRONOUS` attribute, and the proce-
 24 dure declarations in the `mpi_f08` and `mpi` module and the (deprecated) `mpif.h` in-
 25 clude file declare choice buffers with the `ASYNCHRONOUS` attribute, then the value of
 26 `MPI_ASYNC_PROTECTS_NONBLOCKING` can be switched to `.TRUE.` in the module def-
 27 inition and include file. (*End of rationale.*)

28 *Advice to users.* Partial recompilation of user applications when upgrading MPI
 29 implementations is a highly complex and subtle topic. Users are strongly advised to
 30 consult their MPI implementation’s documentation to see exactly what is—and what
 31 is not—supported. (*End of advice to users.*)

32 Within the `mpi_f08` and `mpi` modules and (deprecated) `mpif.h` include file, for all
 33 MPI procedures, a second procedure with the same calling conventions shall be supplied,
 34 except that the name is modified by prefixing with the letter “P”, e.g., `PMPI_Isend`. The
 35 specific procedure names for these `PMPI_Xxxx` procedures must be different from the specific
 36 procedure names for the `MPI_Xxxx` procedures and are not specified by this standard.

37 A user-written or middleware profiling routine should provide the same specific Fortran
 38 procedure names and calling conventions, and therefore can interpose itself as the MPI
 39 library routine. The profiling routine can internally call the matching
 40 `PMPI` routine with any of its existing bindings, except for routines that have callback routine
 41 dummy arguments, choice buffer arguments, or that are attribute caching routines (
 42 `MPI_{COMM|WIN|TYPE}_{SET|GET}_ATTR`). In this case, the profiling software should
 43 invoke the corresponding `PMPI` routine using the same Fortran support method as used in
 44 the calling application program, because the C, `mpi_f08` and `mpi` callback prototypes are
 45 different or the meaning of the choice buffer or `attribute_val` arguments are different.
 46
 47
 48

Advice to users. Although for each support method and MPI routine (e.g., MPI_ISEND in `mpi_f08`), multiple routines may need to be provided to intercept the specific procedures in the MPI library (e.g., `MPI_Isend_f08` and `MPI_Isend_f08ts`), each profiling routine itself uses only one support method (e.g., `mpi_f08`) and calls the real MPI routine through the one PMPI routine defined in this support method (i.e., `PMPI_Isend` in this example). (*End of advice to users.*)

Advice to implementors. If all of the following conditions are fulfilled:

- the handles in the `mpi_f08` module occupy one Fortran numerical storage unit (same as an `INTEGER` handle),
- the internal argument passing mechanism used to pass an actual `ierror` argument to a nonoptional `ierror` dummy argument is binary compatible to passing an actual `ierror` argument to an `ierror` dummy argument that is declared as `OPTIONAL`,
- the internal argument passing mechanism for `ASYNCHRONOUS` and non-`ASYNCHRONOUS` arguments is the same,
- the internal routine call mechanism is the same for the Fortran and the C compilers for which the MPI library is compiled, and
- the compiler does not provide the appropriate features from Fortran 2018,

then the implementor may use the same internal routine implementations for all Fortran support methods but with several different specific procedure names. If the accompanying Fortran compiler supports Fortran 2018 or at least Fortran 2008 with TS 29113, then the new routines are needed only for routines with choice buffer arguments. (*End of advice to implementors.*)

Advice to implementors. In the (deprecated) Fortran support method `mpif.h`, compile-time argument checking can be also implemented for all routines. For `mpif.h`, the argument names are not specified through the MPI standard, i.e., only positional argument lists are defined, and not key-word based lists. Due to the rule that `mpif.h` must be valid for fixed and free source form, the subroutine declaration is restricted to one line with 72 characters. To keep the argument lists short, each argument name can be shortened to a minimum of one character. With this, the three longest subroutine declaration statements are

```
SUBROUTINE PMPI_DIST_GRAPH_CREATE_ADJACENT(a,b,c,d,e,f,g,h,i,j,k)
SUBROUTINE PMPI_NEIGHBOR_ALLTOALLW_INIT(a,b,c,d,e,f,g,h,i,j,k,l)
SUBROUTINE PMPI_NEIGHBOR_ALLTOALLV_INIT(a,b,c,d,e,f,g,h,i,j,k,l)
```

with 71 and 70 characters each. With buffers implemented with Fortran 2018 (or TS 29113), the specific procedure names have an additional postfix. Some of the longest of such interface definitions are

```
INTERFACE PMPI_NEIGHBOR_ALLTOALLW_INIT
SUBROUTINE PMPI_NEIGHBOR_ALLTOALLW_INIT_F(a,b,c,d,e,f,g,h,i,j,j,k)
INTERFACE PMPI_NEIGHBOR_ALLGATHERV_INIT
SUBROUTINE PMPI_NEIGHBOR_ALLGATHERV_INIT_F(a,b,c,d,e,f,g,h,i,j,k)
INTERFACE PMPI_RGET_ACCUMULATE
SUBROUTINE PMPI_RGET_ACCUMULATE_FTS(a,b,c,d,e,f,g,h,i,j,k,l,m,n)
```

with 72, 71, and 70 characters. In principle, continuation lines would be possible in `mpif.h` (spaces in columns 73–131, & in column 132, and in column 6 of the continuation line) but this would not be valid if the source line length is extended with a compiler flag to 132 characters. Column 133 is also not available for the continuation character because lines longer than 132 characters are invalid with some compilers by default.

If an implementation applies the rules of Table 19.1 also for the PMPI interface, then the longest specific procedure name is `PMPI_Reduce_scatter_block_init_c_f08ts` with 38 characters in the `mpi_f08` module.

For example, the interface specifications together with the specific procedure names can be implemented with

```

13  MODULE mpi_f08
14      TYPE, BIND(C) :: MPI_Comm
15      INTEGER :: MPI_VAL
16      END TYPE MPI_Comm
17      ...
18      INTERFACE MPI_Comm_rank ! (as defined in Chapter 6)
19          SUBROUTINE MPI_Comm_rank_f08(comm, rank, ierror)
20              IMPORT :: MPI_Comm
21              TYPE(MPI_Comm),      INTENT(IN)  :: comm
22              INTEGER,              INTENT(OUT) :: rank
23              INTEGER, OPTIONAL,   INTENT(OUT) :: ierror
24          END SUBROUTINE
25      END INTERFACE
26  END MODULE mpi_f08
27
28  MODULE mpi
29      INTERFACE MPI_Comm_rank ! (as defined in Chapter 6)
30          SUBROUTINE MPI_Comm_rank(comm, rank, ierror)
31              INTEGER, INTENT(IN)  :: comm ! The INTENT may be added although
32              INTEGER, INTENT(OUT) :: rank ! it is not defined in the
33              INTEGER, INTENT(OUT) :: ierror ! official routine definition.
34          END SUBROUTINE
35      END INTERFACE
36  END MODULE mpi

```

And if interfaces are provided in `mpif.h`, they might look like this (outside of any module and in fixed source format):

```

37  !23456789012345678901234567890123456789012345678901234567890123456789012
38  INTERFACE MPI_Comm_rank ! (as defined in Chapter 6)
39  SUBROUTINE MPI_Comm_rank(comm, rank, ierror)
40  INTEGER, INTENT(IN)  :: comm ! The argument names may be
41  INTEGER, INTENT(OUT) :: rank ! shortened so that the
42  INTEGER, INTENT(OUT) :: ierror ! subroutine line fits to the
43  END SUBROUTINE ! maximum of 72 characters.
44  END INTERFACE

```

(End of advice to implementors.)

Advice to users. The following is an example of how a user-written or middleware profiling routine can be implemented:

```

45  SUBROUTINE MPI_Isend_f08ts(buf, count, datatype, dest, tag, comm, request, ierror)
46  USE :: mpi_f08, my_noname => MPI_Isend_f08ts

```



```

TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, dest, tag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
! ... some code for the begin of profiling
call PMPI_Isend (buf, count, datatype, dest, tag, comm, request, ierror)
! ... some code for the end of profiling
END SUBROUTINE MPI_Isend_f08ts

```

Note that this routine is used to intercept the existing specific procedure name `MPI_Isend_f08ts` in the MPI library. This routine must not be part of a module. This routine itself calls `PMPI_Isend`. The USE of the `mpi_f08` module is needed for definitions of handle types and the interface for `PMPI_Isend`. However, this module also contains an interface definition for the specific procedure name `MPI_Isend_f08ts` that conflicts with the definition of this profiling routine (i.e., the name is doubly defined). Therefore, the USE here specifically excludes the interface from the module by renaming the unused routine name in the `mpi_f08` module into “`my_noname`” in the scope of this routine. (*End of advice to users.*)

Advice to users. The PMPI interface allows intercepting MPI routines. For example, an additional MPI_ISEND profiling wrapper can be provided that is called by the application and internally calls PMPI_ISEND. There are two typical use cases: a profiling layer that is developed independently from the application and the MPI library, and profiling routines that are part of the application and have access to the application data. With MPI-3.0, new Fortran interfaces and implementation schemes were introduced that have several implications on how Fortran MPI routines are internally implemented and optimized. For profiling layers, these schemes imply that several internal interfaces with different specific procedure names may need to be intercepted, as shown in the example code above. Therefore, for wrapper routines that are part of a Fortran application, it may be more convenient to make the name shift within the application, i.e., to substitute the call to the MPI routine (e.g., `MPI_ISEND`) by a call to a user-written profiling wrapper with a new name (e.g., `X_MPI_ISEND`) and to call the Fortran `MPI_ISEND` from this wrapper, instead of using the PMPI interface. (*End of advice to users.*)

Advice to implementors. An implementation that provides a Fortran interface must provide a combination of MPI library and module or include file that uses the specific procedure names as described in Table 19.1 so that the MPI Fortran routines are interceptable as described above. (*End of advice to implementors.*)

19.1.6 MPI for Different Fortran Standard Versions

This section describes which Fortran interface functionality can be provided for different versions of the Fortran standard.

- *For Fortran 77* with some extensions:
 - MPI identifiers may be up to 30 characters (31 with the profiling interface).
 - MPI identifiers may contain underscores after the first character.

- 1 – An MPI subroutine with a choice argument may be called with different argument
- 2 types.
- 3 – Although not required by the MPI standard, the `INCLUDE` statement should be
- 4 available for including `mpif.h` into the user application source code.

5 Only MPI-1.1, MPI-1.2, and MPI-1.3 can be implemented. The use of absolute ad-
 6 resses from `MPI_ADDRESS` and `MPI_BOTTOM` may cause problems if an address
 7 does not fit into the memory space provided by an `INTEGER`. (In MPI-2.0 this problem
 8 is solved with `MPI_GET_ADDRESS`, but not for Fortran 77.)

9
10 • *For Fortran 90:*

11 The major additional features that are needed from Fortran 90 are:

- 12 – The `MODULE` and `INTERFACE` concept.
- 13 – The `KIND=` and `SELECTED_XXX_KIND` concept.
- 14 – Fortran derived `TYPE`s and the `SEQUENCE` attribute.
- 15 – The `OPTIONAL` attribute for dummy arguments.
- 16 – Cray pointers, which are a nonstandard compiler extension, are needed for the
- 17 use of `MPI_ALLOC_MEM`.

18 With these features, MPI-1.1 – MPI-2.2 can be implemented without restrictions.
 19 MPI-3.0 and later can be implemented with some restrictions. The Fortran support
 20 methods are abbreviated with `S1` = the `mpi_f08` module, `S2` = the `mpi` module, and `S3`
 21 = the `mpif.f` include file. If not stated otherwise, restrictions exist for each method
 22 that prevent implementing the complete semantics of MPI.

- 23 – `MPI_SUBARRAYS_SUPPORTED` equals `.FALSE.`, i.e., subscript triplets and non-
 24 contiguous subarrays cannot be used as buffers in nonblocking routines, RMA,
 25 or split-collective I/O.
- 26 – `S1`, `S2`, and `S3` can be implemented, but for `S1`, only a preliminary implementa-
 27 tion is possible.
- 28 – In this preliminary interface of `S1`, the following changes are necessary:
 29 * `TYPE(*)`, `DIMENSION(..)` is substituted by nonstandardized extensions like
 30 `!$PRAGMA IGNORE_TKR`.
- 31 * The `ASYNCHRONOUS` attribute is omitted.
- 32 * `PROCEDURE(...)` callback declarations are substituted by `EXTERNAL`.
- 33 – The specific procedure names are specified in Section 19.1.5.
- 34 – Due to the rules specified in Section 19.1.5, choice buffer declarations should be
 35 implemented only with nonstandardized extensions like `!$PRAGMA IGNORE_TKR`
 36 (as long as F2008 with TS 29113 or Fortran 2018 is not available).

37 In `S2` and `S3`: Without such extensions, routines with choice buffers should be
 38 provided with an implicit interface, instead of overloading with a different MPI
 39 function for each possible buffer type (as mentioned in Section 19.1.11). Such
 40 overloading would also imply restrictions for passing Fortran derived types as
 41 choice buffer, see also Section 19.1.15.

Only in S1: The implicit interfaces for routines with choice buffer arguments imply that the `ieror` argument cannot be defined as `OPTIONAL`. For this reason, it is recommended not to provide the `mpi_f08` module if such an extension is not available.

- The `ASYNCHRONOUS` attribute can **not** be used in applications to protect buffers in nonblocking MPI calls (S1–S3).
 - The `TYPE(C_PTR)` binding of the `MPI_ALLOC_MEM` and `MPI_WIN_ALLOCATE` routines is not available.
 - In S1 and S2, the definition of the handle types (e.g., `TYPE(MPI_Comm)` and the status type `TYPE(MPI_Status)` must be modified: The `SEQUENCE` attribute must be used instead of `BIND(C)` (which is not available in Fortran 90/95). This restriction implies that the application must be fully recompiled if one switches to an MPI library for Fortran 2003 and later because the internal memory size of the handles may have changed. For this reason, an implementor may choose not to provide the `mpi_f08` module for Fortran 90 compilers. In this case, the `mpi_f08` handle types and all routines, constants and types related to `TYPE(MPI_Status)` (see Section 19.3.5) are also not available in the `mpi` module and `mpif.h`.
- *For Fortran 95:*
The quality of the MPI interface and the restrictions are the same as with Fortran 90.

- *For Fortran 2003:*

The major features that are needed from Fortran 2003 are:

- Interoperability with C, i.e.,
 - * `BIND(C)` derived types.
 - * The `ISO_C_BINDING` intrinsic type `C_PTR` and routine `C_F_POINTER`.
- The ability to define an `ABSTRACT INTERFACE` and to use it for `PROCEDURE` dummy arguments.
- The ability to overload the operators `.EQ.` and `.NE.` to allow the comparison of derived types (used in MPI-3.0 and later for MPI handles).
- The `ASYNCHRONOUS` attribute is available to protect Fortran asynchronous I/O. This feature is not yet used by MPI, but it is the basis for the enhancement for MPI communication in the TS 29113.

With these features (but still without the features of TS 29113), MPI-1.1 – MPI-2.2 can be implemented without restrictions, but with one enhancement:

- The user application can use `TYPE(C_PTR)` together with `MPI_ALLOC_MEM` as long as `MPI_ALLOC_MEM` is defined with an implicit interface because a `C_PTR` and an `INTEGER(KIND=MPI_ADDRESS_KIND)` argument must both map to a `void *` argument.

MPI-3.0 and later can be implemented with the following restrictions:

- `MPI_SUBARRAYS_SUPPORTED` equals `.FALSE.`
- For S1, only a preliminary implementation is possible. The following changes are necessary:

- 1 * TYPE(*), DIMENSION(..) is substituted by nonstandardized extensions like
- 2 !\$PRAGMA IGNORE_TKR.
- 3
- 4 – The specific procedure names are specified in Section 19.1.5.
- 5 – With S1, the ASYNCHRONOUS is required as specified in the second Fortran inter-
- 6 – faces. With S2 and S3 the implementation can also add this attribute if explicit
- 7 – interfaces are used.
- 8 – The ASYNCHRONOUS Fortran attribute can be used in applications to *try to* protect
- 9 – buffers in nonblocking MPI calls, but the protection can work only if the compiler
- 10 – is able to protect asynchronous Fortran I/O and makes no difference between such
- 11 – asynchronous Fortran I/O and MPI communication.
- 12 – The TYPE(C_PTR) binding of the MPI_ALLOC_MEM, MPI_WIN_ALLOCATE,
- 13 – MPI_WIN_ALLOCATE_SHARED, and MPI_WIN_SHARED_QUERY routines can
- 14 – be used only for Fortran types that are C compatible.
- 15 – The same restriction as for Fortran 90 applies if nonstandardized extensions like
- 16 – !\$PRAGMA IGNORE_TKR are not available.
- 17

- 18 • For Fortran 2008 with TS 29113 and later and
- 19 For Fortran 2003 with TS 29113:

20 The major features that are needed from TS 29113 are:

- 21 – TYPE(*), DIMENSION(..) is available.
- 22 – The ASYNCHRONOUS attribute is extended to protect also nonblocking MPI com-
- 23 – munication.
- 24 – The array dummy argument of the ISO_C_BINDING intrinsic C_F_POINTER is not
- 25 – restricted to Fortran types for which a corresponding type in C exists.
- 26
- 27

28 Using these features, MPI-3.0 and later can be implemented without any restrictions.

- 29 – With S1, MPI_SUBARRAYS_SUPPORTED equals .TRUE.. The
- 30 – ASYNCHRONOUS attribute can be used to protect buffers in nonblocking MPI calls.
- 31 – The TYPE(C_PTR) binding of the MPI_ALLOC_MEM, MPI_WIN_ALLOCATE,
- 32 – MPI_WIN_ALLOCATE_SHARED, and MPI_WIN_SHARED_QUERY routines can
- 33 – be used for any Fortran type.
- 34 – With S2 and S3, the value of MPI_SUBARRAYS_SUPPORTED is implementation
- 35 – dependent. A high quality implementation will also provide
- 36 – MPI_SUBARRAYS_SUPPORTED set to .TRUE. and will use the ASYNCHRONOUS at-
- 37 – tribute in the same way as in S1.
- 38 – If nonstandardized extensions like !\$PRAGMA IGNORE_TKR are not available then
- 39 – S2 must be implemented with TYPE(*), DIMENSION(..).
- 40

41 *Advice to implementors.* If MPI_SUBARRAYS_SUPPORTED=.FALSE., the choice ar-

42 gument may be implemented with an explicit interface using compiler directives, for

43 example:

```
44 INTERFACE
45 SUBROUTINE MPI_...(buf, ...)
46   !DEC$ ATTRIBUTES NO_ARG_CHECK :: buf
47   !$PRAGMA IGNORE_TKR buf
48   !DIR$ IGNORE_TKR buf
```

```

!IBM* IGNORE_TKR buf
REAL, DIMENSION(*) :: buf
... ! declarations of the other arguments
END SUBROUTINE
END INTERFACE

```

(*End of advice to implementors.*)

19.1.7 Requirements on Fortran Compilers

MPI-3.0 (and later) compliant Fortran bindings are not only a property of the MPI library itself, but rather a property of an MPI library together with the Fortran compiler suite for which it is compiled.

Advice to users. Users must take appropriate steps to ensure that proper options are specified to compilers. MPI libraries must document these options. Some MPI libraries are shipped together with special compilation scripts (e.g., `mpif90`, `mpicc`) that set these options automatically. (*End of advice to users.*)

An MPI library together with the Fortran compiler suite is only compliant with MPI-3.0 (and later), as referred by `MPI_GET_VERSION`, if all the solutions described in Sections 19.1.11 through 19.1.19 work correctly. Based on this rule, major requirements for all three Fortran support methods (i.e., the `mpi_f08` and `mpi` modules, and `mpif.h`) are:

- The language features assumed-type and assumed-rank from Fortran 2008 TS 29113 [47] are available. This is required only for `mpi_f08`. As long as this requirement is not supported by the compiler, it is valid to build an MPI library that implements the `mpi_f08` module with `MPI_SUBARRAYS_SUPPORTED` set to `.FALSE.`.
- “Simply contiguous” arrays and scalars must be passed to choice buffer dummy arguments of nonblocking routines with call by reference. This is needed only if one of the support methods does not use the `ASYNCHRONOUS` attribute. See Section 19.1.12 for more details.
- `SEQUENCE` and `BIND(C)` derived types are valid as actual arguments passed to choice buffer dummy arguments, and, in the case of `MPI_SUBARRAYS_SUPPORTED` set to `.FALSE.`, they are passed with call by reference, and passed by descriptor in the case of `.TRUE.`.
- All actual arguments that are allowed for a dummy argument in an implicitly defined and separately compiled Fortran routine with the given compiler (e.g., `CHARACTER(LEN=*)` strings and array of strings) must also be valid for choice buffer dummy arguments with all Fortran support methods.
- The array dummy argument of the `ISO_C_BINDING` intrinsic module procedure `C_F_POINTER` is not restricted to Fortran types for which a corresponding type in C exists.
- The Fortran compiler shall not provide `TYPE(*)` unless the `ASYNCHRONOUS` attribute protects MPI communication as described in TS 29113. Specifically, the TS 29113 must be implemented as a whole.

The following rules are required at least as long as the compiler does not provide the extension of the ASYNCHRONOUS attribute as part of TS 29113 and there still exists a Fortran support method with MPI_ASYNC_PROTECTS_NONBLOCKING set to .FALSE.. Observation of these rules by the MPI application developer is especially recommended for backward compatibility of existing applications that use the `mpi` module or the (deprecated) `mpif.h` include file. The rules are as follows:

- Separately compiled empty Fortran routines with implicit interfaces and separately compiled empty C routines with BIND(C) Fortran interfaces (e.g., MPI_F_SYNC_REG on page 826 and Section 19.1.8, and DD on page 827) solve the problems described in Section 19.1.17.
- The problems with temporary data movement (described in detail in Section 19.1.18) are solved as long as the application uses different sets of variables for the nonblocking communication (or nonblocking or split collective I/O) and the computation when overlapping communication and computation.
- Problems caused by automatic and permanent data movement (e.g., within a garbage collection, see Section 19.1.19) are resolved **without** any further requirements on the application program, neither on the usage of the buffers, nor on the declaration of application routines that are involved in invoking MPI procedures.

All of these rules are valid for the `mpi_f08` and `mpi` modules and independently of whether `mpif.h` uses explicit interfaces.

Advice to implementors. Some of these rules are already part of the Fortran 2003 standard, some of these requirements require the Fortran TS 29113 [47], and some of these requirements for MPI are beyond the scope of TS 29113. (*End of advice to implementors.*)

19.1.8 Additional Support for Fortran Register-Memory-Synchronization

As described in Section 19.1.17, a dummy call may be necessary to tell the compiler that registers are to be flushed for a given buffer or that accesses to a buffer may not be moved across a given point in the execution sequence. Only a Fortran binding exists for this call.

MPI_F_SYNC_REG(buf)

INOUT buf initial address of buffer (choice)

Fortran 2008 binding

MPI_F_sync_reg(buf)

TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf

Fortran binding

MPI_F_SYNC_REG(BUF)

<type> BUF(*)

This routine has no executable statements. It must be compiled in the MPI library in such a manner that a Fortran compiler cannot detect in the module that the routine has an empty body. It is used only to force the compiler to flush a cached register value of a variable or buffer back to memory (when necessary), or to invalidate the register value.

Rationale. This function is not available in other languages because it would not be useful. This routine has no `ierror` return argument because there is no operation that can fail. (*End of rationale.*)

Advice to implementors. This routine can be bound to a C routine to minimize the risk that the Fortran compiler can learn that this routine is empty (and that the call to this routine can be removed as part of an optimization). However, it is explicitly allowed to implement this routine within the `mpi_f08` module according to the definition for the `mpi` module or `mpif.h` to circumvent the overhead of building the internal dope vector to handle the assumed-type, assumed-rank argument. (*End of advice to implementors.*)

Rationale. This routine is not defined with `TYPE(*)`, `DIMENSION(*)`, i.e., assumed size instead of assumed rank, because this would restrict the usability to “simply contiguous” arrays and would require overloading with another interface for scalar arguments. (*End of rationale.*)

Advice to users. If only a part of an array (e.g., defined by a subscript triplet) is used in a nonblocking routine, it is recommended to pass the whole array to `MPI_F_SYNC_REG` anyway to minimize the overhead of this no-operation call. Note that this routine need not be called if `MPI_ASYNC_PROTECTS_NONBLOCKING` is `.TRUE.` and the application fully uses the facilities of `ASYNCHRONOUS` arrays. (*End of advice to users.*)

19.1.9 Additional Support for Fortran Numeric Intrinsic Types

MPI provides a small number of named datatypes that correspond to named intrinsic types supported by C and Fortran. These include `MPI_INTEGER`, `MPI_REAL`, `MPI_INT`, `MPI_DOUBLE`, etc., as well as the optional types `MPI_REAL4`, `MPI_REAL8`, etc. There is a one-to-one correspondence between language declarations and MPI types.

Fortran (starting with Fortran 90) provides so-called `KIND`-parameterized types. These types are declared using an intrinsic type (one of `INTEGER`, `REAL`, `COMPLEX`, `LOGICAL`, and `CHARACTER`) with an optional integer `KIND` parameter that selects from among one or more variants. The specific meaning of different `KIND` values themselves are implementation dependent and not specified by the language. Fortran provides the `KIND` selection functions `selected_real_kind` for `REAL` and `COMPLEX` types, and `selected_int_kind` for `INTEGER` types that allow users to declare variables with a minimum precision or number of digits. These functions provide a portable way to declare `KIND`-parameterized `REAL`, `COMPLEX`, and `INTEGER` variables in Fortran. This scheme is backward compatible with Fortran 77. `REAL` and `INTEGER` Fortran variables have a default `KIND` if none is specified. Fortran `DOUBLE PRECISION` variables are of intrinsic type `REAL` with a nondefault `KIND`. The following two declarations are equivalent:

```
double precision x
real(KIND(0.0d0)) x
```

MPI provides two orthogonal methods for handling communication buffers of numeric intrinsic types. The first method (see the following section) can be used when variables have been declared in a portable way—using default `KIND` or using `KIND` parameters obtained with the `selected_int_kind` or `selected_real_kind` functions. With this method, MPI

1 automatically selects the correct data size (e.g., 4 or 8 bytes) and provides representation
 2 conversion in heterogeneous environments. The second method (see “Support for size-
 3 specific MPI Datatypes” on page 810) gives the user complete control over communication
 4 by exposing machine representations.

6 *Parameterized Datatypes with Specified Precision and Exponent Range*

7 MPI provides named datatypes corresponding to standard Fortran 77 numeric types:
 8 MPI_INTEGER, MPI_COMPLEX, MPI_REAL, MPI_DOUBLE_PRECISION and
 9 MPI_DOUBLE_COMPLEX. MPI automatically selects the correct data size and provides rep-
 10 resentation conversion in heterogeneous environments. The mechanism described in this
 11 section extends this model to support portable parameterized numeric types.

12 The model for supporting portable parameterized types is as follows. Real variables
 13 are declared (perhaps indirectly) using `selected_real_kind(p, r)` to determine the KIND
 14 parameter, where `p` is decimal digits of precision and `r` is an exponent range. Implicitly
 15 MPI maintains a two-dimensional array of predefined MPI datatypes `D(p, r)`. `D(p, r)` is
 16 defined for each value of `(p, r)` supported by the compiler, including pairs for which one
 17 value is unspecified. Attempting to access an element of the array with an index `(p, r)` not
 18 supported by the compiler is erroneous. MPI implicitly maintains a similar array of COMPLEX
 19 datatypes. For integers, there is a similar implicit array related to `selected_int_kind` and
 20 indexed by the requested number of digits `r`. Note that the predefined datatypes contained
 21 in these implicit arrays are not the same as the named MPI datatypes MPI_REAL, etc., but
 22 a new set.

23
 24 *Advice to implementors.* The above description is for explanatory purposes only. It
 25 is not expected that implementations will have such internal arrays. (*End of advice*
 26 *to implementors.*)

27
 28 *Advice to users.* `selected_real_kind()` maps a large number of `(p,r)` pairs to a
 29 much smaller number of KIND parameters supported by the compiler. KIND parameters
 30 are not specified by the language and are not portable. From the language point of
 31 view intrinsic types of the same base type and KIND parameter are of the same type. In
 32 order to allow interoperability in a heterogeneous environment, MPI is more stringent.
 33 The corresponding MPI datatypes match if and only if they have the same `(p,r)` value
 34 (REAL and COMPLEX) or `r` value (INTEGER). Thus MPI has many more datatypes than
 35 there are fundamental language types. (*End of advice to users.*)

36
 37
 38 MPI_TYPE_CREATE_F90_REAL(`p`, `r`, `newtype`)

39	IN	<code>p</code>	precision, in decimal digits (integer)
40	IN	<code>r</code>	decimal exponent range (integer)
41	OUT	<code>newtype</code>	the requested MPI datatype (handle)
42			

43 **C binding**

44
 45 `int MPI_Type_create_f90_real(int p, int r, MPI_Datatype *newtype)`

46 **Fortran 2008 binding**

47
 48 `MPI_Type_create_f90_real(p, r, newtype, ierror)`


```

INTEGER, INTENT(IN) :: p, r
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TYPE_CREATE_F90_REAL(P, R, NEWTYPE, IERROR)
INTEGER P, R, NEWTYPE, IERROR

```

This function returns a predefined MPI datatype that matches a REAL variable of KIND `selected_real_kind(p, r)`. In the model described above it returns a handle for the element `D(p, r)`. Either `p` or `r` may be omitted from calls to `selected_real_kind(p, r)` (but not both). Analogously, either `p` or `r` may be set to `MPI_UNDEFINED`. In communication, an MPI datatype `A` returned by `MPI_TYPE_CREATE_F90_REAL` matches a datatype `B` if and only if `B` was returned by `MPI_TYPE_CREATE_F90_REAL` called with the same values for `p` and `r` or `B` is a duplicate of such a datatype. Restrictions on using the returned datatype with the "external32" data representation are given on page 809.

It is erroneous to supply values for `p` and `r` not supported by the compiler.

```

MPI_TYPE_CREATE_F90_COMPLEX(p, r, newtype)

```

IN	<code>p</code>	precision, in decimal digits (integer)
IN	<code>r</code>	decimal exponent range (integer)
OUT	<code>newtype</code>	the requested MPI datatype (handle)

C binding

```

int MPI_Type_create_f90_complex(int p, int r, MPI_Datatype *newtype)

```

Fortran 2008 binding

```

MPI_Type_create_f90_complex(p, r, newtype, ierror)
INTEGER, INTENT(IN) :: p, r
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

Fortran binding

```

MPI_TYPE_CREATE_F90_COMPLEX(P, R, NEWTYPE, IERROR)
INTEGER P, R, NEWTYPE, IERROR

```

This function returns a predefined MPI datatype that matches a COMPLEX variable of KIND `selected_real_kind(p, r)`. Either `p` or `r` may be omitted from calls to `selected_real_kind(p, r)` (but not both). Analogously, either `p` or `r` may be set to `MPI_UNDEFINED`. Matching rules for datatypes created by this function are analogous to the matching rules for datatypes created by `MPI_TYPE_CREATE_F90_REAL`. Restrictions on using the returned datatype with the "external32" data representation are given on page 809.

It is erroneous to supply values for `p` and `r` not supported by the compiler.

```

1 MPI_TYPE_CREATE_F90_INTEGER(r, newtype)
2   IN      r                decimal exponent range, i.e., number of decimal
3                        digits (integer)
4
5   OUT     newtype          the requested MPI datatype (handle)
6

```

C binding

```

8 int MPI_Type_create_f90_integer(int r, MPI_Datatype *newtype)
9

```

Fortran 2008 binding

```

10 MPI_Type_create_f90_integer(r, newtype, ierror)
11   INTEGER, INTENT(IN) :: r
12   TYPE(MPI_Datatype), INTENT(OUT) :: newtype
13   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14

```

Fortran binding

```

15 MPI_TYPE_CREATE_F90_INTEGER(R, NEWTYPE, IERROR)
16   INTEGER R, NEWTYPE, IERROR
17

```

This function returns a predefined MPI datatype that matches an INTEGER variable of KIND selected_int_kind(r). Matching rules for datatypes created by this function are analogous to the matching rules for datatypes created by MPI_TYPE_CREATE_F90_REAL. Restrictions on using the returned datatype with the "external32" data representation are given on page 809.

It is erroneous to supply a value for r that is not supported by the compiler.

Example 19.1. Fortran selected integer and real kind buffers in MPI communications.

```

26 integer      longtype, quadtype
27 integer, parameter :: long = selected_int_kind(15)
28 integer(long) ii(10)
29 real(selected_real_kind(30)) x(10)
30 call MPI_TYPE_CREATE_F90_INTEGER(15, longtype, ierror)
31 call MPI_TYPE_CREATE_F90_REAL(30, MPI_UNDEFINED, quadtype, ierror)
32 ...
33
34 call MPI_SEND(ii, 10, longtype, ...)
35 call MPI_SEND(x, 10, quadtype, ...)
36

```

Advice to users. The datatypes returned by the procedures in Example 19.1 are predefined datatypes. They cannot be freed; they do not need to be committed; they can be used with predefined reduction operations. There are two situations in which they behave differently syntactically, but not semantically, from the MPI named predefined datatypes.

1. MPI_TYPE_GET_ENVELOPE returns special combinators that allow a program to retrieve the values of p and r.
2. Because the datatypes are not named, they cannot be used as compile-time initializers or otherwise accessed before a call to one of the MPI_TYPE_CREATE_F90_XXX routines.

If a variable was declared specifying a nondefault KIND value that was not obtained with `selected_real_kind()` or `selected_int_kind()`, the only way to obtain a matching MPI datatype is to use the size-based mechanism described in the next section. (*End of advice to users.*)

Advice to implementors. An application may often repeat a call to `MPI_TYPE_CREATE_F90_XXX` with the same combination of `(XXX,p,r)`. The application is not allowed to free the returned predefined, unnamed datatype handles. To prevent the creation of a potentially huge amount of handles, a high quality MPI implementation should return the same datatype handle for the same `(REAL/COMPLEX/INTEGER,p,r)` combination. Checking for the combination `(p,r)` in the preceding call to `MPI_TYPE_CREATE_F90_XXX` and using a hash table to find formerly generated handles should limit the overhead of finding a previously generated datatype with same combination of `(XXX,p,r)`. (*End of advice to implementors.*)

Rationale. The `MPI_TYPE_CREATE_F90_REAL/COMPLEX/INTEGER` interface needs as input the original range and precision values to be able to define useful and compiler-independent external (Section 14.5.2) or user-defined (Section 14.5.3) data representations, and in order to be able to perform automatic and efficient data conversions in a heterogeneous environment. (*End of rationale.*)

We now specify how the datatypes described in this section behave when used with the "external32" external data representation described in Section 14.5.2.

The "external32" representation specifies data formats for integer and floating point values. Integer values are represented in two's complement big-endian format. Floating point values are represented by one of three IEEE formats. These are the IEEE "Single," "Double," and "Double Extended" formats, requiring 4, 8, and 16 bytes of storage, respectively. For the IEEE "Double Extended" formats, MPI specifies a Format Width of 16 bytes, with 15 exponent bits, bias = +10383, 112 fraction bits, and an encoding analogous to the "Double" format.

The "external32" representations of the datatypes returned by `MPI_TYPE_CREATE_F90_REAL/COMPLEX/INTEGER` are given by the following rules.

For `MPI_TYPE_CREATE_F90_REAL`:

```

if      (p > 33) or (r > 4931) then external32 representation
                                     is undefined
else if (p > 15) or (r > 307) then external32_size = 16
else if (p > 6) or (r > 37) then external32_size = 8
else                                     external32_size = 4

```

For `MPI_TYPE_CREATE_F90_COMPLEX`: twice the size as for `MPI_TYPE_CREATE_F90_REAL`.

For `MPI_TYPE_CREATE_F90_INTEGER`:

```

if      (r > 38) then external32 representation is undefined
else if (r > 18) then external32_size = 16
else if (r > 9) then external32_size = 8
else if (r > 4) then external32_size = 4
else if (r > 2) then external32_size = 2
else                                     external32_size = 1

```

1 If the "external32" representation of a datatype is undefined, the result of using the datatype
 2 directly or indirectly (i.e., as part of another datatype or through a duplicated datatype)
 3 in operations that require the "external32" representation is undefined. These operations in-
 4 clude MPI_PACK_EXTERNAL, MPI_UNPACK_EXTERNAL, and many MPI_FILE functions,
 5 when the "external32" data representation is used. The ranges for which the "external32"
 6 representation is undefined are reserved for future standardization.

8 *Support for Size-specific MPI Datatypes*

9
 10 MPI provides named datatypes corresponding to optional Fortran 77 numeric types that
 11 contain explicit byte lengths—MPI_REAL4, MPI_INTEGER8, etc. This section describes a
 12 mechanism that generalizes this model to support all Fortran numeric intrinsic types.

13 We assume that for each **typeclass** (integer, real, complex) and each word size **n** there
 14 is a unique machine representation. For every pair (**typeclass**, **n**) supported by a compiler,
 15 MPI must provide a named size-specific datatype. The name of this datatype is of the form
 16 MPI_<TYPECLASS><n> in C and Fortran where <TYPECLASS> is one of REAL, INTEGER,
 17 or COMPLEX, and <n> is the length in bytes of the machine representation. This datatype
 18 locally matches all variables of type (**typeclass**, **n**) in Fortran. The list of names for such
 19 types includes:

20 MPI_REAL4
 21 MPI_REAL8
 22 MPI_REAL16
 23 MPI_COMPLEX8
 24 MPI_COMPLEX16
 25 MPI_COMPLEX32
 26 MPI_INTEGER1
 27 MPI_INTEGER2
 28 MPI_INTEGER4
 29 MPI_INTEGER8
 30 MPI_INTEGER16

31
 32 One datatype is required for each representation supported by the Fortran compiler.

33 *Rationale.* Particularly for the longer floating-point types, C and Fortran may use
 34 different representations. For example, a Fortran compiler may define a 16-byte REAL
 35 type with 33 decimal digits of precision while a C compiler may define a 16-byte long
 36 double type that implements an 80-bit (10 byte) extended precision floating point
 37 value. Both of these types are 16 bytes long, but they are not interoperable. Thus,
 38 these types are defined by Fortran, even though C may define types of the same length.
 39 (*End of rationale.*)

40
 41 To be backward compatible with the interpretation of these types in MPI-1, we as-
 42 sume that the nonstandard declarations REAL*n, INTEGER*n, always create a variable whose
 43 representation is of size **n**. These datatypes may also be used for variables declared with
 44 KIND=INT8/16/32/64 or KIND=REAL32/64/128, which are defined in the ISO_FORTRAN_ENV in-
 45 trinsic module. Note that the MPI datatypes and the REAL*n, INTEGER*n declarations count
 46 bytes whereas the Fortran KIND values count bits. All these datatypes are predefined.

47 The following function allows a user to obtain a size-specific MPI datatype for any
 48 intrinsic Fortran type.

MPI_Type_match_size(typeclass, size, datatype) 1

IN	typeclass	generic type specifier (integer)	2
IN	size	size, in bytes, of representation (integer)	3
OUT	datatype	datatype with correct type, size (handle)	4

C binding 5

```
int MPI_Type_match_size(int typeclass, int size, MPI_Datatype *datatype) 6
```

Fortran 2008 binding 7

```
MPI_Type_match_size(typeclass, size, datatype, ierror) 8
  INTEGER, INTENT(IN) :: typeclass, size 9
  TYPE(MPI_Datatype), INTENT(OUT) :: datatype 10
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror 11
```

Fortran binding 12

```
MPI_TYPE_MATCH_SIZE(TYPECLASS, SIZE, DATATYPE, IERROR) 13
  INTEGER TYPECLASS, SIZE, DATATYPE, IERROR 14
```

typeclass is one of MPI_TYPECLASS_REAL, MPI_TYPECLASS_INTEGER and MPI_TYPECLASS_COMPLEX, corresponding to the desired **typeclass**. The function returns an MPI datatype matching a local variable of type (**typeclass**, **size**).

This function returns a reference (handle) to one of the predefined named datatypes, not a duplicate. This type cannot be freed. MPI_TYPE_MATCH_SIZE can be used to obtain a size-specific type that matches a Fortran numeric intrinsic type by first calling storage_size() in order to compute the variable size in bits, dividing it by eight, and then calling MPI_TYPE_MATCH_SIZE to find a suitable datatype. In C, one can use the C operator sizeof() (which returns the size in bytes) instead of storage_size() (which returns the size in bits). In addition, for variables of default kind the variable's size can be computed by a call to MPI_TYPE_GET_EXTENT, if the typeclass is known. It is erroneous to specify a size not supported by the compiler.

Rationale. This is a convenience function. Without it, it can be tedious to find the correct named type. See note to implementors below. (*End of rationale.*)

Advice to implementors. This function could be implemented as a series of tests.

Example 19.2. Example of an implementation of MPI_Type_match_size. 15

```
int MPI_Type_match_size(int typeclass, int size, MPI_Datatype *rtype) 16
{ 17
  switch(typeclass) { 18
    case MPI_TYPECLASS_REAL: switch(size) { 19
      case 4: *rtype = MPI_REAL4; return MPI_SUCCESS; 20
      case 8: *rtype = MPI_REAL8; return MPI_SUCCESS; 21
      default: error(...); 22
    } 23
    case MPI_TYPECLASS_INTEGER: switch(size) { 24
      case 4: *rtype = MPI_INTEGER4; return MPI_SUCCESS; 25
      case 8: *rtype = MPI_INTEGER8; return MPI_SUCCESS; 26
      default: error(...); 27
    } 28
    ... etc. ... 29
  } 30
} 31
```

```

1      }
2
3      return MPI_SUCCESS;
4  }

```

(End of advice to implementors.)

Communication With Size-specific Types

The usual type matching rules apply to size-specific datatypes: a value sent with datatype `MPI_<TYPECLASS><n>` can be received with this same datatype on another MPI process. Most modern computers use two's complement for integers and IEEE format for floating point. Thus, communication using these size-specific datatypes will not entail loss of precision or truncation errors.

Advice to users. Care is required when communicating in a heterogeneous environment. Consider the following code:

Example 19.3. Unsafe heterogeneous communication due to the use of `MPI_TYPE_MATCH_SIZE`.

```

20  real(selected_real_kind(5)) x(100)
21  size = storage_size(x) / 8
22  call MPI_TYPE_MATCH_SIZE(MPI_TYPECLASS_REAL, size, xtype, ierror)
23  if (myrank .eq. 0) then
24      ... initialize x ...
25      call MPI_SEND(x, xtype, 100, 1, ...)
26  else if (myrank .eq. 1) then
27      call MPI_RECV(x, xtype, 100, 0, ...)
28  endif

```

This may not work in a heterogeneous environment if the value of `size` is not the same on the MPI processes with ranks 0 and 1. There should be no problem in a homogeneous environment. To communicate in a heterogeneous environment, there are at least four options. The first is to declare variables of default type and use the MPI datatypes for these types, e.g., declare a variable of type `REAL` and use `MPI_REAL`. The second is to use `selected_real_kind` or `selected_int_kind` and with the functions of the previous section. The third is to declare a variable that is known to be the same size on all architectures (e.g., `selected_real_kind(12)` on almost all compilers will result in an 8-byte representation). The fourth is to carefully check representation size before communication. This may require explicit conversion to a variable of size that can be communicated and handshaking between sender and receiver to agree on a size.

Note finally that using the "external32" representation for I/O requires explicit attention to the representation sizes. Consider the following code:

Example 19.4. Unsafe heterogeneous MPI file I/O due to the use of `MPI_TYPE_MATCH_SIZE`.

```

47  real(selected_real_kind(5)) x(100)

```

```

size = storage_size(x) / 8
call MPI_TYPE_MATCH_SIZE(MPI_TYPECLASS_REAL, size, xtype, ierror)

if (myrank .eq. 0) then
  call MPI_FILE_OPEN(MPI_COMM_SELF, 'foo',
                    MPI_MODE_CREATE+MPI_MODE_WRONLY,
                    MPI_INFO_NULL, fh, ierror)
  call MPI_FILE_SET_VIEW(fh, zero, xtype, xtype, 'external32', &
                        MPI_INFO_NULL, ierror)
  call MPI_FILE_WRITE(fh, x, 100, xtype, status, ierror)
  call MPI_FILE_CLOSE(fh, ierror)
endif

call MPI_BARRIER(MPI_COMM_WORLD, ierror)

if (myrank .eq. 1) then
  call MPI_FILE_OPEN(MPI_COMM_SELF, 'foo', MPI_MODE_RDONLY,
                    MPI_INFO_NULL, fh, ierror)
  call MPI_FILE_SET_VIEW(fh, zero, xtype, xtype, 'external32', &
                        MPI_INFO_NULL, ierror)
  call MPI_FILE_READ(fh, x, 100, xtype, status, ierror)
  call MPI_FILE_CLOSE(fh, ierror)
endif

```

If the MPI processes with ranks 0 and 1 are on different machines, this code may not work as expected if the `size` is different on the two machines. (*End of advice to users.*)

19.1.10 Problems With Fortran Bindings for MPI

This section discusses a number of problems that may arise when using MPI in a Fortran program. It is intended as advice to users, and clarifies how MPI interacts with Fortran. It is intended to clarify, not add to, this standard.

As noted in the original MPI specification, the interface violates the Fortran standard in several ways. While these may cause few problems for Fortran 77 programs, they become more significant for Fortran 90 programs, so that users must exercise care when using new Fortran 90 features. With Fortran 2008 and the semantics defined in TS 29113, most violations are resolved, and this is hinted at in an addendum to each item. The violations were originally adopted and have been retained because they are important for the usability of MPI. The rest of this section describes the potential problems in detail.

The following MPI features are inconsistent with Fortran 90 and Fortran 77.

1. An MPI subroutine with a choice argument may be called with different argument types. When using the `mpi_f08` module together with a compiler that supports Fortran 2008 with TS 29113, this problem is resolved.
2. An MPI subroutine with an assumed-size dummy argument may be passed an actual scalar argument. This is only solved for choice buffers through the use of `DIMENSION(...)`.
3. Nonblocking and split-collective MPI routines assume that actual arguments are passed by address or descriptor and that arguments and the associated data are not copied

1 on entrance to or exit from the subroutine. This problem is solved with the use of the
2 ASYNCHRONOUS attribute.

- 3
4 4. An MPI implementation may read or modify user data (e.g., communication buffers
5 used by nonblocking communications) concurrently with a user program that is ex-
6 ecuting outside of MPI calls. This problem is resolved by relying on the extended
7 semantics of the ASYNCHRONOUS attribute as specified in TS 29113.
- 8
9 5. Several named “constants,” such as MPI_BOTTOM, MPI_IN_PLACE,
10 MPI_STATUS_IGNORE, MPI_STATUSES_IGNORE, MPI_ERRCODES_IGNORE,
11 MPI_UNWEIGHTED, MPI_WEIGHTS_EMPTY, MPI_ARGV_NULL, and MPI_ARGVS_NULL
12 are not ordinary Fortran constants and require a special implementation. See Sec-
13 tion 2.5.4 for more information.
- 14
15 6. The memory allocation routine MPI_ALLOC_MEM cannot be used from
16 Fortran 77/90/95 without a language extension (for example, Cray pointers) that
17 allows the allocated memory to be associated with a Fortran variable. Therefore,
18 address sized integers were used in MPI-2.0 – MPI-2.2. In Fortran 2003,
19 TYPE(C_PTR) entities were added, which allow a standard-conforming implementation
20 of the semantics of MPI_ALLOC_MEM. In MPI-3.0 and later, MPI_ALLOC_MEM has
21 an additional, overloaded interface to support this language feature. The use of Cray
22 pointers is deprecated. The `mpi_f08` module only supports TYPE(C_PTR) pointers.

23 Additionally, MPI is inconsistent with Fortran 77 in a number of ways, as noted below.

- 24 • MPI identifiers exceed 6 characters.
- 25 • MPI identifiers may contain underscores after the first character.
- 26 • MPI requires an include file, `mpif.h` (deprecated). On systems that do not support
27 include files, the implementation should specify the values of named constants.
- 28 • MPI requires an include file, `mpif.h` (deprecated). On systems that do not support
29 include files, the implementation should specify the values of named constants.
- 30 • Many routines in MPI have KIND-parameterized integers (e.g., MPI_ADDRESS_KIND and
31 MPI_OFFSET_KIND) that hold address information. On systems that do not support
32 Fortran 90-style parameterized types, INTEGER*8 or INTEGER should be used instead.

34 MPI-1 contained several routines that take address-sized information as input or return
35 address-sized information as output. In C such arguments were of type
36 MPI_Aint and in Fortran of type INTEGER. On machines where integers are smaller than
37 addresses, these routines can lose information. In MPI-2 the use of these functions has
38 been deprecated and they have been replaced by routines taking INTEGER arguments of
39 KIND=MPI_ADDRESS_KIND. A number of MPI-2 functions also take INTEGER arguments of
40 nondefault KIND. See Section 2.6 and Section 5.1.1 for more information.

41 Sections 19.1.11 through 19.1.19 describe several problems in detail that concern the
42 interaction of MPI and Fortran as well as their solutions. Some of these solutions require spe-
43 cial capabilities from the compilers. Major requirements are summarized in Section 19.1.7.

45 19.1.11 Problems Due to Strong Typing

46 All MPI functions with choice arguments associate actual arguments of different Fortran
47 datatypes with the same dummy argument. This is not allowed by Fortran 77, and in
48

Fortran 90, it is technically only allowed if the function is overloaded with a different function for each type (see also Section 19.1.6). In C, the use of `void*` formal arguments avoids these problems. Similar to C, with Fortran 2008 with TS 29113 (and later) together with the `mpi_f08` module, the problem is avoided by declaring choice arguments with `TYPE(*)`, `DIMENSION(..)`, i.e., as assumed-type and assumed-rank dummy arguments.

Using `INCLUDE 'mpif.h'` (deprecated), the following code fragment is technically invalid and may generate a compile-time error.

```
integer i(5)
real    x(5)
...
call mpi_send(x, 5, MPI_REAL, ...)
call mpi_send(i, 5, MPI_INTEGER, ...)
```

In practice, it is rare for compilers to do more than issue a warning. When using either the `mpi_f08` or `mpi` module, the problem is usually resolved through the assumed-type and assumed-rank declarations of the dummy arguments, or with a compiler-dependent mechanism that overrides type checking for choice arguments.

It is also technically invalid in Fortran to pass a scalar actual argument to an array dummy argument that is not a choice buffer argument. Thus, when using the `mpi_f08` or `mpi` module, the following code fragment usually generates an error since the `dims` and `periods` arguments to `MPI_CART_CREATE` are declared as assumed size arrays `INTEGER :: DIMS(*)` and `LOGICAL :: PERIODS(*)`.

Example 19.5. It is erroneous to pass a variable instead of an array with one element.

```
! ----- THIS EXAMPLE IS ERRONEOUS -----
USE mpi_f08      ! or USE mpi
INTEGER size
CALL MPI_Cart_create(comm_old, 1, size, .TRUE., .TRUE., comm_cart, ierror)
```

Although this is a nonconforming MPI call, compiler warnings are not expected (but may occur) when using `INCLUDE 'mpif.h'` (deprecated) and this include file does not use Fortran explicit interfaces.

19.1.12 Problems Due to Data Copying and Sequence Association with Subscript Triplets

Arrays with subscript **triplets** describe Fortran subarrays with or without strides, e.g.,

Example 19.6. Fortran subarrays as actual buffer in MPI procedures.

```
REAL a(100,100,100)
CALL MPI_Send(a(11:17, 12:99:3, 1:100), 7*30*100, MPI_REAL, ...)
```

The handling of subscript triplets depends on the value of the constant `MPI_SUBARRAYS_SUPPORTED`:

- If `MPI_SUBARRAYS_SUPPORTED` equals `.TRUE.:`

Choice buffer arguments are declared as `TYPE(*)`, `DIMENSION(..)`. For example, consider the following code fragment:

Example 19.7. Fortran subarrays without restrictions if `MPI_SUBARRAYS_SUPPORTED` equals `.TRUE.`.

```

REAL s(100), r(100)
CALL MPI_Isend(s(1:100:5), 3, MPI_REAL, ..., rq, ierror)
CALL MPI_Wait(rq, status, ierror)
CALL MPI_Irecv(r(1:100:5), 3, MPI_REAL, ..., rq, ierror)
CALL MPI_Wait(rq, status, ierror)

```

In this case, the individual elements `s(1)`, `s(6)`, and `s(11)` are sent between the start of `MPI_ISEND` and the end of `MPI_WAIT` even though the compiled code will not copy `s(1:100:5)` to a real contiguous temporary scratch buffer. Instead, the compiled code will pass a descriptor to `MPI_ISEND` that allows MPI to operate directly on `s(1)`, `s(6)`, `s(11)`, ..., `s(96)`. The called `MPI_ISEND` routine will take only the first three of these elements due to the type signature “3, `MPI_REAL`”.

All nonblocking MPI functions (e.g., `MPI_ISEND`, `MPI_PUT`, `MPI_FILE_WRITE_ALL_BEGIN`) behave as if *the user-specified elements of choice buffers are copied to a contiguous scratch buffer in the MPI runtime environment*. All datatype descriptions (in the example above, “3, `MPI_REAL`”) read and store data from and to this virtual contiguous scratch buffer. Displacements in MPI derived datatypes are relative to the beginning of this virtual contiguous scratch buffer. Upon completion of a nonblocking receive operation (e.g., when `MPI_WAIT` on a corresponding `MPI_Request` returns), it is as if the received data has been copied from the virtual contiguous scratch buffer back to the noncontiguous application buffer. In the example above, `r(1)`, `r(6)`, and `r(11)` are guaranteed to be defined with the received data when `MPI_WAIT` returns.

Note that the above definition does not supercede restrictions about buffers used with nonblocking operations (e.g., those specified in Section 3.7.2).

Advice to implementors. The Fortran descriptor for `TYPE(*)`, `DIMENSION(..)` arguments contains enough information that, if desired, the MPI library can make a real contiguous copy of noncontiguous user buffers when the nonblocking operation is started, and release this buffer not before the nonblocking communication has completed (e.g., the `MPI_WAIT` routine). Efficient implementations may avoid such additional memory-to-memory data copying. (*End of advice to implementors.*)

Rationale. If `MPI_SUBARRAYS_SUPPORTED` equals `.TRUE.`, non-contiguous buffers are handled inside the MPI library instead of by the compiler through argument association conventions. Therefore, the scope of MPI library scratch buffers can be from the beginning of a nonblocking operation until the completion of the operation although beginning and completion are implemented in different routines. (*End of rationale.*)

- If `MPI_SUBARRAYS_SUPPORTED` equals `.FALSE.`:

In this case, the use of Fortran arrays with subscript triplets as actual choice buffer arguments in any nonblocking MPI operation (which also includes persistent request, and split collectives) may cause undefined behavior. They may, however, be used in blocking MPI operations.

Implicit in MPI is the idea of a contiguous chunk of memory accessible through a linear address space. MPI copies data to and from this memory. An MPI program specifies the location of data by providing memory addresses and offsets. In the C language, sequence association rules plus pointers provide all the necessary low-level structure.

In Fortran, array data is not necessarily stored contiguously. For example, the array section `A(1:N:2)` involves only the elements of `A` with indices 1, 3, 5, The same is true for a pointer array whose target is such a section. Most compilers ensure that an array that is a dummy argument is held in contiguous memory if it is declared with an explicit shape (e.g., `B(N)`) or is of assumed size (e.g., `B(*)`). If necessary, they do this by making a copy of the array into contiguous memory.¹

Because MPI dummy buffer arguments are assumed-size arrays if `MPI_SUBARRAYS_SUPPORTED` equals `.FALSE.`, this leads to a serious problem for a nonblocking call: the compiler copies the temporary array back on return but MPI continues to copy data to the memory that held it. For example, consider the following code fragment:

Example 19.8. Fortran subarrays cannot be used if `MPI_SUBARRAYS_SUPPORTED` equals `.FALSE.`.

```
!-- THIS EXAMPLE IS ERRONEOUS if MPI_SUBARRAYS_SUPPORTED==.FALSE. --
real a(100)
call MPI_IRECV(a(1:100:2), MPI_REAL, 50, ...)
```

Since the first dummy argument to `MPI_IRECV` is an assumed-size array (`<type> buf(*)`), the array section `a(1:100:2)` is copied to a temporary before being passed to `MPI_IRECV`, so that it is contiguous in memory. `MPI_IRECV` returns immediately, and data is copied from the temporary back into the array `a`. Sometime later, MPI may write to the address of the deallocated temporary. Copying is also a problem for `MPI_ISEND` since the temporary array may be deallocated before the data has all been sent from it.

Most Fortran 90 compilers do not make a copy if the actual argument is the whole of an explicit-shape or assumed-size array or is a “simply contiguous” section such as `A(1:N)` of such an array. (“Simply contiguous” is defined in the next paragraph.) Also, many compilers treat allocatable arrays the same as they treat explicit-shape arrays in this regard (though we know of one that does not). However, the same is not true for assumed-shape and pointer arrays; since they may be discontinuous, copying is often done. It is this copying that causes problems for MPI as described in the previous paragraph.

According to the Fortran 2008 Standard, Section 6.5.4, a “simply contiguous” array section is

```
name ( [:,]... [<subscript>]:<subscript> [,<subscript>]... )
```

That is, there are zero or more dimensions that are selected in full, then one dimension selected without a stride, then zero or more dimensions that are selected with a simple

¹Technically, the Fortran standard is worded to allow noncontiguous storage of any array data, unless the dummy argument has the `CONTIGUOUS` attribute.

subscript. The compiler can detect from analyzing the source code that the array is contiguous. Examples are

```
A(1:N), A(:,N), A(:,1:N,1), A(1:6,N), A(:, :, 1:N)
```

Because of Fortran’s column-major ordering, where the first index varies fastest, a “simply contiguous” section of a contiguous array will also be contiguous.

The same problem can occur with a scalar argument. A compiler may make a copy of scalar dummy arguments within a called procedure when passed as an actual argument to a choice buffer routine. That this can cause a problem is illustrated by the example

Example 19.9. Problem with scalar arguments.

```
real :: a
call user1(a,rq)
call MPI_WAIT(rq,status,ierr)
write (*,*) a

subroutine user1(buf,request)
call MPI_IRECV(buf,...,request,...)
end
```

If `a` is copied, `MPI_IRECV` will alter the copy when it completes the communication and will not alter `a` itself.

Note that copying will almost certainly occur for an argument that is a nontrivial expression (one with at least one operator or function call), a section that does not select a contiguous part of its parent (e.g., `A(1:n:2)`), a pointer whose target is such a section, or an assumed-shape array that is (directly or indirectly) associated with such a section.

If a compiler option exists that inhibits copying of arguments, in either the calling or called procedure, this must be employed.

If a compiler makes copies in the calling procedure of arguments that are explicit-shape or assumed-size arrays, “simply contiguous” array sections of such arrays, or scalars, and if no compiler option exists to inhibit such copying, then the compiler cannot be used for applications that use `MPI_GET_ADDRESS`, or any nonblocking MPI routine. If a compiler copies scalar arguments in the called procedure and there is no compiler option to inhibit this, then this compiler cannot be used for applications that use memory references across subroutine calls as in the example above.

19.1.13 Problems Due to Data Copying and Sequence Association with Vector Subscripts

Fortran arrays with **vector** subscripts describe subarrays containing a possibly irregular set of elements

Example 19.10. Fortran irregular subarrays through using vector subscripts.

```
REAL a(100)
CALL MPI_Send(A((/7,9,23,81,82/)), 5, MPI_REAL, ...)
```

Fortran arrays with a vector subscript must not be used as actual choice buffer arguments in any nonblocking or split collective MPI operations. They may, however, be used in blocking MPI operations.

19.1.14 Special Constants

MPI requires a number of special “constants” that cannot be implemented as normal Fortran constants, e.g., `MPI_BOTTOM`. The complete list can be found in Section 2.5.4. In C, these are implemented as constant pointers, usually as `NULL` and are used where the function prototype calls for a pointer to a variable, not the variable itself.

In Fortran, using special values for the constants (e.g., by defining them through parameter statements) is not possible because an implementation cannot distinguish these values from valid data. Typically these constants are implemented as predefined static variables (e.g., a variable in an MPI-declared `COMMON` block), relying on the fact that the target compiler passes data by address. Inside the subroutine, the address of the actual choice buffer argument can be compared with the address of such a predefined static variable.

These special constants also cause an exception with the usage of Fortran `INTENT`: with `USE mpi_f08`, the attributes `INTENT(IN)`, `INTENT(OUT)`, and `INTENT(INOUT)` are used in the Fortran interface. In most cases, `INTENT(IN)` is used if the C interface uses call-by-value. For all buffer arguments and for dummy arguments that may be modified and allow one of these special constants as input, an `INTENT` is not specified.

19.1.15 Fortran Derived Types

MPI supports passing Fortran entities of `BIND(C)` and `SEQUENCE` derived types to choice dummy arguments, provided no type component has the `ALLOCATABLE` or `POINTER` attribute.

The following code fragment shows some possible ways to send scalars or arrays of interoperable derived types in Fortran. The example assumes that all data is passed by address.

Example 19.11. Fortran array of derived Fortran types: the struct MPI derived type should be resized.

```

type, BIND(C) :: mytype
  integer :: i
  real :: x
  double precision :: d
  logical :: l
end type mytype

type(mytype) :: foo, fooarr(5)
integer :: blocklen(4), dtype(4)
integer(KIND=MPI_ADDRESS_KIND) :: disp(4), base, lb, extent

call MPI_GET_ADDRESS(foo%i, disp(1), ierr)
call MPI_GET_ADDRESS(foo%x, disp(2), ierr)
call MPI_GET_ADDRESS(foo%d, disp(3), ierr)
call MPI_GET_ADDRESS(foo%l, disp(4), ierr)

base = disp(1)
disp(1) = disp(1) - base

```

```

1  disp(2) = disp(2) - base
2  disp(3) = disp(3) - base
3  disp(4) = disp(4) - base
4
5  blocklen(1) = 1
6  blocklen(2) = 1
7  blocklen(3) = 1
8  blocklen(4) = 1
9
10 dtype(1) = MPI_INTEGER
11 dtype(2) = MPI_REAL
12 dtype(3) = MPI_DOUBLE_PRECISION
13 dtype(4) = MPI_LOGICAL
14
15 call MPI_TYPE_CREATE_STRUCT(4, blocklen, disp, dtype, newtype, ierr)
16 call MPI_TYPE_COMMIT(newtype, ierr)
17
18 call MPI_SEND(foo%i, 1, newtype, dest, tag, comm, ierr)
19 ! or
20 call MPI_SEND(foo, 1, newtype, dest, tag, comm, ierr)
21 ! expects that base == address(foo%i) == address(foo)
22
23 call MPI_GET_ADDRESS(fooarr(1), disp(1), ierr)
24 call MPI_GET_ADDRESS(fooarr(2), disp(2), ierr)
25 extent = disp(2) - disp(1)
26 lb = 0
27 call MPI_TYPE_CREATE_RESIZED(newtype, lb, extent, newarrtype, ierr)
28 call MPI_TYPE_COMMIT(newarrtype, ierr)
29
30 call MPI_SEND(fooarr, 5, newarrtype, dest, tag, comm, ierr)

```

Using the derived type variable `foo` instead of its first basic type element `foo%i` may be impossible if the MPI library implements choice buffer arguments through overloading instead of using `TYPE(*)`, `DIMENSION(..)`, or through a nonstandardized extension such as `!$PRAGMA IGNORE_TKR`; see Section 19.1.6.

To use a derived type in an array requires a correct extent of the datatype handle to take care of the alignment rules applied by the compiler. These alignment rules may imply that there are gaps between the components of a derived type, and also between the subsequent elements of an array of a derived type. The extent of an interoperable derived type (i.e., defined with `BIND(C)`) and a `SEQUENCE` derived type with the same content may be different because C and Fortran may apply different alignment rules. As recommended in the advice to users in Section 5.1.6, one should add an additional fifth structure element with one numerical storage unit at the end of this structure to force in most cases that the array of structures is contiguous. Even with such an additional element, one should keep this resizing due to the special alignment rules that can be used by the compiler for structures, as also mentioned in this advice.

Using the extended semantics defined in TS 29113, it is also possible to use entities or derived types without either the `BIND(C)` or the `SEQUENCE` attribute as choice buffer arguments; some additional constraints must be observed, e.g., no `ALLOCATABLE` or `POINTER` type components may exist. In this case, the base address in the example must be changed to become the address of `foo` instead of `foo%i`, because the Fortran compiler may rearrange

Table 19.2: Occurrence of Fortran optimization problems in several usage areas

Optimization may cause a problem in following usage areas			
	Nonbl.	1-sided	Split	Bottom
Code movement and register optimization	yes	yes	no	yes
Temporary data movement	yes	yes	yes	no
Permanent data movement	yes	yes	yes	yes

type components or add padding. Sending the structure `foo` should then also be performed by providing it (and not `foo%i`) as actual argument for `MPI_Send`.

19.1.16 Optimization Problems, an Overview

MPI provides operations that may be hidden from the user code and run concurrently with it, accessing the same memory as user code. Examples include the data transfer for an `MPI_IRECV`. The optimizer of a compiler will assume that it can recognize periods when a copy of a variable can be kept in a register without reloading from or storing to memory. When the user code is working with a register copy of some variable while the hidden operation reads or writes the memory copy, problems occur. These problems are independent of the Fortran support method; i.e., they occur with the `mpi_f08` module, the `mpi` module, and the (deprecated) `mpif.h` include file.

This section shows four problematic usage areas (the abbreviations in parentheses are used in the table below):

- Use of nonblocking routines or persistent requests (*Nonbl.*).
- Use of one-sided routines (*1-sided*).
- Use of MPI parallel file I/O split collective operations (*Split*).
- Use of `MPI_BOTTOM` together with absolute displacements in MPI datatypes, or relative displacements between two variables in such datatypes (*Bottom*).

The following compiler optimization strategies (valid for serial code) may cause problems in MPI applications:

- Code movement and register optimization problems; see Section 19.1.17.
- Temporary data movement and temporary memory modifications; see Section 19.1.18.
- Permanent data movement (e.g., through garbage collection); see Section 19.1.19.

Table 19.2 shows the only usage areas where these optimization problems may occur.

The solutions in the following sections are based on compromises:

- to minimize the burden for the application programmer, e.g., as shown in Sections Solutions through The (Poorly Performing) Fortran VOLATILE Attribute on pages 824–828,

- to minimize the drawbacks on compiler based optimization, and
- to minimize the requirements defined in Section 19.1.7.

19.1.17 Problems with Code Movement and Register Optimization

Nonblocking Operations

If a variable is local to a Fortran subroutine (i.e., not in a module or a COMMON block), the compiler will assume that it cannot be modified by a called subroutine unless it is an actual argument of the call. In the most common linkage convention, the subroutine is expected to save and restore certain registers. Thus, the optimizer will assume that a register that held a valid copy of such a variable before the call will still hold a valid copy on return.

Example 19.12. Fortran 90 register optimization—extreme.

Source	compiled as	or compiled as
REAL :: buf, b1	REAL :: buf, b1	REAL :: buf, b1
call MPI_IRecv(buf,..req)	call MPI_IRecv(buf,..req)	call MPI_IRecv(buf,..req)
	register = buf	b1 = buf
call MPI_WAIT(req,..)	call MPI_WAIT(req,..)	call MPI_WAIT(req,..)
b1 = buf	b1 = register	

Example 19.12 shows extreme, but allowed, possibilities. MPI_WAIT on a concurrent thread modifies buf between the invocation of MPI_IRecv and the completion of MPI_WAIT. But the compiler cannot see any possibility that buf can be changed after MPI_IRecv has returned, and may schedule the load of buf earlier than typed in the source. The compiler has no reason to avoid using a register to hold buf across the call to MPI_WAIT. It also may reorder the instructions as illustrated in the rightmost column.

Example 19.13. Similar example with MPI_ISEND

Source	compiled as	with a possible MPI-internal execution sequence
REAL :: buf, copy	REAL :: buf, copy	REAL :: buf, copy
buf = val	buf = val	buf = val
call MPI_ISEND(buf,..req)	call MPI_ISEND(buf,..req)	addr = &buf
copy = buf	copy = buf	copy = buf
	buf = val_overwrite	buf = val_overwrite
call MPI_WAIT(req,..)	call MPI_WAIT(req,..)	call send(*addr) ! within
		! MPI_WAIT
buf = val_overwrite		

Due to valid compiler code movement optimizations in Example 19.13, the content of buf may already have been overwritten by the compiler when the content of buf is sent. The code movement is permitted because the compiler cannot detect a possible access to buf in MPI_WAIT (or in a second thread between the start of MPI_ISEND and the end of MPI_WAIT).

Such register optimization is based on moving code; here, the access to buf was moved from after MPI_WAIT to before MPI_WAIT. Note that code movement may also occur across subroutine boundaries when subroutines or functions are inlined.

This register optimization/code movement problem for nonblocking operations does not occur with MPI parallel file I/O split collective operations, because in the MPI_XXX_BEGIN and MPI_XXX_END calls, the same buffer has to be provided as an actual

argument. The register optimization / code movement problem for MPI_BOTTOM and derived MPI datatypes may occur in each blocking and nonblocking communication call, as well as in each parallel file I/O operation.

Persistent Operations

With persistent requests, the buffer argument is hidden from the MPI_START and MPI_STARTALL calls, i.e., the Fortran compiler may move buffer accesses across the MPI_START or MPI_STARTALL call, similar to the MPI_WAIT call as described in the Nonblocking Operations subsection in Section 19.1.17.

One-sided Communication

An example with instruction reordering due to register optimization can be found in Section 12.7.4.

MPI_BOTTOM and Combining Independent Variables in Datatypes

This section is only relevant if the MPI program uses a buffer argument to an MPI_SEND, MPI_RECV, etc., that hides the actual variables involved in the communication. MPI_BOTTOM with an MPI_Datatype containing *absolute addresses* is one example. Creating a datatype that uses one variable as an anchor and brings along others by using MPI_GET_ADDRESS to determine their offsets from the anchor is another. The anchor variable would be the only one referenced in the call. Also attention must be paid if MPI operations are used that run in parallel with the user's application.

Example 19.14 shows what Fortran compilers are allowed to do.

Example 19.14. Fortran 90 register optimization.

This source	can be compiled as
<code>call MPI_GET_ADDRESS(buf, bufaddr, ierror)</code>	<code>call MPI_GET_ADDRESS(buf, ...)</code>
<code>call MPI_TYPE_CREATE_STRUCT(1, 1, bufaddr, MPI_REAL, dtype, ierror)</code>	<code>call MPI_TYPE_CREATE_STRUCT(...)</code>
<code>call MPI_TYPE_COMMIT(dtype, ierror)</code> <code>val_old = buf</code>	<code>call MPI_TYPE_COMMIT(...)</code> <code>register = buf</code> <code>val_old = register</code>
<code>call MPI_RECV(MPI_BOTTOM, 1, dtype, ...)</code> <code>val_new = buf</code>	<code>call MPI_RECV(MPI_BOTTOM, ...)</code> <code>val_new = register</code>

In Example 19.14, the compiler does not invalidate the register because it cannot see that MPI_RECV changes the value of buf. The access to buf is hidden by the use of MPI_GET_ADDRESS and MPI_BOTTOM.

Example 19.15. Similar example with MPI_SEND

This source	can be compiled as
<code>! buf contains val_old</code> <code>buf = val_new</code>	<code>! buf contains val_old</code>
<code>call MPI_SEND(MPI_BOTTOM, 1, dtype, ...)</code> <code>! with buf as a displacement in dtype</code>	<code>call MPI_SEND(...)</code> <code>! i.e., val_old is sent</code>

```

1           !
2           ! buf=val_new  is moved to here
3           ! and detected as dead code
4           ! and therefore removed
5           !
6 buf = val_overwrite           buf = val_overwrite

```

In Example 19.15, several successive assignments to the same variable `buf` can be combined in a way such that only the last assignment is executed. “Successive” means that no interfering load access to this variable occurs between the assignments. The compiler cannot detect that the call to `MPI_SEND` statement is interfering because the load access to `buf` is hidden by the usage of `MPI_BOTTOM`.

Solutions

The following sections show in detail how the problems with code movement and register optimization can be portably solved. Application writers can partially or fully avoid these compiler optimization problems by using one or more of the special Fortran declarations with the send and receive buffers used in nonblocking operations, or in operations in which `MPI_BOTTOM` is used, or if datatype handles that combine several variables are used:

- Use of the Fortran `ASYNCHRONOUS` attribute.
- Use of the helper routine `MPI_F_SYNC_REG`, or an equivalent user-written dummy routine.
- Declare the buffer as a Fortran module variable or within a Fortran common block.
- Use of the Fortran `VOLATILE` attribute.

Example 19.16. Protecting nonblocking communication with the `ASYNCHRONOUS` attribute.

```

31 USE mpi_f08
32 REAL, ASYNCHRONOUS :: b(0:101) ! elements 0 and 101 are halo cells
33 REAL :: bnew(0:101)           ! elements 1 and 100 are newly computed
34 TYPE(MPI_Request) :: req(4)
35 INTEGER :: left, right, i
36 CALL MPI_Cart_shift(...,left,right,...)
37 CALL MPI_Irecv(b( 0), ..., left, ..., req(1), ...)
38 CALL MPI_Irecv(b(101), ..., right, ..., req(2), ...)
39 CALL MPI_Isend(b( 1), ..., left, ..., req(3), ...)
40 CALL MPI_Isend(b(100), ..., right, ..., req(4), ...)
41
42 #ifdef WITHOUT_OVERLAPPING_COMMUNICATION_AND_COMPUTATION
43 ! Case (a)
44 CALL MPI_Waitall(4, req, ...)
45 DO i=1,100 ! compute all new local data
46   bnew(i) = function(b(i-1), b(i), b(i+1))
47 END DO
48 #endif
49
50 #ifdef WITH_OVERLAPPING_COMMUNICATION_AND_COMPUTATION

```

```

1  ! Case (b)
2  DO i=2,99 ! compute only elements for which halo data is not needed
3      bnew(i) = function(b(i-1), b(i), b(i+1))
4  END DO
5  CALL MPI_Waitall(4, req, ...)
6  i=1 ! compute leftmost element
7      bnew(i) = function(b(i-1), b(i), b(i+1))
8  i=100 ! compute rightmost element
9      bnew(i) = function(b(i-1), b(i), b(i+1))
10 #endif

```

Each of these methods solves the problems of code movement and register optimization, but may incur various degrees of performance impact, and may not be usable in every application context. These methods may not be guaranteed by the Fortran standard, but they must be guaranteed by a MPI-3.0 (and later) compliant MPI library and associated compiler suite according to the requirements listed in Section 19.1.7. The performance impact of using `MPI_F_SYNC_REG` is expected to be low, that of using module variables or the `ASYNCHRONOUS` attribute is expected to be low to medium, and that of using the `VOLATILE` attribute is expected to be high or very high. Note that there is one attribute that cannot be used for this purpose: the Fortran `TARGET` attribute does not solve code movement problems in MPI applications.

The Fortran ASYNCHRONOUS Attribute

Declaring an actual buffer argument with the `ASYNCHRONOUS` Fortran attribute in a scoping unit (or `BLOCK`) informs the compiler that any statement in the scoping unit may be executed while the buffer is affected by a pending asynchronous Fortran input/output operation (since Fortran 2003) or by an asynchronous communication (TS 29113 extension). Without the extensions specified in TS 29113, a Fortran compiler may totally ignore this attribute if the Fortran compiler implements asynchronous Fortran input/output operations with blocking I/O. The `ASYNCHRONOUS` attribute protects the buffer accesses from optimizations through code movements across routine calls, and the buffer itself from temporary and permanent data movements. If the choice buffer dummy argument of a nonblocking MPI routine is declared with `ASYNCHRONOUS` (which is mandatory for the `mpi_f08` module, with allowable exceptions listed in Section 19.1.6), then the compiler has to guarantee call by reference and should report a compile-time error if call by reference is impossible, e.g., if vector subscripts are used. The `MPI_ASYNC_PROTECTS_NONBLOCKING` is set to `.TRUE.` if both the protection of the actual buffer argument through `ASYNCHRONOUS` according to the TS 29113 extension and the declaration of the dummy argument with `ASYNCHRONOUS` in the Fortran support method is guaranteed for all nonblocking routines, otherwise it is set to `.FALSE.`

The `ASYNCHRONOUS` attribute has some restrictions. Section 5.4.2 of the TS 29113 specifies:

“Asynchronous communication for a Fortran variable occurs through the action of procedures defined by means other than Fortran. It is initiated by execution of an asynchronous communication initiation procedure and completed by execution of an asynchronous communication completion procedure. Between the execution of the initiation and completion procedures, any variable of which any part is associated with any part of the asynchronous communication variable is a

pending communication affector. Whether a procedure is an asynchronous communication initiation or completion procedure is processor dependent.

Asynchronous communication is either input communication or output communication. For input communication, a *pending communication affector* shall not be referenced, become defined, become undefined, become associated with a dummy argument that has the VALUE attribute, or have its pointer association status changed. For output communication, a *pending communication affector* shall not be redefined, become undefined, or have its pointer association status changed.”

In Example 19.16 Case (a) on page 824, the read accesses to `b` within `function(b(i-1), b(i), b(i+1))` cannot be moved by compiler optimizations to before the wait call because `b` was declared as `ASYNCHRONOUS`. Note that only the elements 0, 1, 100, and 101 of `b` are involved in asynchronous communication but by definition, the total variable `b` is the *pending communication affector* and is usable for input and output asynchronous communication between the `MPI_IXXX` routines and `MPI_Waitall`. Case (a) works fine because the read accesses to `b` occur after the communication has completed.

In Case (b), the read accesses to `b(1:100)` in the loop `i=2,99` are read accesses to a *pending communication affector* while input communication (i.e., the two `MPI_Irecv` calls) is *pending*. This is a contradiction to the rule that *for input communication, a pending communication affector shall not be referenced*. The problem can be solved by using separate variables for the halos and the inner array, or by splitting a common array into disjoint subarrays that are passed through different dummy arguments into a subroutine, as shown in Example 19.21.

If one does not overlap communication and computation on the same variable, then all optimization problems can be solved through the `ASYNCHRONOUS` attribute.

The problems with `MPI_BOTTOM`, as shown in Example 19.14 and Example 19.15, can also be solved by declaring the buffer `buf` with the `ASYNCHRONOUS` attribute.

In some MPI routines, a buffer dummy argument is defined as `ASYNCHRONOUS` to guarantee passing by reference, provided that the actual argument is also defined as `ASYNCHRONOUS`.

Calling `MPI_F_SYNC_REG`

The compiler may be prevented from moving a reference to a buffer across a call to an MPI subroutine by surrounding the call by calls to an external subroutine with the buffer as an actual argument. The MPI library provides the `MPI_F_SYNC_REG` routine for this purpose; see Section 19.1.8.

- The problems illustrated by the Examples 19.12 and 19.13 can be solved by calling `MPI_F_SYNC_REG(buf)` once immediately after `MPI_WAIT`.

Example 19.12 can be solved with

```
call MPI_IRecv(buf, .. req)

call MPI_WAIT(req, ..)
call MPI_F_SYNC_REG(buf)
b1 = buf
```

Example 19.13 can be solved with

```
buf = val
call MPI_ISEND(buf, .. req)
copy = buf
call MPI_WAIT(req, ..)
call MPI_F_SYNC_REG(buf)
buf = val_overwrite
```

The call to `MPI_F_SYNC_REG(buf)` prevents moving the last line before the `MPI_WAIT` call. Further calls to `MPI_F_SYNC_REG(buf)` are not needed because it is still correct if the additional read access `copy=buf` is moved below `MPI_WAIT` and before `buf=val_overwrite`.

- The problems illustrated by the Examples 19.14 and 19.15 can be solved with two additional `MPI_F_SYNC_REG(buf)` statements; one directly before `MPI_RECV/MPI_SEND`, and one directly after this communication operation.

Example 19.14 can be solved with

```
call MPI_F_SYNC_REG(buf)
call MPI_RECV(MPI_BOTTOM, ...)
call MPI_F_SYNC_REG(buf)
```

Example 19.15 can be solved with

```
call MPI_F_SYNC_REG(buf)
call MPI_SEND(MPI_BOTTOM, ...)
call MPI_F_SYNC_REG(buf)
```

The first call to `MPI_F_SYNC_REG(buf)` is needed to finish all load and store references to `buf` prior to `MPI_RECV/MPI_SEND`; the second call is needed to assure that any subsequent access to `buf` is not moved before `MPI_RECV/MPI_SEND`.

- In the Example 12.14 in Section 12.7.4, two asynchronous accesses must be protected: in Process 1, the access to `bbbb` must be protected similar to Example 19.12, i.e., a call to `MPI_F_SYNC_REG(bbbb)` is needed after the second `MPI_WIN_FENCE` to guarantee that further accesses to `bbbb` are not moved ahead of the call to `MPI_WIN_FENCE`. In Process 2, both calls to `MPI_WIN_FENCE` together act as a communication call with `MPI_BOTTOM` as the buffer. That is, before the first fence and after the second fence, a call to `MPI_F_SYNC_REG(buff)` is needed to guarantee that accesses to `buff` are not moved after or ahead of the calls to `MPI_WIN_FENCE`. Using `MPI_GET` instead of `MPI_PUT`, the same calls to `MPI_F_SYNC_REG` are necessary.

Example 19.17. Solution for the Fortran register optimization problems with one-sided communication in Example 12.14.

Source of Process 1

```
bbbb = 777

call MPI_WIN_FENCE
call MPI_PUT(bbbb
into buff of process 2)

call MPI_WIN_FENCE
call MPI_F_SYNC_REG(bbbb)
```

Source of Process 2

```
buff = 999
call MPI_F_SYNC_REG(buff)
call MPI_WIN_FENCE

call MPI_WIN_FENCE
call MPI_F_SYNC_REG(buff)
ccc = buff
```

- The temporary memory modification problem, i.e., Example 19.18, can **not** be solved with this method.

A User Defined Routine Instead of MPI_F_SYNC_REG

Instead of `MPI_F_SYNC_REG`, one can also use a user defined external subroutine, which is separately compiled:

```

1  subroutine DD(buf)
2      integer buf
3  end

```

Note that if the `INTENT` is declared in an explicit interface for the external subroutine, it must be `OUT` or `INOUT`. The subroutine itself may have an empty body, but the compiler does not know this and has to assume that the buffer may be altered. For example, a call to `MPI_RECV` with `MPI_BOTTOM` as buffer might be replaced by

```

8  call DD(buf)
9  call MPI_RECV(MPI_BOTTOM, ...)
10 call DD(buf)

```

Such a user-defined routine was introduced in MPI-2.0 and is still included here to document such usage in existing application programs although new applications should prefer `MPI_F_SYNC_REG` or one of the other possibilities. In an existing application, calls to such a user-written routine should be substituted by a call to `MPI_F_SYNC_REG` because the user-written routine may not be implemented in accordance with the rules specified in Section 19.1.7.

Module Variables and COMMON Blocks

An alternative to the previously mentioned methods is to put the buffer or variable into a module or a common block and access it through a `USE` or `COMMON` statement in each scope where it is referenced, defined or appears as an actual argument in a call to an MPI routine. The compiler will then have to assume that the MPI procedure may alter the buffer or variable, provided that the compiler cannot infer that the MPI procedure does not reference the module or common block.

- This method solves problems of instruction reordering, code movement, and register optimization related to nonblocking and one-sided communication, or related to the usage of `MPI_BOTTOM` and derived datatype handles.
- Unfortunately, this method does **not** solve problems caused by asynchronous accesses between the start and end of a nonblocking or one-sided communication. Specifically, problems caused by temporary memory modifications are not solved.

The (Poorly Performing) Fortran VOLATILE Attribute

The `VOLATILE` attribute gives the buffer or variable the properties needed to avoid register optimization or code movement problems, but it may inhibit optimization of any code containing references or definitions of the buffer or variable. On many modern systems, the performance impact will be large because not only register, but also cache optimizations will not be applied. Therefore, use of the `VOLATILE` attribute to enforce correct execution of MPI programs is discouraged.

The Fortran TARGET Attribute

The `TARGET` attribute does not solve the code movement problem because it is not specified for the choice buffer dummy arguments of nonblocking routines. If the compiler detects that the application program specifies the `TARGET` attribute for an actual buffer argument used

in the call to a nonblocking routine, the compiler may ignore this attribute if no pointer reference to this buffer exists.

Rationale. The Fortran standardization body decided to extend the ASYNCHRONOUS attribute within the TS 29113 to protect buffers in nonblocking calls from all kinds of optimization, instead of extending the TARGET attribute. (*End of rationale.*)

19.1.18 Temporary Data Movement and Temporary Memory Modification

The compiler is allowed to temporarily modify data in memory. Normally, this problem may occur only when overlapping communication and computation, as in Example 19.16, Case (b) on page 824. Example 19.18 also shows a possibility that could be problematic.

Example 19.18. Overlapping Communication and Computation.

```
USE mpi_f08
REAL :: buf(100,100)
CALL MPI_Irecv(buf(1,1:100),..., req,...)
DO j=1,100
  DO i=2,100
    buf(i,j)=...
  END DO
END DO
CALL MPI_Wait(req,...)
```

Example 19.19. The compiler may substitute the nested loops through loop fusion.

```
REAL :: buf(100,100), buf_1dim(10000)
EQUIVALENCE (buf(1,1), buf_1dim(1))
CALL MPI_Irecv(buf(1,1:100),..., req,...)
tmp(1:100) = buf(1,1:100)
DO j=1,10000
  buf_1dim(h)=...
END DO
buf(1,1:100) = tmp(1:100)
CALL MPI_Wait(req,...)
```

Example 19.20. Another optimization is based on the usage of a separate memory storage area, e.g., in a GPU.

```
REAL :: buf(100,100), local_buf(100,100)
CALL MPI_Irecv(buf(1,1:100),..., req,...)
local_buf = buf
DO j=1,100
  DO i=2,100
    local_buf(i,j)=...
  END DO
END DO
buf = local_buf ! may overwrite asynchronously received
                ! data in buf(1,1:100)
CALL MPI_Wait(req,...)
```

1 In the compiler-generated, possible optimization in Example 19.19, `buf(100,100)` from
2 Example 19.18 is equivalenced with the 1-dimensional array `buf_1dim(10000)`. The non-
3 blocking receive may asynchronously receive the data in the boundary `buf(1,1:100)` while
4 the fused loop is temporarily using this part of the buffer. When the `tmp` data is written
5 back to `buf`, the previous data of `buf(1,1:100)` is restored and the received data is lost. The
6 principle behind this optimization is that the receive buffer data `buf(1,1:100)` was temporar-
7 ily moved to `tmp`.

8 Example 19.20 shows a second possible optimization. The whole array is temporarily
9 moved to `local_buf`.

10 When storing `local_buf` back to the original location `buf`, then this implies overwriting
11 the section of `buf` that serves as a receive buffer in the nonblocking MPI call, i.e., this
12 storing back of `local_buf` is therefore likely to interfere with asynchronously received data
13 in `buf(1,1:100)`.

14 Note that this problem may also occur:

- 15
- 16 • With the local buffer at the origin process, between an RMA communication call and
17 the ensuing synchronization call; see Chapter 12.
- 18
- 19 • With the window buffer at the target process between two ensuing RMA synchroniza-
20 tion calls.
- 21
- 22 • With the local buffer in MPI parallel file I/O split collective operations between the
23 `MPI_XXX_BEGIN` and `MPI_XXX_END` calls; see Section 14.4.5.

24 As already mentioned in Section The Fortran ASYNCHRONOUS Attribute on
25 page 825 of Section 19.1.17, the ASYNCHRONOUS attribute can prevent compiler optimization
26 with temporary data movement, but only if the receive buffer and the local references are
27 separated into different variables, as shown in Example 19.21 and in Example 19.22.

28 Note also that the methods

- 29 • calling `MPI_F_SYNC_REG` (or such a user-defined routine),
- 30
- 31 • using module variables and COMMON blocks, and
- 32
- 33 • the TARGET attribute

34 cannot be used to prevent such temporary data movement. These methods influence com-
35 piler optimization when library routines are called. They cannot prevent the optimizations
36 of the code fragments shown in Example 19.18 and 19.19.

37 Note also that compiler optimization with temporary data movement should **not** be
38 prevented by declaring `buf` as VOLATILE because the VOLATILE implies that all accesses to
39 any storage unit (word) of `buf` must be directly done in the main memory exactly in the
40 sequence defined by the application program. The VOLATILE attribute prevents all register
41 and cache optimizations. Therefore, VOLATILE may cause a huge performance degradation.

42 Instead of solving the problem, it is better to **prevent** the problem: when overlapping
43 communication and computation, the nonblocking communication (or nonblocking or split
44 collective I/O) and the computation should be executed **on different variables**, and the
45 communication should be *protected* with the ASYNCHRONOUS attribute. In this case, the
46 temporary memory modifications are done only on the variables used in the computation
47 and cannot have any side effect on the data used in the nonblocking MPI operations.
48

Rationale. This is a strong restriction for application programs. To weaken this restriction, a new or modified asynchronous feature in the Fortran language would be necessary: an asynchronous attribute that can be used on parts of an array and together with asynchronous operations outside the scope of Fortran. If such a feature becomes available in a future edition of the Fortran standard, then this restriction also may be weakened in a later version of the MPI standard. (*End of rationale.*)

In Example 19.21 (which is a solution for the problem shown in Example 19.16 and in Example 19.22 (which is a solution for the problem shown in Example 19.20), the array is split into inner and halo part and both disjoint parts are passed to a subroutine `separated_sections`. This routine overlaps the receiving of the halo data and the calculations on the inner part of the array. In a second step, the whole array is used to do the calculation on the elements where inner+halo is needed. Note that the halo and the inner area are strided arrays. Those can be used in nonblocking communication only with a Fortran 2018 (or TS 29113) based MPI library.

19.1.19 Permanent Data Movement

A Fortran compiler may implement permanent data movement during the execution of a Fortran program. This would require that pointers to such data are appropriately updated. An implementation with automatic garbage collection is one use case. Such permanent data movement is in conflict with MPI in several areas:

- MPI datatype handles with absolute addresses in combination with `MPI_BOTTOM`.
- All nonblocking MPI operations if the internally used pointers to the buffers are not updated by the Fortran runtime, or if within an MPI process, the data movement is executed in parallel with the MPI operation.

This problem can be also solved by using the `ASYNCHRONOUS` attribute for such buffers. This MPI standard requires that the problems with permanent data movement do not occur by imposing suitable restrictions on the MPI library together with the compiler used; see Section 19.1.7.

Example 19.21. Using separated variables for overlapping communication and computation to allow the protection of nonblocking communication with the `ASYNCHRONOUS` attribute.

```

USE mpi_f08
REAL :: b(0:101)      ! elements 0 and 101 are halo cells
REAL :: bnew(0:101)  ! elements 1 and 100 are newly computed
INTEGER :: i
CALL separated_sections(b(0), b(1:100), b(101), bnew(0:101))
i=1 ! compute leftmost element
  bnew(i) = function(b(i-1), b(i), b(i+1))
i=100 ! compute rightmost element
  bnew(i) = function(b(i-1), b(i), b(i+1))
END

SUBROUTINE separated_sections(b_lefthalo, b_inner, b_righthalo, bnew)
USE mpi_f08
REAL, ASYNCHRONOUS :: b_lefthalo(0:0), b_inner(1:100), b_righthalo(101:101)
REAL :: bnew(0:101) ! elements 1 and 100 are newly computed
TYPE(MPI_Request) :: req(4)

```

```

1  INTEGER :: left, right, i
2  CALL MPI_Cart_shift(...,left, right,...)
3  CALL MPI_Irecv(b_lefthalo( 0), ..., left, ..., req(1), ...)
4  CALL MPI_Irecv(b_righthalo(101), ..., right, ..., req(2), ...)
5  ! b_lefthalo and b_righthalo is written asynchronously.
6  ! There is no other concurrent access to b_lefthalo and b_righthalo.
7  CALL MPI_Isend(b_inner( 1), ..., left, ..., req(3), ...)
8  CALL MPI_Isend(b_inner(100), ..., right, ..., req(4), ...)
9
10 DO i=2,99 ! compute only elements for which halo data is not needed
11     bnew(i) = function(b_inner(i-1), b_inner(i), b_inner(i+1))
12     ! b_inner is read and sent at the same time.
13     ! This is allowed based on the rules for ASYNCHRONOUS.
14 END DO
15 CALL MPI_Waitall(4, req,...)
16 END SUBROUTINE

```

19.1.20 Comparison with C

In C, subroutines that modify variables that are not in the argument list will not cause register optimization problems. This is because taking pointers to storage objects by using the & operator and later referencing the objects by indirection on the pointer is an integral part of the language. A C compiler understands the implications, so that the problem should not occur, in general. However, some compilers do offer optional aggressive optimization levels that may not be safe. Problems due to temporary memory modifications can also occur in C. As above, the best advice is to avoid the problem: use different variables for buffers in nonblocking MPI operations and computation that is executed while a nonblocking communication operation is *pending*.

Example 19.22. Protecting GPU optimizations with the ASYNCHRONOUS attribute.

```

29 USE mpi_f08
30 REAL :: buf(100,100)
31 CALL separated_sections(buf(1:1,1:100), buf(2:100,1:100))
32 END
33
34 SUBROUTINE separated_sections(buf_halo, buf_inner)
35 REAL, ASYNCHRONOUS :: buf_halo(1:1,1:100)
36 REAL :: buf_inner(2:100,1:100)
37 REAL :: local_buf(2:100,100)
38
39 CALL MPI_Irecv(buf_halo(1,1:100),..., req,...)
40 local_buf = buf_inner
41 DO j=1,100
42     DO i=2,100
43         local_buf(i,j)=...
44     END DO
45 END DO
46 buf_inner = local_buf ! buf_halo is not touched!!!
47
48 CALL MPI_Wait(req,...)

```

19.2 Support for Large Count and Large Byte Displacement in MPI Language Bindings

The following types, which were used prior to MPI-4.0, have been deemed too small to hold values that applications wish to use:

- The C `int` type and the Fortran `INTEGER` type were used for *count* parameters.
- The C `int` type and the Fortran `INTEGER` type were used for some parameters that represent *byte displacement* in memory.
- The C `MPI_Aint` type and the Fortran `INTEGER(KIND=MPI_ADDRESS_KIND)` type were used for some parameters that represent *byte displacement* in files (e.g., in constructors of MPI datatypes that can be used with files).

In order to avoid breaking backwards compatibility, this version of MPI supports larger types via separate additional MPI procedures in C (suffixed with “_c”) and via interface polymorphism in Fortran when using `USE mpi_f08`. For better readability, all Fortran large count procedure declarations are marked with a comment “!(`_c`)”. No polymorphic support for larger types is provided in Fortran when using `mpif.h` and `use mpi`.

For the large count versions of three datatype constructors, `MPI_TYPE_CREATE_HINDEXED`, `MPI_TYPE_CREATE_HINDEXED_BLOCK`, and `MPI_TYPE_CREATE_STRUCT`, absolute addresses shall not be used to specify byte displacements since the parameter is of type `MPI_COUNT` instead of type `MPI_AINT` (see Section 2.5.8).

In addition, the functions `MPI_TYPE_GET_ENVELOPE` and `MPI_TYPE_GET_CONTENTS` also support large count types via *additional parameters* in separate additional MPI procedures in C (suffixed with “_c”) and interface polymorphism in Fortran when using `USE mpi_f08` (see Section 5.1.13).

Further, the callbacks of type `MPI_User_function` and `MPI_Datarep_conversion_function` also support large count types via separate additional callback prototypes in C (suffixed with “_c”) and multiple abstract interfaces in Fortran when using `USE mpi_f08` (see Sections 6.9.5 and 14.5.3, respectively). An additional large count predefined callback function `MPI_CONVERSION_FN_NULL_C` is provided within each of these two language bindings.

In C bindings, for each MPI procedure that had at least one *count* or *byte displacement* parameter that used the `int` and/or `MPI_Aint` types prior to MPI-4.0, an additional MPI procedure is provided, with the same name but suffixed by “_c”. The MPI procedure without the “_c” token has the same name and parameter types as versions prior to MPI-4.0. The “_c” suffixed MPI procedure has `MPI_Count` for all *count* parameters, `MPI_Aint` for parameters that represent *byte displacement* in memory, `MPI_Offset` for parameters that represent *byte displacement* in files, and `MPI_Count` for parameters that may represent *byte displacement* in both memory and files.

In Fortran, when using `USE mpi_f08`, for each MPI procedure that had at least one *count* or *byte displacement* parameter that used the `INTEGER` or `INTEGER(KIND=MPI_ADDRESS_KIND)` types prior to MPI-4.0, a polymorphic interface containing two specific procedures is provided. One of the specific procedures has the same name and dummy parameter types as in versions prior to MPI-4.0. `INTEGER` and/or `INTEGER(KIND=MPI_ADDRESS_KIND)` for *count* and *byte displacement* parameters. The other specific procedure has the same name followed by “_c”, and then suffixed by the token specified in Table 19.1 for `USE mpi_f08`.

1 It also has `INTEGER(KIND=MPI_COUNT_KIND)` for all *count* parameters,
 2 `INTEGER(KIND=MPI_ADDRESS_KIND)` for parameters that represent *byte displacement* in mem-
 3 ory, `INTEGER(KIND=MPI_OFFSET_KIND)` for parameters that represent *byte displacement* in
 4 files, and `INTEGER(KIND=MPI_COUNT_KIND)` for parameters that may represent *byte displace-*
 5 *ment* in both memory and files (for more details on specific Fortran procedure names and
 6 related calling conventions, refer to Table 19.1 in Section 19.1.5). There is one exception:
 7 if the type signatures of the two specific procedures are identical (e.g., if
 8 `INTEGER(KIND=MPI_COUNT_KIND)` is the same type as `INTEGER(KIND=MPI_ADDRESS_KIND)`),
 9 then the implementation shall not provide the “_c” specific procedure.

10 It is erroneous to directly invoke the “_c” specific procedures in the Fortran `mpi_f08`
 11 module with the exception of the following procedures: `MPI_Op_create_c` and
 12 `MPI_Register_datarep_c`.

13 In older Fortran bindings (`mpif.h` (deprecated) and `use mpi`), no new interfaces and
 14 no new specific procedures for larger types are provided beyond what existed in MPI-3.1;
 15 all MPI procedures have the same types as in the versions prior to MPI-4.0.

17 19.3 Language Interoperability

18 19.3.1 Introduction

19 It is not uncommon for library developers to use one language to develop an application
 20 library that may be called by an application program written in a different language. MPI
 21 currently supports ISO (previously ANSI) C and Fortran bindings. It should be possible
 22 for applications in any of the supported languages to call MPI-related functions in another
 23 language.
 24 language.

25 Moreover, MPI allows the development of client-server code, with MPI communication
 26 used between a parallel client and a parallel server. It should be possible to code the server
 27 in one language and the clients in another language. To do so, communications should be
 28 possible between applications written in different languages.

29 There are several issues that need to be addressed in order to achieve interoperability.

30 **Initialization:** We need to specify how the MPI environment is initialized for all languages.

31 **Interlanguage passing of MPI opaque objects:** We need to specify how MPI object
 32 handles are passed between languages. We also need to specify what happens when
 33 an MPI object is accessed in one language, to retrieve information (e.g., attributes)
 34 set in another language.
 35

36 **Interlanguage communication:** We need to specify how messages sent in one language
 37 can be received in another language.

38 It is highly desirable that the solution for interlanguage interoperability be extensible
 39 to new languages, should MPI bindings be defined for such languages.

40 19.3.2 Assumptions

41 We assume that conventions exist for programs written in one language to call routines
 42 written in another language. These conventions specify how to link routines in different
 43 languages into one program, how to call functions in a different language, how to pass
 44 arguments between languages, and the correspondence between basic datatypes in different
 45

languages. In general, these conventions will be implementation dependent. Furthermore, not every basic datatype may have a matching type in other languages. For example, C character strings may not be compatible with Fortran CHARACTER variables. However, we assume that a Fortran INTEGER, as well as a (sequence associated) Fortran array of INTEGERS, can be passed to a C program. We also assume that Fortran and C have address-sized integers. This does not mean that the default-size integers are the same size as default-sized pointers, but only that there is some way to hold (and pass) a C address in a Fortran integer. It is also assumed that INTEGER(KIND=MPI_OFFSET_KIND) can be passed from Fortran to C as MPI_Offset.

19.3.3 Initialization

Two approaches are available for initializing MPI: the World Model(Section 11.2) , and the Sessions Model(Section 11.3).

Concerns specific to the World Model

A call to MPI_INIT or MPI_INIT_THREAD, from any language, initializes MPI for execution in all languages.

Advice to users. Certain implementations use the (inout) argc, argv arguments of the C version of MPI_INIT in order to propagate values for argc and argv to all executing MPI processes. Use of the Fortran version of MPI_INIT to initialize MPI may result in a loss of this ability. (*End of advice to users.*)

The function MPI_INITIALIZED returns the same answer in all languages.

The function MPI_FINALIZE finalizes the MPI environments for all languages.

The function MPI_FINALIZED returns the same answer in all languages.

The MPI environment is initialized in the same manner for all languages by MPI_INIT. E.g., MPI_COMM_WORLD carries the same information regardless of language: same MPI processes, same environmental attributes, same error handlers.

Advice to users. The use of several languages in one MPI program may require the use of special options at compile and/or link time. (*End of advice to users.*)

Advice to implementors. Implementations may selectively link language specific MPI libraries only to codes that need them, so as not to increase the size of binaries for codes that use only one language. The MPI initialization code needs to perform initialization for a language only if that language library is loaded. (*End of advice to implementors.*)

Concerns specific to the Sessions Model

A call to MPI_SESSION_INIT from any language initializes a session that can be used from all languages.

A call to MPI_SESSION_FINALIZE from any language finalizes the session for all languages.

1 *Concerns common to both the World Model and the Sessions Model*

2 The function `MPI_ABORT` kills MPI processes in the group of the supplied communicator,
3 irrespective of the language used by the caller or by the MPI processes killed.

4 Information can be added to info objects in one language and retrieved in another.
5

6 19.3.4 Transfer of Handles

7 Handles are passed between Fortran and C by using an explicit C wrapper to convert Fortran
8 handles to C handles. There is no direct access to C handles in Fortran.

9 The type definition `MPI_Fint` is provided in C for an integer of the size that matches a
10 Fortran `INTEGER`; usually, `MPI_Fint` will be equivalent to `int`. With the Fortran `mpi` module
11 or the (deprecated) `mpif.h` include file, a Fortran handle is a Fortran `INTEGER` value that can
12 be used in the following conversion functions. With the Fortran `mpi_f08` module, a Fortran
13 handle is a `BIND(C)` derived type that contains an `INTEGER` component named `MPI_VAL`.
14 This `INTEGER` value can be used in the following conversion functions.
15

16 The following functions are provided in C to convert from a Fortran communicator
17 handle (which is an integer) to a C communicator handle, and vice versa.
18

19 **C binding**

20 `MPI_Comm MPI_Comm_f2c(MPI_Fint comm)`

21 If `comm` is a valid Fortran handle to a communicator, then `MPI_Comm_f2c` returns a
22 valid C handle to that same communicator; if `comm = MPI_COMM_NULL` (Fortran value),
23 then `MPI_Comm_f2c` returns a null C handle; if `comm` is an invalid Fortran handle, then
24 `MPI_Comm_f2c` returns an invalid C handle.
25

26 `MPI_Fint MPI_Comm_c2f(MPI_Comm comm)`

27 The function `MPI_Comm_c2f` translates a C communicator handle into a Fortran handle
28 to the same communicator; it maps a null handle into a null handle and an invalid handle
29 into an invalid handle.

30 Similar functions are provided for the other types of opaque objects.

31 `MPI_Datatype MPI_Type_f2c(MPI_Fint datatype)`

32 `MPI_Fint MPI_Type_c2f(MPI_Datatype datatype)`

33 `MPI_Group MPI_Group_f2c(MPI_Fint group)`

34 `MPI_Fint MPI_Group_c2f(MPI_Group group)`

35 `MPI_Request MPI_Request_f2c(MPI_Fint request)`

36 `MPI_Fint MPI_Request_c2f(MPI_Request request)`

37 `MPI_File MPI_File_f2c(MPI_Fint file)`

38 `MPI_Fint MPI_File_c2f(MPI_File file)`

39 `MPI_Win MPI_Win_f2c(MPI_Fint win)`

40 `MPI_Fint MPI_Win_c2f(MPI_Win win)`

41 `MPI_Op MPI_Op_f2c(MPI_Fint op)`

42 `MPI_Fint MPI_Op_c2f(MPI_Op op)`
43
44
45
46
47
48

```

MPI_Info MPI_Info_f2c(MPI_Fint info)
MPI_Fint MPI_Info_c2f(MPI_Info info)
MPI_Errhandler MPI_Errhandler_f2c(MPI_Fint errhandler)
MPI_Fint MPI_Errhandler_c2f(MPI_Errhandler errhandler)
MPI_Message MPI_Message_f2c(MPI_Fint message)
MPI_Fint MPI_Message_c2f(MPI_Message message)
MPI_Session MPI_Session_f2c(MPI_Fint session)
MPI_Fint MPI_Session_c2f(MPI_Session session)

```

Example 19.23. The example below illustrates how the Fortran MPI function `MPI_TYPE_COMMIT` can be implemented by wrapping the C MPI function `MPI_Type_commit` with a C wrapper to do handle conversions. In this example a Fortran-C interface is assumed where a Fortran function is all upper case when referred to from C and arguments are passed by addresses.

```

! FORTRAN PROCEDURE
SUBROUTINE MPI_TYPE_COMMIT(DATATYPE, IERR)
INTEGER :: DATATYPE, IERR
CALL MPI_X_TYPE_COMMIT(DATATYPE, IERR)
RETURN
END

/* C wrapper */

void MPI_X_TYPE_COMMIT(MPI_Fint *f_handle, MPI_Fint *ierr)
{
    MPI_Datatype datatype;

    datatype = MPI_Type_f2c(*f_handle);
    *ierr = (MPI_Fint)MPI_Type_commit(&datatype);
    *f_handle = MPI_Type_c2f(datatype);
    return;
}

```

The same approach can be used for all other MPI functions. The call to `MPI_XXX_f2c` (resp. `MPI_XXX_c2f`) can be omitted when the handle is an OUT (resp. IN) argument, rather than INOUT.

Rationale. The design here provides a convenient solution for the prevalent case, where a C wrapper is used to allow Fortran code to call a C library, or C code to call a Fortran library. The use of C wrappers is much more likely than the use of Fortran wrappers, because it is much more likely that a variable of type `INTEGER` can be passed to C, than a C handle can be passed to Fortran.

Returning the converted value as a function value rather than through the argument list allows the generation of efficient inlined code when these functions are simple (e.g., the identity). The conversion function in the wrapper does not catch an invalid handle argument. Instead, an invalid handle is passed below to the library function, which, presumably, checks its input arguments. (*End of rationale.*)

19.3.5 Status

The following two procedures are provided in C to convert from a Fortran (with the `mpi` module or deprecated `mpif.h`) status (which is an array of integers) to a C status (which is a structure), and vice versa. The conversion occurs on all the information in status, including that which is hidden. That is, no status information is lost in the conversion.

```
int MPI_Status_f2c(const MPI_Fint *f_status, MPI_Status *c_status)
```

If `f_status` is a valid Fortran status, but not the Fortran value of `MPI_STATUS_IGNORE` or `MPI_STATUSES_IGNORE`, then `MPI_Status_f2c` returns in `c_status` a valid C status with the same content. If `f_status` is the Fortran value of `MPI_STATUS_IGNORE` or `MPI_STATUSES_IGNORE`, or if `f_status` is not a valid Fortran status, then the call is erroneous.

In C, such an `f_status` array can be defined with `MPI_Fint f_status[MPI_F_STATUS_SIZE]`. Within this array, one can use in C the indexes `MPI_F_SOURCE`, `MPI_F_TAG`, and `MPI_F_ERROR`, to access the same elements as in Fortran with `MPI_SOURCE`, `MPI_TAG` and `MPI_ERROR`. The C indexes are 1 less than the corresponding indexes in Fortran due to the different default array start indexes in both languages.

The C status has the same source, tag and error code values as the Fortran status, and returns the same answers when queried for count, elements, and cancellation. The conversion function may be called with a Fortran status argument that has an undefined error field, in which case the value of the error field in the C status argument is undefined.

Two global variables of type `MPI_Fint*`, `MPI_F_STATUS_IGNORE` and `MPI_F_STATUSES_IGNORE` are declared in `mpi.h`. They can be used to test, in C, whether `f_status` is the Fortran value of `MPI_STATUS_IGNORE` or `MPI_STATUSES_IGNORE` defined in the `mpi` module or (deprecated) `mpif.h`. These are global variables, not C constant expressions and cannot be used in places where C requires constant expressions. Their value is defined only between the calls to `MPI_INIT` and `MPI_FINALIZE` and should not be changed by user code.

To do the conversion in the other direction, we have the following:

```
int MPI_Status_c2f(const MPI_Status *c_status, MPI_Fint *f_status)
```

This call converts a C status into a Fortran status, and has a behavior similar to `MPI_Status_f2c`. That is, the value of `c_status` must not be either `MPI_STATUS_IGNORE` or `MPI_STATUSES_IGNORE`.

Advice to users. There exists no separate conversion function for arrays of statuses, since one can simply loop through the array, converting each status with the routines in Figure 19.1. (*End of advice to users.*)

Rationale. The handling of `MPI_STATUS_IGNORE` is required in order to layer libraries with only a C wrapper: if the Fortran call has passed `MPI_STATUS_IGNORE`, then the C wrapper must handle this correctly. Note that this constant need not have the same value in Fortran and C. If `MPI_Status_f2c` were to handle `MPI_STATUS_IGNORE`, then the type of its result would have to be `MPI_Status**`, which was considered an inferior solution. (*End of rationale.*)

Using the `mpi_f08` Fortran module, a status is declared as `TYPE(MPI_Status)`. The C type `MPI_F08_status` can be used to pass a Fortran `TYPE(MPI_Status)` argument into a C routine. Figure 19.1 illustrates all status conversion routines. Some are only available in

C, some in both C and the Fortran `mpi` and `mpi_f08` interfaces (but not in the deprecated `mpif.h` include file).

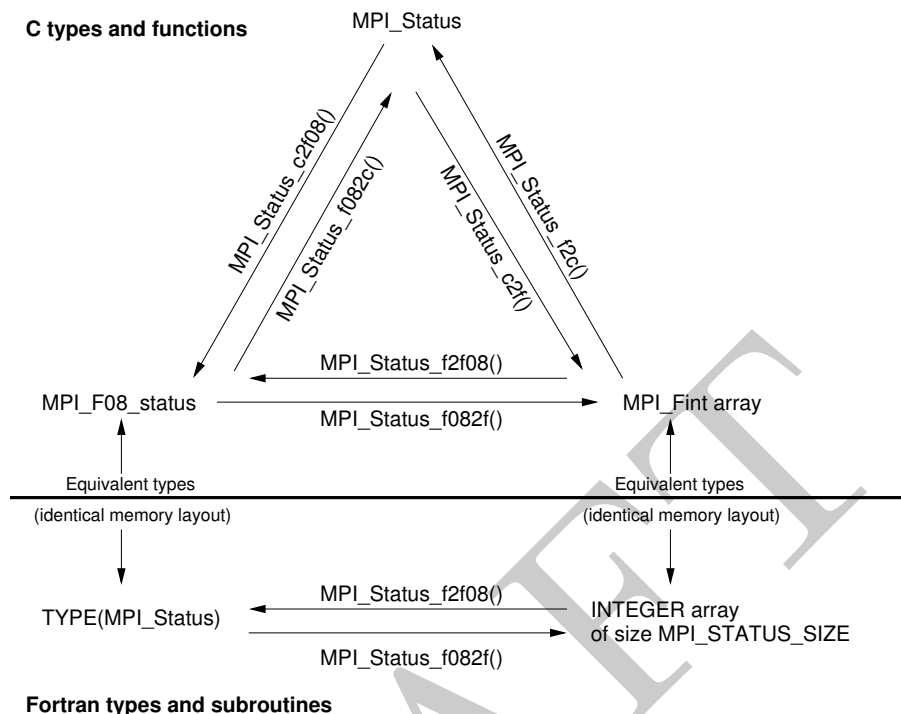


Figure 19.1: Status conversion routines

```
int MPI_Status_f082c(const MPI_F08_status *f08_status, MPI_Status *c_status)
```

This C routine converts a Fortran `mpi_f08` `TYPE(MPI_Status)` into a C `MPI_Status`.

```
int MPI_Status_c2f08(const MPI_Status *c_status, MPI_F08_status *f08_status)
```

This C routine converts a C `MPI_Status` into a Fortran `mpi_f08` `TYPE(MPI_Status)`. Two global variables of type `MPI_F08_status*`, `MPI_F08_STATUS_IGNORE` and `MPI_F08_STATUSES_IGNORE` are declared in `mpi.h`. They can be used to test, in C, whether `f_status` is the Fortran value of `MPI_STATUS_IGNORE` or `MPI_STATUSES_IGNORE` defined in the `mpi_f08` module. These are global variables, not C constant expressions and cannot be used in places where C requires constant expressions. Their value is defined only between the calls to `MPI_INIT` and `MPI_FINALIZE` and should not be changed by user code.

Conversion between the two Fortran versions of a status can be done with:

```
MPI_STATUS_F2F08(f_status, f08_status)
```

IN `f_status` status object declared as array (status)

OUT `f08_status` status object declared as named type (status)

C binding

```
int MPI_Status_f2f08(const MPI_Fint *f_status, MPI_F08_status *f08_status)
```

Fortran 2008 binding

```

1 MPI_Status_f2f08(f_status, f08_status, ierror)
2     INTEGER, INTENT(IN) :: f_status(MPI_STATUS_SIZE)
3     TYPE(MPI_Status), INTENT(OUT) :: f08_status
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6

```

Fortran binding (the following procedure is not available with mpif.h)

```

7 MPI_STATUS_F2F08(F_STATUS, F08_STATUS, IERROR)
8     INTEGER :: F_STATUS(MPI_STATUS_SIZE), IERROR
9     TYPE(MPI_Status) :: F08_STATUS
10
11

```

This routine converts a Fortran INTEGER, DIMENSION(MPI_STATUS_SIZE) status array into a Fortran mpi_f08 TYPE(MPI_Status).

```

12
13
14 MPI_STATUS_F082F(f08_status, f_status)
15

```

16	IN	f08_status	status object declared as named type (status)
17	OUT	f_status	status object declared as array (status)

C binding

```

18
19
20 int MPI_Status_f082f(const MPI_F08_status *f08_status, MPI_Fint *f_status)
21
22

```

Fortran 2008 binding

```

23 MPI_Status_f082f(f08_status, f_status, ierror)
24     TYPE(MPI_Status), INTENT(IN) :: f08_status
25     INTEGER, INTENT(OUT) :: f_status(MPI_STATUS_SIZE)
26     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28

```

Fortran binding (the following procedure is not available with mpif.h)

```

29 MPI_STATUS_F082F(F08_STATUS, F_STATUS, IERROR)
30     TYPE(MPI_Status) :: F08_STATUS
31     INTEGER :: F_STATUS(MPI_STATUS_SIZE), IERROR
32
33

```

This routine converts a Fortran mpi_f08 TYPE(MPI_Status) into a Fortran INTEGER, DIMENSION(MPI_STATUS_SIZE) status array.

19.3.6 MPI Opaque Objects

Unless said otherwise, opaque objects are “the same” in all languages: they carry the same information, and have the same meaning in both languages. The mechanism described in the previous section can be used to pass references to MPI objects from language to language. An object created in one language can be accessed, modified or freed in another language.

We examine below in more detail issues that arise for each type of MPI object.

Datatypes

Datatypes encode the same information in all languages. E.g., a datatype accessor like MPI_TYPE_GET_EXTENT will return the same information in all languages. If a datatype defined in one language is used for a communication call in another language, then the message sent will be identical to the message that would be sent from the first language:

the same communication buffer is accessed, and the same representation conversion is performed, if needed. All predefined datatypes can be used in datatype constructors in any language. If a datatype is committed, it can be used for communication in any language.

The function `MPI_GET_ADDRESS` returns the same value in all languages. Note that we do not require that the constant `MPI_BOTTOM` have the same value in all languages (see Section 19.3.9).

Example 19.24. Absolute addresses and the conversion of datatype handles in a mixed Fortran/C program.

```

! FORTRAN CODE
REAL :: R(5)
INTEGER :: DTYPE, IERR, AOBLLEN(1), AOTYPE(1)
INTEGER(KIND=MPI_ADDRESS_KIND) :: AODISP(1)

! create an absolute datatype for array R
AOBLLEN(1) = 5
CALL MPI_GET_ADDRESS(R, AODISP(1), IERR)
AOTYPE(1) = MPI_REAL
CALL MPI_TYPE_CREATE_STRUCT(1, AOBLLEN, AODISP, AOTYPE, DTYPE, IERR)
CALL C_ROUTINE(DTYPE)

```

```

/* C code */

void C_ROUTINE(MPI_Fint *ftype)
{
    int count = 5;
    int lens[2] = {1,1};
    MPI_Aint displs[2];
    MPI_Datatype types[2], newtype;

    /* create an absolute datatype for buffer that consists */
    /* of count, followed by R(5) */

    MPI_Get_address(&count, &displs[0]);
    displs[1] = 0;
    types[0] = MPI_INT;
    types[1] = MPI_Type_f2c(*ftype);
    MPI_Type_create_struct(2, lens, displs, types, &newtype);
    MPI_Type_commit(&newtype);

    MPI_Send(MPI_BOTTOM, 1, newtype, 1, 0, MPI_COMM_WORLD);
    /* the message sent contains an int count of 5, followed */
    /* by the 5 REAL entries of the Fortran array R. */
}

```

Advice to implementors. The following implementation can be used: MPI addresses, as returned by `MPI_GET_ADDRESS`, will have the same value in all languages. One obvious choice is that MPI addresses be identical to regular addresses. The address is stored in the datatype, when datatypes with absolute addresses are constructed. When a send or receive operation is performed, then addresses stored in a datatype are interpreted as displacements that are all augmented by a base address. This base

1 address is (the address of) `buf`, or zero, if `buf = MPI_BOTTOM`. Thus, if `MPI_BOTTOM`
2 is zero then a send or receive call with `buf = MPI_BOTTOM` is implemented exactly as
3 a call with a regular buffer argument: in both cases the base address is `buf`. On the
4 other hand, if `MPI_BOTTOM` is not zero, then the implementation has to be slightly
5 different. A test is performed to check whether `buf = MPI_BOTTOM`. If true, then the
6 base address is zero, otherwise it is `buf`. In particular, if `MPI_BOTTOM` does not have
7 the same value in Fortran and C, then an additional test for `buf = MPI_BOTTOM` is
8 needed in at least one of the languages.

9 It may be desirable to use a value other than zero for `MPI_BOTTOM` even in C, so as
10 to distinguish it from a NULL pointer. If `MPI_BOTTOM = c` then one can still avoid
11 the test `buf = MPI_BOTTOM`, by using the displacement from `MPI_BOTTOM`, i.e., the
12 regular address - `c`, as the MPI address returned by `MPI_GET_ADDRESS` and stored
13 in absolute datatypes. (*End of advice to implementors.*)
14

15 *Callback Functions*

16 MPI calls may associate callback functions with MPI objects: error handlers are associated
17 with communicators, files, windows, and sessions; attribute copy and delete functions are
18 associated with attribute keys; reduce operations are associated with operation objects, etc.
19 In a multilanguage environment, a function passed in an MPI call in one language may be
20 invoked by an MPI call in another language. MPI implementations must make sure that
21 such invocation will use the calling convention of the language the function is bound to.
22

23
24 *Advice to implementors.* Callback functions need to have a language tag. This
25 tag is set when the callback function is passed in by the library function (which is
26 presumably different for each language and language support method), and is used
27 to generate the right calling sequence when the callback function is invoked. (*End of*
28 *advice to implementors.*)
29

30 *Advice to users.* If a subroutine written in one language or Fortran support method
31 wants to pass a callback routine including the predefined Fortran functions (e.g.,
32 `MPI_COMM_NULL_COPY_FN`) to another application routine written in another lan-
33 guage or Fortran support method, then it must be guaranteed that both routines use
34 the callback interface definition that is defined for the argument when passing the
35 callback to an MPI routine (e.g., `MPI_COMM_CREATE_KEYVAL`); see also the advice
36 to users on page 364. (*End of advice to users.*)
37

38 *Error Handlers*

39 *Advice to implementors.* Error handlers, have, in C, a variable length argument list.
40 It might be useful to provide to the handler information on the language environment
41 where the error occurred. (*End of advice to implementors.*)
42

43 *Reduce Operations*

44 All predefined named and unnamed datatypes as listed in Section 6.9.2 can be used in the
45 listed predefined operations independent of the programming language from which the MPI
46 routine is called.
47
48

Advice to users. Reduce operations receive as one of their arguments the datatype of the operands. Thus, one can define “polymorphic” reduce operations that work for C and Fortran datatypes. (*End of advice to users.*)

19.3.7 Attributes

Attribute keys can be allocated in one language and freed in another. Similarly, attribute values can be set in one language and accessed in another. To achieve this, attribute keys will be allocated in an integer range that is valid all languages. The same holds true for system-defined attribute values (such as `MPI_TAG_UB`, `MPI_WTIME_IS_GLOBAL`, etc.).

Attribute keys declared in one language are associated with copy and delete functions in that language (the functions provided by the `MPI_XXX_CREATE_KEYVAL` call). When a communicator is duplicated, for each attribute, the corresponding copy function is called, using the right calling convention for the language of that function; and similarly, for the delete callback function.

Advice to implementors. This requires that attributes be tagged either as “C” or “Fortran” and that the language tag be checked in order to use the right calling convention for the callback function. (*End of advice to implementors.*)

The attribute manipulation functions described in Section 7.7 defines attributes arguments to be of type `void*` in C, and of type `INTEGER`, in Fortran. On some systems, `INTEGER`s will have 32 bits, while C pointers will have 64 bits. This is a problem if communicator attributes are used to move information from a Fortran caller to a C callee, or vice-versa.

MPI behaves as if it stores, internally, address sized attributes. If Fortran `INTEGER`s are smaller, then the (deprecated) Fortran function `MPI_ATTR_GET` will return the least significant part of the attribute word; the (deprecated) Fortran function `MPI_ATTR_PUT` will set the least significant part of the attribute word, which will be sign extended to the entire word. (These two functions may be invoked explicitly by user code, or implicitly, by attribute copying callback functions.)

As for addresses, new functions are provided that manipulate Fortran address sized attributes, and have the same functionality as the old functions in C. These functions are described in Section 7.7. Users are encouraged to use these new functions.

MPI supports two types of attributes: address-valued (pointer) attributes, and integer-valued attributes. C attribute functions put and get address-valued attributes. Fortran attribute functions put and get integer-valued attributes. When an integer-valued attribute is accessed from C, then `MPI_XXX_get_attr` will return the address of (a pointer to) the integer-valued attribute, which is a pointer to `MPI_Aint` if the attribute was stored with Fortran `MPI_XXX_SET_ATTR`, and a pointer to `int` if it was stored with the deprecated Fortran `MPI_ATTR_PUT`. When an address-valued attribute is accessed from Fortran, then `MPI_XXX_GET_ATTR` will convert the address into an integer and return the result of this conversion. This conversion is lossless if new style attribute functions are used, and an integer of kind `MPI_ADDRESS_KIND` is returned. The conversion may cause truncation if deprecated attribute functions are used. In C, the deprecated routines `MPI_Attr_put` and `MPI_Attr_get` behave identical to `MPI_Comm_set_attr` and `MPI_Comm_get_attr`.

Example 19.25. Setting an attribute in C and reading in C or Fortran.
A. Setting an attribute value in C

```

1  int set_val = 3;
2  struct foo set_struct;
3
4  /* Set a value that is a pointer to an int */
5
6  MPI_Comm_set_attr(MPI_COMM_WORLD, keyval1, &set_val);
7  /* Set a value that is a pointer to a struct */
8  MPI_Comm_set_attr(MPI_COMM_WORLD, keyval2, &set_struct);
9  /* Set an integer value */
10 MPI_Comm_set_attr(MPI_COMM_WORLD, keyval3, (void *) 17);

```

B. Reading the attribute value in C

```

12 int flag, *get_val;
13 struct foo *get_struct;
14
15 /* Upon successful return, get_val == &set_val
16    (and therefore *get_val == 3) */
17 MPI_Comm_get_attr(MPI_COMM_WORLD, keyval1, &get_val, &flag);
18 /* Upon successful return, get_struct == &set_struct */
19 MPI_Comm_get_attr(MPI_COMM_WORLD, keyval2, &get_struct, &flag);
20 /* Upon successful return, get_val == (void*) 17 */
21 /*      i.e., (MPI_Aint) get_val == 17 */
22 MPI_Comm_get_attr(MPI_COMM_WORLD, keyval3, &get_val, &flag);

```

C. Reading the attribute value with (deprecated) Fortran MPI-1 calls

```

24 LOGICAL FLAG
25 INTEGER IERR, GET_VAL, GET_STRUCT
26
27 ! Upon successful return, GET_VAL == &set_val, possibly truncated
28 CALL MPI_ATTR_GET(MPI_COMM_WORLD, KEYVAL1, GET_VAL, FLAG, IERR)
29 ! Upon successful return, GET_STRUCT == &set_struct, possibly truncated
30 CALL MPI_ATTR_GET(MPI_COMM_WORLD, KEYVAL2, GET_STRUCT, FLAG, IERR)
31 ! Upon successful return, GET_VAL == 17
32 CALL MPI_ATTR_GET(MPI_COMM_WORLD, KEYVAL3, GET_VAL, FLAG, IERR)

```

D. Reading the attribute value with Fortran MPI-2 calls

```

34 LOGICAL FLAG
35 INTEGER IERR
36 INTEGER(KIND=MPI_ADDRESS_KIND) GET_VAL, GET_STRUCT
37
38 ! Upon successful return, GET_VAL == &set_val
39 CALL MPI_COMM_GET_ATTR(MPI_COMM_WORLD, KEYVAL1, GET_VAL, FLAG, IERR)
40 ! Upon successful return, GET_STRUCT == &set_struct
41 CALL MPI_COMM_GET_ATTR(MPI_COMM_WORLD, KEYVAL2, GET_STRUCT, FLAG, IERR)
42 ! Upon successful return, GET_VAL == 17
43 CALL MPI_COMM_GET_ATTR(MPI_COMM_WORLD, KEYVAL3, GET_VAL, FLAG, IERR)

```

Example 19.26. Setting an attribute in Fortran and reading in C or Fortran.

A. Setting an attribute value with the (deprecated) Fortran MPI-1 call

```

47 INTEGER IERR, VAL

```

```

VAL = 7
CALL MPI_ATTR_PUT(MPI_COMM_WORLD, KEYVAL, VAL, IERR)

```

B. Reading the attribute value in C

```

int flag;
int *value;

/* Upon successful return, value points to internal MPI storage and
   *value == (int) 7 */
MPI_Comm_get_attr(MPI_COMM_WORLD, keyval, &value, &flag);

```

C. Reading the attribute value with (deprecated) Fortran MPI-1 calls

```

LOGICAL FLAG
INTEGER IERR, VALUE

! Upon successful return, VALUE == 7
CALL MPI_ATTR_GET(MPI_COMM_WORLD, KEYVAL, VALUE, FLAG, IERR)

```

D. Reading the attribute value with Fortran MPI-2 calls

```

LOGICAL FLAG
INTEGER IERR
INTEGER(KIND=MPI_ADDRESS_KIND) VALUE

! Upon successful return, VALUE == 7 (sign extended)
CALL MPI_COMM_GET_ATTR(MPI_COMM_WORLD, KEYVAL, VALUE, FLAG, IERR)

```

Example 19.27. Setting an attribute in Fortran and reading in C or Fortran.

A. Setting an attribute value via a Fortran MPI-2 call

```

INTEGER IERR
INTEGER(KIND=MPI_ADDRESS_KIND) VALUE1
INTEGER(KIND=MPI_ADDRESS_KIND) VALUE2
VALUE1 = 42
VALUE2 = INT(2, KIND=MPI_ADDRESS_KIND) ** 40

CALL MPI_COMM_SET_ATTR(MPI_COMM_WORLD, KEYVAL1, VALUE1, IERR)
CALL MPI_COMM_SET_ATTR(MPI_COMM_WORLD, KEYVAL2, VALUE2, IERR)

```

B. Reading the attribute value in C

```

int flag;
MPI_Aint *value1, *value2;

/* Upon successful return, value1 points to internal MPI storage and
   *value1 == 42 */
MPI_Comm_get_attr(MPI_COMM_WORLD, keyval1, &value1, &flag);
/* Upon successful return, value2 points to internal MPI storage and
   *value2 == 2^40 */
MPI_Comm_get_attr(MPI_COMM_WORLD, keyval2, &value2, &flag);

```

C. Reading the attribute value with (deprecated) Fortran MPI-1 calls

```

LOGICAL FLAG

```

```

1  INTEGER IERR, VALUE1, VALUE2
2
3  ! Upon successful return, VALUE1 == 42
4  CALL MPI_ATTR_GET(MPI_COMM_WORLD, KEYVAL1, VALUE1, FLAG, IERR)
5  ! Upon successful return, VALUE2 == 2^40, or 0 if truncation
6  ! needed (i.e., the least significant part of the attribute word)
7  CALL MPI_ATTR_GET(MPI_COMM_WORLD, KEYVAL2, VALUE2, FLAG, IERR)

```

D. Reading the attribute value with Fortran MPI-2 calls

```

9  LOGICAL FLAG
10 INTEGER IERR
11 INTEGER(KIND=MPI_ADDRESS_KIND) VALUE1, VALUE2
12
13 ! Upon successful return, VALUE1 == 42
14 CALL MPI_COMM_GET_ATTR(MPI_COMM_WORLD, KEYVAL1, VALUE1, FLAG, IERR)
15 ! Upon successful return, VALUE2 == 2^40
16 CALL MPI_COMM_GET_ATTR(MPI_COMM_WORLD, KEYVAL2, VALUE2, FLAG, IERR)

```

The predefined MPI attributes can be integer valued or address-valued. Predefined integer valued attributes, such as `MPI_TAG_UB`, behave as if they were put by a call to the deprecated Fortran routine `MPI_ATTR_PUT`, i.e., in Fortran, `MPI_COMM_GET_ATTR(MPI_COMM_WORLD, MPI_TAG_UB, val, flag, ierr)` will return in `val` the upper bound for tag value; in C, `MPI_Comm_get_attr(MPI_COMM_WORLD, MPI_TAG_UB, &p, &flag)` will return in `p` a pointer to an int containing the upper bound for tag value.

Address-valued predefined attributes, such as `MPI_WIN_BASE` behave as if they were put by a C call, i.e., in Fortran, `MPI_WIN_GET_ATTR(win, MPI_WIN_BASE, val, flag, ierr)` will return in `val` the base address of the window, converted to an integer. In C, `MPI_Win_get_attr(win, MPI_WIN_BASE, &p, &flag)` will return in `p` a pointer to the window base, cast to `(void *)`.

Rationale. The design is consistent with the behavior specified for predefined attributes, and ensures that no information is lost when attributes are passed from language to language. Because the language interoperability for predefined attributes was defined based on `MPI_ATTR_PUT`, this definition is kept for compatibility reasons although the routine itself is now deprecated. (*End of rationale.*)

Advice to implementors. Implementations should tag attributes either as (1) address attributes, (2) as `INTEGER(KIND=MPI_ADDRESS_KIND)` attributes or (3) as `INTEGER` attributes, according to whether they were set in (1) C (with `MPI_Attr_put` or `MPI_XXX_set_attr`), (2) in Fortran with `MPI_XXX_SET_ATTR` or (3) with the deprecated Fortran routine `MPI_ATTR_PUT`. Thus, the right choice can be made when the attribute is retrieved. (*End of advice to implementors.*)

19.3.8 Extra-State

Extra-state should not be modified by the copy or delete callback functions. (This is obvious from the C binding, but not obvious from the Fortran binding). However, these functions may update state that is indirectly accessed via extra-state. E.g., in C, extra-state can be a pointer to a data structure that is modified by the copy or callback functions; in Fortran,

extra-state can be an index into an entry in a COMMON array that is modified by the copy or callback functions. In a multithreaded environment, users should be aware that distinct threads may invoke the same callback function concurrently: if this function modifies state associated with extra-state, then mutual exclusion code must be used to protect updates and accesses to the shared state.

19.3.9 Constants

MPI constants have the same value in all languages, unless specified otherwise. This does not apply to constant handles (MPI_INT, MPI_COMM_WORLD, MPI_ERRORS_RETURN, MPI_SUM, etc.) These handles need to be converted, as explained in Section 19.3.4. Constants that specify maximum lengths of strings (see Section A.1.1 for a listing) have a value one less in Fortran than C since in C the length includes the null terminating character. Thus, these constants represent the amount of space that must be allocated to hold the largest possible such string, rather than the maximum number of printable characters the string could contain.

Advice to users. This definition means that it is safe in C to allocate a buffer to receive a string using a declaration like

```
char name [MPI_MAX_OBJECT_NAME];
```

(*End of advice to users.*)

Also constant “addresses,” i.e., special values for reference arguments that are not handles, such as MPI_BOTTOM or MPI_STATUS_IGNORE may have different values in different languages.

Rationale. The current MPI standard specifies that MPI_BOTTOM can be used in initialization expressions in C, but not in Fortran. Since Fortran does not normally support call by value, then MPI_BOTTOM in Fortran must be the name of a predefined static variable, e.g., a variable in an MPI declared COMMON block. On the other hand, in C, it is natural to take MPI_BOTTOM = 0 (Caveat: Defining MPI_BOTTOM = 0 implies that NULL pointer cannot be distinguished from MPI_BOTTOM; it may be that MPI_BOTTOM = 1 is better. See the advice to implementors in the [Datatypes](#) subsection in Section 19.3.6) Requiring that the Fortran and C values be the same will complicate the initialization process. (*End of rationale.*)

19.3.10 Interlanguage Communication

The type matching rules for communication in MPI are not changed: the datatype specification for each item sent should match, in type signature, the datatype specification used to receive this item (unless one of the types is MPI_PACKED). Also, the type of a message item should match the type declaration for the corresponding communication buffer location, unless the type is MPI_BYTE or MPI_PACKED. Interlanguage communication is allowed if it complies with these rules.

Example 19.28. In the example below, a Fortran array is sent from Fortran and received in C.

```
! FORTRAN CODE
```

```

1  SUBROUTINE MYEXAMPLE()
2  USE mpi_f08
3  REAL :: R(5)
4  INTEGER :: IERR, MYRANK, AOBLEN(1)
5  TYPE(MPI_Datatype) :: DTYPE, AOTYPE(1)
6  INTEGER(KIND=MPI_ADDRESS_KIND) :: AODISP(1)
7
8  ! create an absolute datatype for array R
9  AOBLEN(1) = 5
10 CALL MPI_GET_ADDRESS(R, AODISP(1), IERR)
11 AOTYPE(1) = MPI_REAL
12 CALL MPI_TYPE_CREATE_STRUCT(1, AOBLEN, AODISP, AOTYPE, DTYPE, IERR)
13 CALL MPI_TYPE_COMMIT(DTYPE, IERR)
14
15 CALL MPI_COMM_RANK(MPI_COMM_WORLD, MYRANK, IERR)
16 IF (MYRANK .EQ. 0) THEN
17     CALL MPI_SEND(MPI_BOTTOM, 1, DTYPE, 1, 0, MPI_COMM_WORLD, IERR)
18 ELSE
19     CALL C_ROUTINE(DTYPE%MPI_VAL)
20 END IF
21 END SUBROUTINE

```

```

22 /* C code */

```

```

23 void C_ROUTINE(MPI_Fint *fhandle)
24 {
25     MPI_Datatype type;
26     MPI_Status status;
27
28     type = MPI_Type_f2c(*fhandle);
29
30     MPI_Recv(MPI_BOTTOM, 1, type, 0, 0, MPI_COMM_WORLD, &status);
31 }

```

MPI implementors may weaken these type matching rules, and allow messages to be sent with Fortran types and received with C types, and vice versa, when those types match. I.e., if the Fortran type `INTEGER` is identical to the C type `int`, then an MPI implementation may allow data to be sent with datatype `MPI_INTEGER` and be received with datatype `MPI_INT`. However, such code is not portable.

Appendix A

Language Bindings Summary

In this section we summarize the specific bindings for C and Fortran. First we present the constants, type definitions, info values and keys. Then we present the routine prototypes separately for each binding. Listings are alphabetical within chapter.

A.1 Defined Values and Handles

A.1.1 Defined Constants

The C and Fortran names are listed below. Constants described as “integer constant expression” may be implemented as literal integer constants of the specified integer type substituted by the preprocessor or (where possible) as enum members.

Error classes

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_SUCCESS

MPI_ERR_BUFFER

MPI_ERR_COUNT

MPI_ERR_TYPE

MPI_ERR_TAG

MPI_ERR_COMM

MPI_ERR_RANK

MPI_ERR_REQUEST

MPI_ERR_ROOT

MPI_ERR_GROUP

MPI_ERR_OP

MPI_ERR_TOPOLOGY

MPI_ERR_DIMS

MPI_ERR_ARG

MPI_ERR_UNKNOWN

MPI_ERR_TRUNCATE

MPI_ERR_OTHER

MPI_ERR_INTERN

MPI_ERR_PENDING

(Continued on next page)

	Error classes (continued)
1	
2	<hr/> C type: integer constant expression of type int
3	Fortran type: INTEGER
4	<hr/> MPI_ERR_IN_STATUS
5	MPI_ERR_ACCESS
6	MPI_ERR_AMODE
7	MPI_ERR_ASSERT
8	MPI_ERR_BAD_FILE
9	MPI_ERR_BASE
10	MPI_ERR_CONVERSION
11	MPI_ERR_DISP
12	MPI_ERR_DUP_DATAREP
13	MPI_ERR_ERRHANDLER
14	MPI_ERR_FILE_EXISTS
15	MPI_ERR_FILE_IN_USE
16	MPI_ERR_FILE
17	MPI_ERR_INFO_KEY
18	MPI_ERR_INFO_NOKEY
19	MPI_ERR_INFO_VALUE
20	MPI_ERR_INFO
21	MPI_ERR_IO
22	MPI_ERR_KEYVAL
23	MPI_ERR_LOCKTYPE
24	MPI_ERR_NAME
25	MPI_ERR_NO_MEM
26	MPI_ERR_NOT_SAME
27	MPI_ERR_NO_SPACE
28	MPI_ERR_NO_SUCH_FILE
29	MPI_ERR_PORT
30	MPI_ERR_PROC_ABORTED
31	MPI_ERR_QUOTA
32	MPI_ERR_READ_ONLY
33	MPI_ERR_RMA_ATTACH
34	MPI_ERR_RMA_CONFLICT
35	MPI_ERR_RMA_RANGE
36	MPI_ERR_RMA_SHARED
37	MPI_ERR_RMA_SYNC
38	MPI_ERR_RMA_FLAVOR
39	MPI_ERR_SERVICE
40	MPI_ERR_SESSION
41	MPI_ERR_SIZE
42	MPI_ERR_SPAWN
43	MPI_ERR_UNSUPPORTED_DATAREP
44	MPI_ERR_UNSUPPORTED_OPERATION
45	MPI_ERR_VALUE_TOO_LARGE
46	MPI_ERR_WIN

(Continued on next page)

48

Error classes (continued)

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_T_ERR_CANNOT_INIT

MPI_T_ERR_NOT_ACCESSIBLE

MPI_T_ERR_NOT_INITIALIZED

MPI_T_ERR_NOT_SUPPORTED

MPI_T_ERR_MEMORY

MPI_T_ERR_INVALID

MPI_T_ERR_INVALID_INDEX

MPI_T_ERR_INVALID_ITEM

MPI_T_ERR_INVALID_SESSION

MPI_T_ERR_INVALID_HANDLE

MPI_T_ERR_INVALID_NAME

MPI_T_ERR_OUT_OF_HANDLES

MPI_T_ERR_OUT_OF_SESSIONS

MPI_T_ERR_CVAR_SET_NOT_NOW

MPI_T_ERR_CVAR_SET_NEVER

MPI_T_ERR_PVAR_NO_WRITE

MPI_T_ERR_PVAR_NO_STARTSTOP

MPI_T_ERR_PVAR_NO_ATOMIC

MPI_ERR_LASTCODE

Buffer Address Constants

C type: void * const

Fortran type: (predefined memory location)¹

MPI_BOTTOM

MPI_BUFFER_AUTOMATIC

MPI_IN_PLACE

¹ Note that in Fortran these constants are not usable for initialization expressions or assignment. See Section 2.5.4.

Assorted Constants

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_PROC_NULL

MPI_ANY_SOURCE

MPI_ANY_TAG

MPI_UNDEFINED

MPI_BSEND_OVERHEAD

MPI_KEYVAL_INVALID

MPI_LOCK_EXCLUSIVE

MPI_LOCK_SHARED

MPI_ROOT

No Process Message Handle

C type: MPI_Message

Fortran type: INTEGER or TYPE(MPI_Message)

MPI_MESSAGE_NO_PROC

Fortran Support Method Specific Constants

Fortran type: LOGICAL

MPI_SUBARRAYS_SUPPORTED (Fortran only)

MPI_ASYNC_PROTECTS_NONBLOCKING (Fortran only)

Status array size and reserved index values (Fortran only)

Fortran type: INTEGER

MPI_STATUS_SIZE

MPI_SOURCE

MPI_TAG

MPI_ERROR

Fortran status array size and reserved index values (C only)

C type: integer constant expression of type int

MPI_F_STATUS_SIZE

MPI_F_SOURCE

MPI_F_TAG

MPI_F_ERROR

Variable Address Size (Fortran only)

Fortran type: INTEGER

MPI_ADDRESS_KIND

MPI_COUNT_KIND

MPI_INTEGER_KIND

MPI_OFFSET_KIND

Error-handling specifiers

C type: MPI_Errhandler

Fortran type: INTEGER or TYPE(MPI_Errhandler)

MPI_ERRORS_ARE_FATAL

MPI_ERRORS_ABORT

MPI_ERRORS_RETURN

Maximum Sizes for Strings

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_MAX_DATAREP_STRING

MPI_MAX_ERROR_STRING

MPI_MAX_INFO_KEY

MPI_MAX_INFO_VAL

MPI_MAX_LIBRARY_VERSION_STRING

MPI_MAX_OBJECT_NAME

MPI_MAX_PORT_NAME

MPI_MAX_PROCESSOR_NAME

MPI_MAX_STRINGTAG_LEN

MPI_MAX_PSET_NAME_LEN

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

DRAFT

Named Predefined Datatypes	C types
C type: MPI_Datatype Fortran type: INTEGER or TYPE(MPI_Datatype)	
MPI_CHAR	char (treated as printable character)
MPI_SHORT	signed short int
MPI_INT	signed int
MPI_LONG	signed long
MPI_LONG_LONG_INT	signed long long
MPI_LONG_LONG (as a synonym)	signed long long
MPI_SIGNED_CHAR	signed char (treated as integral value)
MPI_UNSIGNED_CHAR	unsigned char (treated as integral value)
MPI_UNSIGNED_SHORT	unsigned short
MPI_UNSIGNED	unsigned int
MPI_UNSIGNED_LONG	unsigned long
MPI_UNSIGNED_LONG_LONG	unsigned long long
MPI_FLOAT	float
MPI_DOUBLE	double
MPI_LONG_DOUBLE	long double
MPI_WCHAR	wchar_t (defined in <stddef.h> (treated as printable character)
MPI_C_BOOL	_Bool
MPI_INT8_T	int8_t
MPI_INT16_T	int16_t
MPI_INT32_T	int32_t
MPI_INT64_T	int64_t
MPI_UINT8_T	uint8_t
MPI_UINT16_T	uint16_t
MPI_UINT32_T	uint32_t
MPI_UINT64_T	uint64_t
MPI_AINT	MPI_Aint
MPI_COUNT	MPI_Count
MPI_OFFSET	MPI_Offset
MPI_C_COMPLEX	float _Complex
MPI_C_FLOAT_COMPLEX (as a synonym)	float _Complex
MPI_C_DOUBLE_COMPLEX	double _Complex
MPI_C_LONG_DOUBLE_COMPLEX	long double _Complex
MPI_BYTE	(any C type)
MPI_PACKED	(any C type)

Named Predefined Datatypes	Fortran types
C type: MPI_Datatype Fortran type: INTEGER or TYPE(MPI_Datatype)	
MPI_INTEGER	INTEGER
MPI_REAL	REAL
MPI_DOUBLE_PRECISION	DOUBLE PRECISION
MPI_COMPLEX	COMPLEX
MPI_LOGICAL	LOGICAL
MPI_CHARACTER	CHARACTER(1)
MPI_AINT	INTEGER(KIND=MPI_ADDRESS_KIND)
MPI_COUNT	INTEGER(KIND=MPI_COUNT_KIND)
MPI_OFFSET	INTEGER(KIND=MPI_OFFSET_KIND)
MPI_BYTE	(any Fortran type)
MPI_PACKED	(any Fortran type)

Named Predefined Datatypes ¹	C++ types
C type: MPI_Datatype Fortran type: INTEGER or TYPE(MPI_Datatype)	
MPI_CXX_BOOL	bool
MPI_CXX_FLOAT_COMPLEX	std::complex<float>
MPI_CXX_DOUBLE_COMPLEX	std::complex<double>
MPI_CXX_LONG_DOUBLE_COMPLEX	std::complex<long double>

¹ If an accompanying C++ compiler is missing, then the MPI datatypes in this table are not defined.

Optional datatypes (Fortran)	Fortran types
C type: MPI_Datatype Fortran type: INTEGER or TYPE(MPI_Datatype)	
MPI_DOUBLE_COMPLEX	DOUBLE COMPLEX
MPI_INTEGER1	INTEGER*1
MPI_INTEGER2	INTEGER*2
MPI_INTEGER4	INTEGER*4
MPI_INTEGER8	INTEGER*8
MPI_INTEGER16	INTEGER*16
MPI_REAL2	REAL*2
MPI_REAL4	REAL*4
MPI_REAL8	REAL*8
MPI_REAL16	REAL*16
MPI_COMPLEX4	COMPLEX*4
MPI_COMPLEX8	COMPLEX*8
MPI_COMPLEX16	COMPLEX*16
MPI_COMPLEX32	COMPLEX*32

Datatypes for reduction functions (C)

C type: MPI_Datatype

Fortran type: INTEGER or TYPE(MPI_Datatype)

MPI_FLOAT_INT

MPI_DOUBLE_INT

MPI_LONG_INT

MPI_2INT

MPI_SHORT_INT

MPI_LONG_DOUBLE_INT

Datatypes for reduction functions (Fortran)

C type: MPI_Datatype

Fortran type: INTEGER or TYPE(MPI_Datatype)

MPI_2REAL

MPI_2DOUBLE_PRECISION

MPI_2INTEGER

Reserved communicators

C type: MPI_Comm

Fortran type: INTEGER or TYPE(MPI_Comm)

MPI_COMM_WORLD

MPI_COMM_SELF

Communicator split type constants

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_COMM_TYPE_SHARED

MPI_COMM_TYPE_HW_UNGUIDED

MPI_COMM_TYPE_HW_GUIDED

MPI_COMM_TYPE_RESOURCE_GUIDED

Results of communicator and group comparisons

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_IDENT

MPI_CONGRUENT

MPI_SIMILAR

MPI_UNEQUAL

Environmental inquiry info key

C type: MPI_Info

Fortran type: INTEGER or TYPE(MPI_Info)

MPI_INFO_ENV

Environmental inquiry keys	1
<hr/>	
C type: integer constant expression of type int	2
Fortran type: INTEGER	3
<hr/>	
MPI_TAG_UB	4
MPI_IO	5
MPI_HOST (deprecated)	6
MPI_WTIME_IS_GLOBAL	7
<hr/>	
	8
Collective Operations	9
<hr/>	
C type: MPI_Op	10
Fortran type: INTEGER or TYPE(MPI_Op)	11
<hr/>	
MPI_MAX	12
MPI_MIN	13
MPI_SUM	14
MPI_PROD	15
MPI_MAXLOC	16
MPI_MINLOC	17
MPI_BAND	18
MPI_BOR	19
MPI_BXOR	20
MPI_LAND	21
MPI_LOR	22
MPI_LXOR	23
MPI_REPLACE	24
MPI_NO_OP	25
<hr/>	
	26
	27
	28
	29
	30
	31
	32
	33
	34
	35
	36
	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

Null Handles

C/Fortran name
C type / Fortran type
MPI_GROUP_NULL
MPI_Group / INTEGER or TYPE(MPI_Group)
MPI_COMM_NULL
MPI_Comm / INTEGER or TYPE(MPI_Comm)
MPI_DATATYPE_NULL
MPI_Datatype / INTEGER or TYPE(MPI_Datatype)
MPI_REQUEST_NULL
MPI_Request / INTEGER or TYPE(MPI_Request)
MPI_OP_NULL
MPI_Op / INTEGER or TYPE(MPI_Op)
MPI_ERRHANDLER_NULL
MPI_Errhandler / INTEGER or TYPE(MPI_Errhandler)
MPI_FILE_NULL
MPI_File / INTEGER or TYPE(MPI_File)
MPI_INFO_NULL
MPI_Info / INTEGER or TYPE(MPI_Info)
MPI_SESSION_NULL
MPI_Session / INTEGER or TYPE(MPI_Session)
MPI_WIN_NULL
MPI_Win / INTEGER or TYPE(MPI_Win)
MPI_MESSAGE_NULL
MPI_Message / INTEGER or TYPE(MPI_Message)

Empty group

C type: MPI_Group
Fortran type: INTEGER or TYPE(MPI_Group)
MPI_GROUP_EMPTY

Topologies

C type: integer constant expression of type int
Fortran type: INTEGER
MPI_GRAPH
MPI_CART
MPI_DIST_GRAPH

Predefined functions		1
C/Fortran name		2
C type		3
/ Fortran type with <code>mpi</code> module	/ Fortran type with <code>mpi_f08</code> module	4
MPI_COMM_NULL_COPY_FN		5
<code>MPI_Comm_copy_attr_function</code>		6
/ <code>COMM_COPY_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Comm_copy_attr_function)</code> ¹⁾	7
MPI_COMM_DUP_FN		8
<code>MPI_Comm_copy_attr_function</code>		9
/ <code>COMM_COPY_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Comm_copy_attr_function)</code> ¹⁾	10
MPI_COMM_NULL_DELETE_FN		11
<code>MPI_Comm_delete_attr_function</code>		12
/ <code>COMM_DELETE_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Comm_delete_attr_function)</code> ¹⁾	13
MPI_WIN_NULL_COPY_FN		14
<code>MPI_Win_copy_attr_function</code>		15
/ <code>WIN_COPY_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Win_copy_attr_function)</code> ¹⁾	16
MPI_WIN_DUP_FN		17
<code>MPI_Win_copy_attr_function</code>		18
/ <code>WIN_COPY_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Win_copy_attr_function)</code> ¹⁾	19
MPI_WIN_NULL_DELETE_FN		20
<code>MPI_Win_delete_attr_function</code>		21
/ <code>WIN_DELETE_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Win_delete_attr_function)</code> ¹⁾	22
MPI_TYPE_NULL_COPY_FN		23
<code>MPI_Type_copy_attr_function</code>		24
/ <code>TYPE_COPY_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Type_copy_attr_function)</code> ¹⁾	25
MPI_TYPE_DUP_FN		26
<code>MPI_Type_copy_attr_function</code>		27
/ <code>TYPE_COPY_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Type_copy_attr_function)</code> ¹⁾	28
MPI_TYPE_NULL_DELETE_FN		29
<code>MPI_Type_delete_attr_function</code>		30
/ <code>TYPE_DELETE_ATTR_FUNCTION</code>	/ <code>PROCEDURE(MPI_Type_delete_attr_function)</code> ¹⁾	31
MPI_CONVERSION_FN_NULL		32
<code>MPI_Datarep_conversion_function</code>		33
/ <code>DATAREP_CONVERSION_FUNCTION</code>	/ <code>PROCEDURE(MPI_Datarep_conversion_function)</code> ¹⁾	34
MPI_CONVERSION_FN_NULL_C		35
<code>MPI_Datarep_conversion_function_c</code>		36
/ (n/a)	/ <code>PROCEDURE(MPI_Datarep_conversion_function_c)</code>	37

¹ See the advice to implementors (on page 363) and advice to users (on page 364) on the predefined Fortran functions `MPI_COMM_NULL_COPY_FN`, ... in Section 7.7.2.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

Deprecated predefined functions

C/Fortran name
C type / Fortran type with mpi module
MPI_NULL_COPY_FN
MPI_Copy_function / COPY_FUNCTION
MPI_DUP_FN
MPI_Copy_function / COPY_FUNCTION
MPI_NULL_DELETE_FN
MPI_Delete_function / DELETE_FUNCTION

Predefined Attribute Keys

C type: integer constant expression of type int
Fortran type: INTEGER
MPI_APPNUM
MPI_LASTUSED CODE
MPI_UNIVERSE_SIZE
MPI_WIN_BASE
MPI_WIN_DISP_UNIT
MPI_WIN_SIZE
MPI_WIN_CREATE_FLAVOR
MPI_WIN_MODEL

MPI Window Create Flavors

C type: integer constant expression of type int
Fortran type: INTEGER
MPI_WIN_FLAVOR_CREATE
MPI_WIN_FLAVOR_ALLOCATE
MPI_WIN_FLAVOR_DYNAMIC
MPI_WIN_FLAVOR_SHARED

MPI Window Models

C type: integer constant expression of type int
Fortran type: INTEGER
MPI_WIN_SEPARATE
MPI_WIN_UNIFIED

Mode Constants	1
<hr/>	
C type: integer constant expression of type int	2
Fortran type: INTEGER	3
<hr/>	
MPI_MODE_APPEND	4
MPI_MODE_CREATE	5
MPI_MODE_DELETE_ON_CLOSE	6
MPI_MODE_EXCL	7
MPI_MODE_NOCHECK	8
MPI_MODE_NOPRECEDE	9
MPI_MODE_NOPUT	10
MPI_MODE_NOSTORE	11
MPI_MODE_NOSUCCEED	12
MPI_MODE_RDONLY	13
MPI_MODE_RDWR	14
MPI_MODE_SEQUENTIAL	15
MPI_MODE_UNIQUE_OPEN	16
MPI_MODE_WRONLY	17
<hr/>	
Datatype Decoding Constants	18
<hr/>	
C type: integer constant expression of type int	19
Fortran type: INTEGER	20
<hr/>	
MPI_COMBINER_CONTIGUOUS	21
MPI_COMBINER_DARRAY	22
MPI_COMBINER_DUP	23
MPI_COMBINER_F90_COMPLEX	24
MPI_COMBINER_F90_INTEGER	25
MPI_COMBINER_F90_REAL	26
MPI_COMBINER_HINDEXED	27
MPI_COMBINER_HVECTOR	28
MPI_COMBINER_INDEXED_BLOCK	29
MPI_COMBINER_HINDEXED_BLOCK	30
MPI_COMBINER_INDEXED	31
MPI_COMBINER_NAMED	32
MPI_COMBINER_RESIZED	33
MPI_COMBINER_STRUCT	34
MPI_COMBINER_SUBARRAY	35
MPI_COMBINER_VECTOR	36
<hr/>	
Threads Constants	37
<hr/>	
C type: integer constant expression of type int	38
Fortran type: INTEGER	39
<hr/>	
MPI_THREAD_FUNNELED	40
MPI_THREAD_MULTIPLE	41
MPI_THREAD_SERIALIZED	42
MPI_THREAD_SINGLE	43
<hr/>	
	44
	45
	46
	47
	48

File Operation Constants, Part 1

C type: integer constant expression of type MPI_Offset

Fortran type: INTEGER(KIND=MPI_OFFSET_KIND)

MPI_DISPLACEMENT_CURRENT

File Operation Constants, Part 2

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_DISTRIBUTE_BLOCK

MPI_DISTRIBUTE_CYCLIC

MPI_DISTRIBUTE_DFLT_DARG

MPI_DISTRIBUTE_NONE

MPI_ORDER_C

MPI_ORDER_FORTRAN

MPI_SEEK_CUR

MPI_SEEK_END

MPI_SEEK_SET

F90 Datatype Matching Constants

C type: integer constant expression of type int

Fortran type: INTEGER

MPI_TYPECLASS_COMPLEX

MPI_TYPECLASS_INTEGER

MPI_TYPECLASS_REAL

Constants Specifying Empty or Ignored Input

C/Fortran name

C type / Fortran type¹

MPI_ARGVS_NULL

char*** / 2-dim. array of CHARACTER*(*)

MPI_ARGV_NULL

char** / array of CHARACTER*(*)

MPI_ERRCODES_IGNORE

int* / INTEGER array

MPI_STATUSES_IGNORE

MPI_Status* / INTEGER, DIMENSION(MPI_STATUS_SIZE,*)

or TYPE(MPI_Status), DIMENSION(*)

MPI_STATUS_IGNORE

MPI_Status* / INTEGER, DIMENSION(MPI_STATUS_SIZE)

or TYPE(MPI_Status)

MPI_UNWEIGHTED

int* / INTEGER array

MPI_WEIGHTS_EMPTY

int* / INTEGER array

¹ Note that in Fortran these constants are not usable for initialization expressions or assignment. See Section 2.5.4.

C Constants Specifying Ignored Input (no Fortran)

C constant (type: MPI_Fint*)	is equivalent to the Fortran constant	1
MPI_F_STATUSES_IGNORE	MPI_STATUSES_IGNORE in mpi / mpif.h	2
MPI_F_STATUS_IGNORE	MPI_STATUS_IGNORE in mpi / mpif.h	3
C constant (type: MPI_F08_status*)	is equivalent to the Fortran constant	4
MPI_F08_STATUSES_IGNORE	MPI_STATUSES_IGNORE in mpi_f08	5
MPI_F08_STATUS_IGNORE	MPI_STATUS_IGNORE in mpi_f08	6

C preprocessor Constants and Fortran Parameters

C type: C-preprocessor macro that expands to an int value

Fortran type: INTEGER

MPI_SUBVERSION

MPI_VERSION

Null handles used in the MPI tool information interface

MPI_T_ENUM_NULL

MPI_T_enum

MPI_T_CVAR_HANDLE_NULL

MPI_T_cvar_handle

MPI_T_PVAR_HANDLE_NULL

MPI_T_pvar_handle

MPI_T_PVAR_SESSION_NULL

MPI_T_pvar_session

Verbosity Levels in the MPI tool information interface

C type: integer constant expression of type int

MPI_T_VERBOSITY_USER_BASIC

MPI_T_VERBOSITY_USER_DETAIL

MPI_T_VERBOSITY_USER_ALL

MPI_T_VERBOSITY_TUNER_BASIC

MPI_T_VERBOSITY_TUNER_DETAIL

MPI_T_VERBOSITY_TUNER_ALL

MPI_T_VERBOSITY_MPIDEV_BASIC

MPI_T_VERBOSITY_MPIDEV_DETAIL

MPI_T_VERBOSITY_MPIDEV_ALL

1 **Constants to identify associations of variables**
 2 **in the MPI tool information interface**

3 C type: integer constant expression of type int

4 MPI_T_BIND_NO_OBJECT
 5 MPI_T_BIND_MPI_COMM
 6 MPI_T_BIND_MPI_DATATYPE
 7 MPI_T_BIND_MPI_ERRHANDLER
 8 MPI_T_BIND_MPI_FILE
 9 MPI_T_BIND_MPI_GROUP
 10 MPI_T_BIND_MPI_OP
 11 MPI_T_BIND_MPI_REQUEST
 12 MPI_T_BIND_MPI_WIN
 13 MPI_T_BIND_MPI_MESSAGE
 14 MPI_T_BIND_MPI_INFO
 15 MPI_T_BIND_MPI_SESSION

16
 17 **Constants describing the scope of a control variable**
 18 **in the MPI tool information interface**

19 C type: integer constant expression of type int

20 MPI_T_SCOPE_CONSTANT
 21 MPI_T_SCOPE_READONLY
 22 MPI_T_SCOPE_LOCAL
 23 MPI_T_SCOPE_GROUP
 24 MPI_T_SCOPE_GROUP_EQ
 25 MPI_T_SCOPE_ALL
 26 MPI_T_SCOPE_ALL_EQ

27
 28 **Additional constants used**
 29 **by the MPI tool information interface**

30 C type: MPI_T_pvar_handle

31 MPI_T_PVAR_ALL_HANDLES

32
 33 **Performance variables classes used by the**
 34 **MPI tool information interface**

35 C type: integer constant expression of type int

36 MPI_T_PVAR_CLASS_STATE
 37 MPI_T_PVAR_CLASS_LEVEL
 38 MPI_T_PVAR_CLASS_SIZE
 39 MPI_T_PVAR_CLASS_PERCENTAGE
 40 MPI_T_PVAR_CLASS_HIGHWATERMARK
 41 MPI_T_PVAR_CLASS_LOWWATERMARK
 42 MPI_T_PVAR_CLASS_COUNTER
 43 MPI_T_PVAR_CLASS_AGGREGATE
 44 MPI_T_PVAR_CLASS_TIMER
 45 MPI_T_PVAR_CLASS_GENERIC

**Source event ordering guarantees in the
MPI tool information interface**

C type: MPI_T_source_order

MPI_T_SOURCE_ORDERED

MPI_T_SOURCE_UNORDERED

**Callback safety requirement levels used in the
MPI tool information interface**

C type: MPI_T_cb_safety

MPI_T_CB_REQUIRE_NONE

MPI_T_CB_REQUIRE_MPI_RESTRICTED

MPI_T_CB_REQUIRE_THREAD_SAFE

MPI_T_CB_REQUIRE_ASYNC_SIGNAL_SAFE

A.1.2 Types

The following are defined C type definitions included in the file `mpi.h`.

```
/* C opaque types */
```

```
MPI_Aint
```

```
MPI_Count
```

```
MPI_Fint
```

```
MPI_Offset
```

```
MPI_Status
```

```
MPI_F08_status
```

```
/* C handles to assorted structures */
```

```
MPI_Comm
```

```
MPI_Datatype
```

```
MPI_Errhandler
```

```
MPI_File
```

```
MPI_Group
```

```
MPI_Info
```

```
MPI_Message
```

```
MPI_Op
```

```
MPI_Request
```

```
MPI_Session
```

```
MPI_Win
```

```
/* Types for the MPI_T interface */
```

```
MPI_T_enum
```

```
MPI_T_cvar_handle
```

```
MPI_T_pvar_handle
```

```
MPI_T_pvar_session
```

```
MPI_T_event_instance
```

```
MPI_T_event_registration
```

```
MPI_T_source_order
```

```
MPI_T_cb_safety
```

The following are defined Fortran type definitions included in the `mpi_f08` and `mpi` modules.

```

5 ! Fortran opaque types in the mpi_f08 and mpi modules
6 TYPE(MPI_Status)
7
8 ! Fortran handles in the mpi_f08 and mpi modules
9 TYPE(MPI_Comm)
10 TYPE(MPI_Datatype)
11 TYPE(MPI_Errhandler)
12 TYPE(MPI_File)
13 TYPE(MPI_Group)
14 TYPE(MPI_Info)
15 TYPE(MPI_Message)
16 TYPE(MPI_Op)
17 TYPE(MPI_Request)
18 TYPE(MPI_Session)
19 TYPE(MPI_Win)
20

```

A.1.3 Prototype Definitions

C Bindings

The following are defined C typedefs for user-defined functions, also included in the file `mpi.h`.

```

27 /* prototypes for user-defined functions */
28
29 typedef void MPI_User_function(void *invec, void *inoutvec, int *len,
30                               MPI_Datatype *datatype);
31
32 typedef void MPI_User_function_c(void *invec, void *inoutvec, MPI_Count *len,
33                                 MPI_Datatype *datatype);
34
35 typedef int MPI_Comm_copy_attr_function(MPI_Comm oldcomm, int comm_keyval,
36                                       void *extra_state, void *attribute_val_in,
37                                       void *attribute_val_out, int *flag);
38
39 typedef int MPI_Comm_delete_attr_function(MPI_Comm comm, int comm_keyval,
40                                         void *attribute_val, void *extra_state);
41
42 typedef int MPI_Win_copy_attr_function(MPI_Win oldwin, int win_keyval,
43                                       void *extra_state, void *attribute_val_in,
44                                       void *attribute_val_out, int *flag);
45
46 typedef int MPI_Win_delete_attr_function(MPI_Win win, int win_keyval,
47                                         void *attribute_val, void *extra_state);
48
49 typedef int MPI_Type_copy_attr_function(MPI_Datatype oldtype, int type_keyval,
50                                       void *extra_state, void *attribute_val_in,
51                                       void *attribute_val_out, int *flag);

```

```

typedef int MPI_Type_delete_attr_function(MPI_Datatype datatype,      1
                                         int type_keyval, void *attribute_val, void *extra_state);  2
typedef void MPI_Comm_errhandler_function(MPI_Comm *comm, int *error_code,  3
                                         . . . );  4
typedef void MPI_Win_errhandler_function(MPI_Win *win, int *error_code, . . . );  5
typedef void MPI_File_errhandler_function(MPI_File *file, int *error_code,  6
                                         . . . );  7
typedef void MPI_Session_errhandler_function(MPI_Session *session,  8
                                             int *error_code, . . . );  9
typedef int MPI_Grequest_query_function(void *extra_state, MPI_Status *status);  10
typedef int MPI_Grequest_free_function(void *extra_state);  11
typedef int MPI_Grequest_cancel_function(void *extra_state, int complete);  12
typedef int MPI_Datarep_extent_function(MPI_Datatype datatype,  13
                                         MPI_Aint *extent, void *extra_state);  14
typedef int MPI_Datarep_conversion_function(void *userbuf,  15
                                             MPI_Datatype datatype, int count, void *filebuf,  16
                                             MPI_Offset position, void *extra_state);  17
typedef int MPI_Datarep_conversion_function_c(void *userbuf,  18
                                             MPI_Datatype datatype, MPI_Count count, void *filebuf,  19
                                             MPI_Offset position, void *extra_state);  20
typedef void MPI_T_event_cb_function(MPI_T_event_instance event_instance,  21
                                     MPI_T_event_registration event_registration,  22
                                     MPI_T_cb_safety cb_safety, void *user_data);  23
typedef void MPI_T_event_free_cb_function(  24
    MPI_T_event_registration event_registration,  25
    MPI_T_cb_safety cb_safety, void *user_data);  26
typedef void MPI_T_event_dropped_cb_function(MPI_Count count,  27
    MPI_T_event_registration event_registration, int source_index,  28
    MPI_T_cb_safety cb_safety, void *user_data);  29

```

Fortran 2008 Bindings with the mpi_f08 Module

The callback prototypes when using the Fortran `mpi_f08` module are shown below:

The user-function argument to `MPI_Op_create` and `MPI_Op_create_c` should be declared according to:

```

ABSTRACT INTERFACE  37
  SUBROUTINE MPI_User_function(invec, inoutvec, len, datatype)  38
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR  39
    TYPE(C_PTR), VALUE :: invec, inoutvec  40
    INTEGER :: len  41
    TYPE(MPI_Datatype) :: datatype  42
  ABSTRACT INTERFACE  43

```

```

1  SUBROUTINE MPI_User_function_c(invec, inoutvec, len, datatype) !(_c)
2  USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
3  TYPE(C_PTR), VALUE :: invec, inoutvec
4  INTEGER(KIND=MPI_COUNT_KIND) :: len
5  TYPE(MPI_Datatype) :: datatype

```

The copy and delete function arguments to MPI_Comm_create_keyval should be declared according to:

```

6  ABSTRACT INTERFACE
7
8  SUBROUTINE MPI_Comm_copy_attr_function(oldcomm, comm_keyval, extra_state,
9  attribute_val_in, attribute_val_out, flag, ierror)
10
11  TYPE(MPI_Comm) :: oldcomm
12  INTEGER :: comm_keyval, ierror
13  INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
14  attribute_val_out
15  LOGICAL :: flag
16

```

```

17 ABSTRACT INTERFACE
18 SUBROUTINE MPI_Comm_delete_attr_function(comm, comm_keyval, attribute_val,
19 extra_state, ierror)
20 TYPE(MPI_Comm) :: comm
21 INTEGER :: comm_keyval, ierror
22 INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
23

```

The copy and delete function arguments to MPI_Win_create_keyval should be declared according to:

```

24 ABSTRACT INTERFACE
25
26 SUBROUTINE MPI_Win_copy_attr_function(oldwin, win_keyval, extra_state,
27 attribute_val_in, attribute_val_out, flag, ierror)
28
29 TYPE(MPI_Win) :: oldwin
30 INTEGER :: win_keyval, ierror
31 INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
32 attribute_val_out
33 LOGICAL :: flag
34

```

```

35 ABSTRACT INTERFACE
36 SUBROUTINE MPI_Win_delete_attr_function(win, win_keyval, attribute_val,
37 extra_state, ierror)
38 TYPE(MPI_Win) :: win
39 INTEGER :: win_keyval, ierror
40 INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
41

```

The copy and delete function arguments to MPI_Type_create_keyval should be declared according to:

```

42 ABSTRACT INTERFACE
43
44 SUBROUTINE MPI_Type_copy_attr_function(oldtype, type_keyval, extra_state,
45 attribute_val_in, attribute_val_out, flag, ierror)
46
47 TYPE(MPI_Datatype) :: oldtype
48 INTEGER :: type_keyval, ierror
49 INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
50 attribute_val_out

```

```

LOGICAL :: flag
1
ABSTRACT INTERFACE
2
  SUBROUTINE MPI_Type_delete_attr_function(datatype, type_keyval,
3
    attribute_val, extra_state, ierror)
4
    TYPE(MPI_Datatype) :: datatype
5
    INTEGER :: type_keyval, ierror
6
    INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
7
8
  The handler-function argument to MPI_Comm_create_errhandler should be declared like
9
  this:
10
ABSTRACT INTERFACE
11
  SUBROUTINE MPI_Comm_errhandler_function(comm, error_code)
12
    TYPE(MPI_Comm) :: comm
13
    INTEGER :: error_code
14
15
  The handler-function argument to MPI_Win_create_errhandler should be declared like
16
  this:
17
ABSTRACT INTERFACE
18
  SUBROUTINE MPI_Win_errhandler_function(win, error_code)
19
    TYPE(MPI_Win) :: win
20
    INTEGER :: error_code
21
22
  The handler-function argument to MPI_File_create_errhandler should be declared like
23
  this:
24
ABSTRACT INTERFACE
25
  SUBROUTINE MPI_File_errhandler_function(file, error_code)
26
    TYPE(MPI_File) :: file
27
    INTEGER :: error_code
28
29
  The handler-function argument to MPI_Session_create_errhandler should be declared
30
  like this:
31
ABSTRACT INTERFACE
32
  SUBROUTINE MPI_Session_errhandler_function(session, error_code)
33
    TYPE(MPI_Session) :: session
34
    INTEGER :: error_code
35
36
  The query, free, and cancel function arguments to MPI_Grequest_start should be de-
37
  clared according to:
38
ABSTRACT INTERFACE
39
  SUBROUTINE MPI_Grequest_query_function(extra_state, status, ierror)
40
    INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
41
    TYPE(MPI_Status) :: status
42
    INTEGER :: ierror
43
44
ABSTRACT INTERFACE
45
  SUBROUTINE MPI_Grequest_free_function(extra_state, ierror)
46
    INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
47
    INTEGER :: ierror
48
49
ABSTRACT INTERFACE
50
  SUBROUTINE MPI_Grequest_cancel_function(extra_state, complete, ierror)
51

```

```

1     INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
2     LOGICAL :: complete
3     INTEGER :: ierror
4
5     The extent and conversion function arguments to MPI_Register_datarep and
6     MPI_Register_datarep_c should be declared according to:
7     ABSTRACT INTERFACE
8         SUBROUTINE MPI_Datarep_extent_function(datatype, extent, extra_state, ierror)
9             TYPE(MPI_Datatype) :: datatype
10            INTEGER(KIND=MPI_ADDRESS_KIND) :: extent, extra_state
11            INTEGER :: ierror
12
13    ABSTRACT INTERFACE
14        SUBROUTINE MPI_Datarep_conversion_function(userbuf, datatype, count, filebuf,
15            position, extra_state, ierror)
16            USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
17            TYPE(C_PTR), VALUE :: userbuf, filebuf
18            TYPE(MPI_Datatype) :: datatype
19            INTEGER :: count, ierror
20            INTEGER(KIND=MPI_OFFSET_KIND) :: position
21            INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
22
23    ABSTRACT INTERFACE
24        SUBROUTINE MPI_Datarep_conversion_function_c(userbuf, datatype, count,
25            filebuf, position, extra_state, ierror) !(_c)
26            USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
27            TYPE(C_PTR), VALUE :: userbuf, filebuf
28            TYPE(MPI_Datatype) :: datatype
29            INTEGER(KIND=MPI_COUNT_KIND) :: count
30            INTEGER(KIND=MPI_OFFSET_KIND) :: position
31            INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
32            INTEGER :: ierror

```

Fortran Bindings with mpif.h or the mpi Module

With the Fortran `mpi` module or (deprecated) `mpif.h`, here are examples of how each of the user-defined subroutines should be declared.

The user-function argument to `MPI_OP_CREATE` should be declared like this:

```

37    SUBROUTINE USER_FUNCTION(INVEC, INOUTVEC, LEN, DATATYPE)
38        <type> INVEC(LEN), INOUTVEC(LEN)
39        INTEGER LEN, DATATYPE

```

The copy and delete function arguments to `MPI_COMM_CREATE_KEYVAL` should be declared like these:

```

43    SUBROUTINE COMM_COPY_ATTR_FUNCTION(OLDCOMM, COMM_KEYVAL, EXTRA_STATE,
44        ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)
45        INTEGER OLDCOMM, COMM_KEYVAL, IERROR
46        INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
47        ATTRIBUTE_VAL_OUT
48        LOGICAL FLAG

```



```

SUBROUTINE COMM_DELETE_ATTR_FUNCTION(COMM, COMM_KEYVAL, ATTRIBUTE_VAL,
                                     EXTRA_STATE, IERROR)
    INTEGER COMM, COMM_KEYVAL, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE

```

The copy and delete function arguments to MPI_WIN_CREATE_KEYVAL should be declared like these:

```

SUBROUTINE WIN_COPY_ATTR_FUNCTION(OLDWIN, WIN_KEYVAL, EXTRA_STATE,
                                  ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)
    INTEGER OLDWIN, WIN_KEYVAL, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
    ATTRIBUTE_VAL_OUT
    LOGICAL FLAG

```

```

SUBROUTINE WIN_DELETE_ATTR_FUNCTION(WIN, WIN_KEYVAL, ATTRIBUTE_VAL,
                                    EXTRA_STATE, IERROR)
    INTEGER WIN, WIN_KEYVAL, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE

```

The copy and delete function arguments to MPI_TYPE_CREATE_KEYVAL should be declared like these:

```

SUBROUTINE TYPE_COPY_ATTR_FUNCTION(OLDTYPE, TYPE_KEYVAL, EXTRA_STATE,
                                   ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)
    INTEGER OLDTYPE, TYPE_KEYVAL, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
    ATTRIBUTE_VAL_OUT
    LOGICAL FLAG

```

```

SUBROUTINE TYPE_DELETE_ATTR_FUNCTION(DATATYPE, TYPE_KEYVAL, ATTRIBUTE_VAL,
                                     EXTRA_STATE, IERROR)
    INTEGER DATATYPE, TYPE_KEYVAL, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE

```

The handler-function argument to MPI_COMM_CREATE_ERRHANDLER should be declared like this:

```

SUBROUTINE COMM_ERRHANDLER_FUNCTION(COMM, ERROR_CODE)
    INTEGER COMM, ERROR_CODE

```

The handler-function argument to MPI_WIN_CREATE_ERRHANDLER should be declared like this:

```

SUBROUTINE WIN_ERRHANDLER_FUNCTION(WIN, ERROR_CODE)
    INTEGER WIN, ERROR_CODE

```

The handler-function argument to MPI_FILE_CREATE_ERRHANDLER should be declared like this:

```

SUBROUTINE FILE_ERRHANDLER_FUNCTION(FILE, ERROR_CODE)
    INTEGER FILE, ERROR_CODE

```

The handler-function argument to MPI_SESSION_CREATE_ERRHANDLER should be declared like this:

```

SUBROUTINE SESSION_ERRHANDLER_FUNCTION(SESSION, ERROR_CODE)
    INTEGER SESSION, ERROR_CODE

```

1 The query, free, and cancel function arguments to MPI_GREQUEST_START should be
2 declared like these:

```
3 SUBROUTINE GREQUEST_QUERY_FUNCTION(EXTRA_STATE, STATUS, IERROR)
4     INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
5     INTEGER STATUS(MPI_STATUS_SIZE), IERROR
6
7 SUBROUTINE GREQUEST_FREE_FUNCTION(EXTRA_STATE, IERROR)
8     INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
9     INTEGER IERROR
10
11 SUBROUTINE GREQUEST_CANCEL_FUNCTION(EXTRA_STATE, COMPLETE, IERROR)
12     INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
13     LOGICAL COMPLETE
14     INTEGER IERROR
```

14 The extent and conversion function arguments to MPI_REGISTER_DATAREP should
15 be declared like these:

```
16 SUBROUTINE DATAREP_EXTENT_FUNCTION(DATATYPE, EXTENT, EXTRA_STATE, IERROR)
17     INTEGER DATATYPE, IERROR
18     INTEGER(KIND=MPI_ADDRESS_KIND) EXTENT, EXTRA_STATE
19
20 SUBROUTINE DATAREP_CONVERSION_FUNCTION(USERBUF, DATATYPE, COUNT, FILEBUF,
21     POSITION, EXTRA_STATE, IERROR)
22     <TYPE> USERBUF(*), FILEBUF(*)
23     INTEGER DATATYPE, COUNT, IERROR
24     INTEGER(KIND=MPI_OFFSET_KIND) POSITION
25     INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
```

27 A.1.4 Deprecated Prototype Definitions

28 The following are defined C typedefs for deprecated user-defined functions, also included in
29 the file `mpi.h`.

```
31 /* prototypes for user-defined functions */
32
33 typedef int MPI_Copy_function(MPI_Comm oldcomm, int keyval, void *extra_state,
34     void *attribute_val_in, void *attribute_val_out, int *flag);
35
36 typedef int MPI_Delete_function(MPI_Comm comm, int keyval, void *attribute_val,
37     void *extra_state);
```

38 The following are deprecated Fortran user-defined callback subroutine prototypes. The
39 deprecated copy and delete function arguments to MPI_KEYVAL_CREATE should be de-
40 clared like these:

```
41 SUBROUTINE COPY_FUNCTION(OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
42     ATTRIBUTE_VAL_OUT, FLAG, IERR)
43     INTEGER OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT,
44     IERR
45     LOGICAL FLAG
46
47 SUBROUTINE DELETE_FUNCTION(COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR)
48     INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR
```

A.1.5 String Values

Default Communicator Names

The following default communicator names are defined by MPI.

"MPI_COMM_WORLD"
 "MPI_COMM_SELF"
 "MPI_COMM_PARENT"
 "MPI_COMM_NULL"

Default Datatype Names

Named predefined datatypes have the default names of the datatype name. In addition, the following default datatype name is defined by MPI.

"MPI_DATATYPE_NULL"

Default Window Names

The following default window name is defined by MPI.

"MPI_WIN_NULL"

Reserved Data Representations

The following data representations are supported by MPI.

"native"
 "internal"
 "external32"

Process Set Names

Process set name	Comment
"mpi://"	reserved namespace
"mpi://SELF"	mandatory process set name
"mpi://WORLD"	mandatory process set name

Info Keys

The following info keys are reserved. They are strings.

"access_style"
 "accumulate_ops"
 "accumulate_ordering"
 "alloc_shared_noncontig"
 "appnum"
 "arch"
 "argv"
 "cb_block_size"
 "cb_buffer_size"
 "cb_nodes"
 "chunked_item"

1 "chunked_size"
2 "chunked"
3 "collective_buffering"
4 "command"
5 "file"
6 "file_perm"
7 "filename"
8 "host"
9 "io_node_list"
10 "ip_address"
11 "ip_port"
12 "maxprocs"
13 "mpi_accumulate_granularity"
14 "mpi_assert_allow_overtaking"
15 "mpi_assert_exact_length"
16 "mpi_assert_memory_alloc_kinds"
17 "mpi_assert_no_any_source"
18 "mpi_assert_no_any_tag"
19 "mpi_hw_resource_type"
20 "mpi_initial_errhandler"
21 "mpi_memory_alloc_kinds"
22 "mpi_minimum_memory_alignment"
23 "mpi_pset_name"
24 "mpi_size"
25 "nb_proc"
26 "no_locks"
27 "num_io_nodes"
28 "path"
29 "same_disp_unit"
30 "same_size"
31 "soft"
32 "striping_factor"
33 "striping_unit"
34 "thread_level"
35 "wdir"

37 *Info Values*

38
39 The following info values are reserved. They are strings.

40 "alloc_mem"
41 "false"
42 "mpi"
43 "mpi_errors_abort"
44 "mpi_errors_are_fatal"
45 "mpi_errors_return"
46 "mpi_shared_memory"
47 "MPI_THREAD_FUNNELED"
48 "MPI_THREAD_MULTIPLE"

"MPI_THREAD_SERIALIZED"	1
"MPI_THREAD_SINGLE"	2
"none"	3
"random"	4
"rar"	5
"raw"	6
"read_mostly"	7
"read_once"	8
"reverse_sequential"	9
"same_op"	10
"same_op_no_op"	11
"sequential"	12
"system"	13
"true"	14
"war"	15
"waw"	16
"win_allocate"	17
"win_allocate_shared"	18
"write_mostly"	19
"write_once"	20

A.2 Summary of the Semantics of all Operation-Related MPI Procedures

This annex provides the list of MPI procedures that are associated with an MPI operation, or inquiry procedures providing information about an operation.

In many cases, the MPI procedures and their properties are listed under certain constraints, e.g., a call to `MPI_WAIT` that completes either a nonblocking or a persistent operation, or RMA calls in combination with various synchronization methods.

Table Legend:

- **Stages:** i=initialization, s=starting, c=completion, f=freeing. The procedure does at least part of the indicated stage(s).
- **Cpl:** ic=incomplete procedure, c=completing procedure, f=freeing procedure
- **Loc:** l=local procedure, nl=non-local procedure
- *: exceptions, e.g., ic+nl = incomplete+non-local, and c+l = completing+local (both are defined as blocking)
- **Blk:** b=blocking procedure, nb=nonblocking procedure. Note that from a user's view point, this column is only a hint. Relevant is, whether a routine is local or not and which resources are blocked until when. See both previous and last columns.
- ‡: exceptions, e.g., nonblocking procedures without prefix l, or that prefix l only marks immediate return.
- **Op:** part of operation type: b-op = blocking operation, nb-op = nonblocking operation, p-op = persistent operation, pp-op = persistent partitioned operation
- **Collective procedures:**

- 1 – C = all processes of the group must call the procedure
- 2 – sq = in the same sequence
- 3 – S1 = blocking synchronization, i.e., no process shall return from this procedure until all
- 4 processes on the associated process group called this procedure
- 5 – W1 = the implementation is permitted to do S1 but not required to do S1
- 6 – S2 = start-complete-synchronization, i.e., no process shall complete the associated op-
- 7 eration until all processes on the associated process group have called the associated
- 8 starting procedure
- 9 – W2 = the implementation is permitted to do S2 but not required to do S2
- 10 • **Blocked resources:** They are blocked after the call until the end of the subsequent stage
- 11 where this resource is not mentioned further in the table.

13 Table Remarks:

- 15 1. Must not return before the corresponding MPI receive operation is started.
- 16 2. Not related to an MPI operation. Prior to MPI-4.0, MPI_PROBE and MPI_IProbe were also
- 17 described as blocking and nonblocking. From MPI-4.0 onwards, only non-local and local are
- 18 used to describe these procedures.
- 19 3. Usually, MPI_WAIT is non-local, but in this case it is local.
- 20 4. In case of a MPI_(I)BARRIER on an intra-communicator, the S1/S2 synchronization is required
- 21 (instead of W1/W2).
- 22 5. Collective: all processes must complete, but with the free choice of using MPI_WAIT or
- 23 MPI_TEST returning flag = TRUE.
- 24 6. It also may not return until MPI_INIT was called in the children.
- 25 7. Addresses are cached on the request handle.
- 26 8. One of the rare cases that an incomplete call is non-local and therefore blocking.
- 27 9. One shall not free or deallocate the buffer before the operation is freed, that is
- 28 MPI_REQUEST_FREE returned.
- 29 10. For MPI_WAIT and MPI_TEST, see corresponding lines for a) MPI_BSEND, or b)
- 30 MPI_IBCAST.
- 31 11. The prefix l marks only that this procedure returns immediately.
- 32 12. One of the exceptions that a completing and therefore blocking operation-related procedure
- 33 is local.
- 34 13. MPI_(I)MProbe initializes the operation through generating the message handle whereas
- 35 MPI_(I)MRECV initializes the receive buffer (i.e., two MPI procedures together implement the
- 36 initialization stage).
- 37 14. Nonblocking procedure without an l prefix.
- 38 15. Initialization stage (“i”) only if flag = TRUE is returned else no operation is progressed.
- 39 16. Collective: all processes must start, but with the free choice of using MPI_START or
- 40 MPI_STARTALL for a given persistent request handle (i.e., if one process starts a persistent
- 41 request handle then all processes of the associated process group must start their corresponding
- 42 request handle, and if any process starts then all processes must complete their handles).
- 43 17. In a correct MPI program, a call to MPI_(I)RSEND requires that the receiver has already
- 44 started the corresponding receive. Under this assumption, the call to MPI_RSEND and the
- 45 call to MPI_WAIT with an (active) ready send request handle are local.
- 46
- 47
- 48

18. Based on their semantics, when called using an intra-communicator, `MPI_ALLGATHER`, `MPI_ALLTOALL`, and their V and W variants, `MPI_ALLREDUCE`, `MPI_REDUCE_SCATTER`, and `MPI_REDUCE_SCATTER_BLOCK` must synchronize (i.e., S1/S2 instead of W1/W2) provided that all counts and the size of all datatypes are larger than zero. 1
19. `MPI_COMM_FREE` may return before any pending communication has finished and the communicator is deallocated. In contrast, `MPI_COMM_DISCONNECT` waits for pending communication to finish and deallocates the communicator before it returns. 2
20. The request handle is in the “active” state after `MPI_START`, i.e., `MPI_REQUEST_FREE` is now forbidden. But the starting stage is not yet finished, and the contents of the buffer are not yet “blocked.” An additional `MPI_PREADY` and variants `MPI_PREADY_RANGE` `MPI_PREADY_LIST` are required to activate each partition of the send buffer to finish the starting stage. 3
21. As part of the completion stage, the user is allowed to read part of the output buffer after returning from `MPI_PARRIVED` with `flag = TRUE` before completing the whole operation with a `MPI_WAIT`/`MPI_TEST` procedure. 4
22. It initializes the attached buffer as completely free. 5
23. It uses the attached buffer and performs all four stages on the send buffer. It occupies the needed part of the attached buffer. 6
24. It waits until the attached buffer is empty, i.e., all messages have been transmitted, and then releases the attached buffer. 7
25. Although in case of `flag = TRUE` the operation is completed, a subsequent call to test, wait, or free must be executed for deallocating or inactivating the request handle as final part of the stages c and f. It is listed only in this scenario, but can be used everywhere, where `MPI_TEST` can be called. 8
26. It frees the request handle. If the related communication operation is still ongoing then the completion and freeing stage can take place after the procedure returned. 9
27. Cancelling a send request is deprecated. 10
28. Can also be applied to activ persistent requests. 11
29. As an exception, `MPI_WAIT` is local and `MPI_TEST` repeatedly called will eventually return `flag=true`. The cancelled send or receive operation is completed and the buffer can be reused. Whether the message is sent out from the buffer or received in the buffer, this part of the completion stage is only executed if a subsequent `MPI_TEST_CANCELLED` for the returned status would return `flag = FALSE`. The freeing stage will be performed only for non-persistent requests. 12
30. In some cases, more than one MPI procedure may be needed to implement one stage of an MPI one-sided operation. For details on the semantics of one-sided operations, see Chapter 12. 13
31. Local completion only (at origin). 14
32. Local completion only (at target). 15
33. Completion at target and locally at origin. 16
34. Return from `MPI_WIN_START` and these subsequent procedures at the origin process may be delayed until `MPI_WIN_POST` has been called at the target process (see Section 12.5 and Example 12.4). 17
35. Return from `MPI_WIN_LOCK` and these subsequent procedures may be delayed until other origin processes have released their lock (see Section 12.5 and Example 12.5). 18
36. The init and freeing stages and the buffer address of the target window only apply to MPI processes in the role of a target of an RMA operation. 19

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

- 37. The freeing stage applies to operations only and does not apply to any request.
- 38. The same procedure call may serve different stages for different operations, i.e., the completion of a previous RMA and/or exposure epoch and/or the start of a next RMA and/or exposure epoch.
- 39. In addition to the completion and freeing of the RMA operations prior the the flush call (stages “c+f”), this call initializes the next RMA epoch (stage “i”).
- 40. The stages represent the invocation as part of an RMA operation. As collective procedure itself, it is a blocking procedure with all stages.

DRAFT

Procedure	Stages	Cpl	Loc	Blk	Op	Collective	Blocked resources	and remarks
						C sq S/W		
Chapter 3: Point-to-Point Communication								
MPI_SEND	i-s-c-f	c+f	nl	b	b-op	-		
MPI_SSEND	i-s-c-f	c+f	nl	b	b-op	-		1)
MPI_RSEND	i-s-c-f	c+f	l*	b	b-op	-		12)* 17)
MPI_BUFFER_ATTACH	i-----		l		b-op	-		attached buffer 22)
MPI_BSEND	i-s-c-f	c+f	l*	b	b-op	-		attached buffer 23) 12)*
MPI_BUFFER_DETACH	-----f	f	nl	b	b-op	-		24)
MPI_RECV	i-s-c-f	c+f	nl	b	b-op	-		
MPI_ISEND, MPI_ISSEND	i-s----	ic	l	nb	nb-op	-		buffer
MPI_IRECV	i-s----	ic	l	nb	nb-op	-		buffer
corresp. MPI_REQUEST_GET_STATUS ret. flag=FALSE	-----		l		nb-op	-		25)
corresp. MPI_REQUEST_GET_STATUS ret. flag=TRUE	----c-		l		nb-op	-		25)
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op	-		buffer 7)
corresponding MPI_TEST returning flag=TRUE	---c-f	c+f	l		nb-op	-		
corresponding MPI_WAIT	---c-f	c+f	nl		nb-op	-		
corresponding MPI_REQUEST_FREE (for active request)	-----		l		nb-op	-		buffer 26)
MPI_IBSEND, MPI_IRSEND	i-s----	ic	l	nb	nb-op	-		buffer
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op	-		buffer 7)
corresponding MPI_TEST returning flag=TRUE	---c-f	c+f	l		nb-op	-		
corresponding MPI_WAIT	---c-f	c+f	l*		nb-op	-		3)* 17)
corresponding MPI_REQUEST_FREE (for active request)	-----		l		nb-op	-		buffer 26)
MPI_ISEND, MPI_ISSEND, MPI_IBSEND, MPI_IRSEND	i-s----	ic	l	nb	nb-op	-		buffer 27)
MPI_IRECV	i-s----	ic	l	nb	nb-op	-		buffer
corresponding MPI_CANCEL	-----		l		nb-op	-		buffer 28)
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op	-		buffer 7)
corresponding MPI_TEST returning flag=TRUE	---c-f	c+f	l		nb-op	-		29)
corresponding MPI_WAIT	---c-f	c+f	l*		nb-op	-		29)
corresponding MPI_TEST_CANCELLED	-----		l		nb-op	-		
MPI_PROBE	-----	c	nl			-		2)
MPI_IPROBE	-----	c	l†			-		2) 11)†
MPI_MPROBE	i-----	ic	nl*	b	b-op	-		message 8)*
MPI_IMPROBE	i-----	ic	l	nb	b-op	-		message 15)
MPI_MRECV of a probed message	i-s-c-f	c+f	l	b	b-op	-		13)
MPI_IMRECV of a probed message	i-s----	ic	l	nb	nb-op	-		buffer 13)
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op	-		buffer 7)
corresponding MPI_TEST returning flag=TRUE	---c-f	c+f	l		nb-op	-		
corresponding MPI_WAIT	---c-f	c+f	l		nb-op	-		
corresponding MPI_REQUEST_FREE (for active request)	-----		l		nb-op	-		buffer 26)
MPI_[S B R]SEND_INIT, MPI_RECV_INIT	i-----	ic	l	nb†	p-op	-		buffer address 9) 14)†
corresponding MPI_START, MPI_STARTALL	--s----	ic	l	nb†	p-op	-		buffer address+contents 7) 14)†
corresponding MPI_TEST returning flag=FALSE	-----		l		p-op	-		buffer address+contents 7)
corresponding MPI_TEST returning flag=TRUE	----c-	c	l		p-op	-		buffer address 7) 9)
corresponding MPI_WAIT (for MPI_(B R)SEND_INIT req.)	----c-	c	l*		p-op	-		buffer address 3)* 7) 9) 17)
corresponding MPI_WAIT (for other request)	----c-	c	nl		p-op	-		buffer address 7) 9)
corresponding MPI_REQUEST_FREE (for active request)	-----		l		nb-op	-		buffer 26)
corresponding MPI_REQUEST_FREE (for inactive request)	-----f	f	l		p-op	-		
MPI_CANCEL of nonblocking/persistent pt-to-pt			l		p-op	-		
MPI_SENDRECV[_REPLACE]	i-s-c-f	c+f	nl	b	b-op	-		
MPI_ISENDRECV[_REPLACE]	i-s----	ic	l	nb	nb-op	-		buffer
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op	-		buffer 7)
corresponding MPI_TEST returning flag=TRUE	---c-f	c+f	l		nb-op	-		
corresponding MPI_WAIT	---c-f	c+f	nl		nb-op	-		
corresponding MPI_REQUEST_FREE (for active request)	-----		l		nb-op	-		buffer 26)
Chapter 4: Partitioned Point-to-Point Communication								
MPI_PSEND_INIT	i-----	ic	l	nb†	pp-op	-		buffer address 9) 14)†
corresponding MPI_START, MPI_STARTALL	--s----	ic	l	nb†	pp-op	-		buffer address 7) 9) 14)† 20)
corresponding MPI_PREADY and variants	--s----	ic	l	nb†	pp-op	-		buffer address+contents 7) 9) 14)† 20)
corresponding MPI_TEST returning flag=FALSE	-----		l		pp-op	-		buffer address+contents 7) 9)
corresponding MPI_TEST returning flag=TRUE	----c-	c	l		pp-op	-		buffer address 7) 9)
corresponding MPI_WAIT	----c-	c	nl		pp-op	-		buffer address 7) 9)
corresponding MPI_REQUEST_FREE (for inactive request)	-----f	f	l		pp-op	-		
MPI_PRECV_INIT	i-----	ic	l	nb†	pp-op	-		buffer address 9) 14)†
corresponding MPI_START, MPI_STARTALL	--s----	ic	l	nb†	pp-op	-		buffer address+contents 7) 9) 14)†
MPI_PARRIVED returning flag=TRUE	---c-	ic	l		pp-op	-		buffer address+contents 7) 9) 21)
MPI_PARRIVED returning flag=FALSE	-----	ic	l		pp-op	-		buffer address+contents 7) 9)
corresponding MPI_TEST returning flag=FALSE	-----		l		pp-op	-		buffer address+contents 7) 9)
corresponding MPI_TEST returning flag=TRUE	----c-	c	l		pp-op	-		buffer address 7) 9) 21)
corresponding MPI_WAIT	----c-	c	nl		pp-op	-		buffer address 7) 9) 21)
corresponding MPI_REQUEST_FREE (for inactive request)	-----f	f	l		pp-op	-		

Procedure	Stages	Cpl	Loc	Blk	Op	Collective	Blocked resources	and remarks
						C sq S/W		
Chapter 6: Collective Communication								
MPI_BCAST, MPI_BARRIER, MPI_GATHER, MPI_GATHERV, MPI_SCATTER, MPI_SCATTERV, MPI_ALLGATHER, MPI_ALLGATHERV, MPI_ALLTOALL, MPI_ALLTOALLV, MPI_ALLTOALLW, MPI_REDUCE, MPI_ALLREDUCE, MPI_REDUCE_SCATTER_BLOCK, MPI_REDUCE_SCATTER, MPI_SCAN, MPI_EXSCAN	i-s-c-f	c+f	nl	b	b-op	C sq W1		4) 18)
MPI_IBCAST, MPI_IBARRIER, MPI_IGATHER, MPI_ISCATTER, MPI_IALLGATHER, MPI_IALLTOALL, MPI_IREDUCE, MPI_IALLREDUCE, MPI_IREDUCE_SCATTER_BLOCK, MPI_ISCAN, MPI_IEXSCAN	i-s----	ic	l	nb	nb-op	C sq		buffer
MPI_IGATHERV, MPI_ISCATTERV, MPI_IALLGATHERV, MPI_IALLTOALLV, MPI_IALLTOALLW, MPI_IREDUCE_SCATTER	i-s----	ic	l	nb	nb-op	C sq		buffer, array arguments
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op			buffer, array arguments 7)
corresponding MPI_TEST returning flag=TRUE	----c-f	c+f	l		nb-op	C W2		4) 5) 18)
corresponding MPI_WAIT	----c-f	c+f	nl		nb-op	C W2		4) 5) 18)
MPI_BCAST_INIT, MPI_BARRIER_INIT, MPI_GATHER_INIT, MPI_SCATTER_INIT, MPI_ALLGATHER_INIT, MPI_ALLTOALL_INIT, MPI_REDUCE_INIT, MPI_ALLREDUCE_INIT, MPI_REDUCE_SCATTER_BLOCK_INIT, MPI_SCAN_INIT, MPI_EXSCAN_INIT	i-----	ic	nl*	b	p-op	C sq W1		buffer address 8)* 9)
MPI_GATHERV_INIT, MPI_SCATTERV_INIT, MPI_ALLGATHERV_INIT, MPI_ALLTOALLV_INIT, MPI_ALLTOALLW_INIT, MPI_REDUCE_SCATTER_INIT	i-----	ic	nl*	b	p-op	C sq W1		buffer address, array arguments 8)* 9)
corresponding MPI_START, MPI_STARTALL	--s----	ic	l	nb [‡]	p-op	C		buf.addr.+contents 7) 14) [‡] 16)
corresponding MPI_TEST returning flag=FALSE	-----		l		p-op			buf addr+contents & arr-args 7)
corresponding MPI_TEST returning flag=TRUE	----c--	c	l		p-op	C W2		buf-addr & arr-args 4) 5) 7) 9) 18)
corresponding MPI_WAIT	----c--	c	nl		p-op	C W2		buffer address and array arguments cached on the request handle 4) 5) 7) 9) 18)
corresponding MPI_REQUEST_FREE	-----f	f	l		p-op			
Chapter 7: Groups, Contexts, Communicators, and Caching								
MPI_COMM_CREATE, MPI_COMM_DUP, MPI_COMM_DUP_WITH_INFO, MPI_COMM_SPLIT, MPI_COMM_SPLIT_TYPE, MPI_COMM_SET_INFO	i-s-c-f	c+f	nl	b	b-op	C sq W1		coll. over comm arg.
MPI_COMM_CREATE_GROUP	i-s-c-f	c+f	nl	b	b-op	C sq W1		coll. over group arg.
MPI_INTERCOMM_CREATE, MPI_INTERCOMM_MERGE	i-s-c-f	c+f	nl	b	b-op	C sq W1		coll. over union of local & remote group
MPI_COMM_IDUP	i-s----	ic	l	nb	nb-op	C sq		communicator handle
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op			
corresponding MPI_TEST returning flag=TRUE	----c-f	c+f	l		nb-op	C W2		5)
corresponding MPI_WAIT	----c-f	c+f	nl		nb-op	C W2		5)
MPI_COMM_FREE	i-s----	ic	nl	b	nb-op	C sq W1		19)
MPI_COMM_DISCONNECT	i-s-c-f	c+f	nl	b	b-op	C sq W1		19)

Procedure	Stages	Cpl	Loc	Blk	Op	Collective	Blocked resources		and remarks
						C	sq	S/W	
Chapter 8: Process Topologies									
MPI_CART_CREATE, MPI_CART_SUB, MPI_GRAPH_CREATE, MPI_DIST_GRAPH_CREATE_ADJACENT, MPI_DIST_GRAPH_CREATE	i-s-c-f	c	nl	b	b-op	C	sq	W1	coll. over comm arg.
MPI_NEIGHBOR_ALLGATHER, MPI_NEIGHBOR_ALLGATHERV, MPI_NEIGHBOR_ALLTOALL, MPI_NEIGHBOR_ALLTOALLV, MPI_NEIGHBOR_ALLTOALLW	i-s-c-f	c+f	nl	b	b-op	C	sq	W1	18)
MPI_INEIGHBOR_ALLGATHER, MPI_INEIGHBOR_ALLTOALL	i-s----	ic	l	nb	nb-op	C	sq		buffer
MPI_INEIGHBOR_ALLGATHERV, MPI_INEIGHBOR_ALLTOALLV, MPI_INEIGHBOR_ALLTOALLW	i-s----	ic	l	nb	nb-op	C	sq		buffer, array arguments
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op				buffer, array arguments 7)
corresponding MPI_TEST returning flag=TRUE	----c-f	c+f	l		nb-op	C		W2	5) 18)
corresponding MPI_WAIT	----c-f	c+f	nl		nb-op	C		W2	5) 18)
MPI_NEIGHBOR_ALLGATHER_INIT, MPI_NEIGHBOR_ALLTOALL_INIT	i-----	ic	nl*	b	p-op	C	sq	W1	buffer address 8)* 9)
MPI_NEIGHBOR_ALLGATHERV_INIT, MPI_NEIGHBOR_ALLTOALLV_INIT, MPI_NEIGHBOR_ALLTOALLW_INIT	i-----	ic	nl*	b	p-op	C	sq	W1	buffer address, array arguments 8)* 9)
corresponding MPI_START, MPI_STARTALL	--s----	ic	l	nb [†]	p-op	C			buf.addr.+contents 7) 14) [‡] 16)
corresponding MPI_TEST returning flag=FALSE	-----		l		p-op				buf addr+contents & arr-args 7)
corresponding MPI_TEST returning flag=TRUE	----c--	c	l		p-op	C		W2	buf-addr & arr-args 5) 7) 9) 18)
corresponding MPI_WAIT	----c--	c	nl		p-op	C		W2	buffer address and array arguments cached on the request handle 5) 7) 9) 18)
corresponding MPI_REQUEST_FREE	-----f	f	l		p-op				
Chapter 11: Process Initialization, Creation, and Management									
MPI_INIT, MPI_INIT_THREAD	i-s-c-f	c+f	nl	b	b-op	C	sq	W1	collective over MPI_COMM_WORLD
MPI_FINALIZE	i-s-c-f	e+f	nl	b	b-op	C	sq	W1	collective over all connected processes
MPI_SESSION_INIT	-----		l						2)
MPI_SESSION_FINALIZE	i-s-c-f	c+f	nl	b	b-op	C	sq	W1	collective over connected processes scoped by the session
MPI_COMM_SPAWN, ... _MULTIPLE	i-s-c-f	c+f	nl	b	b-op	C	sq	W1	collective over comm 6)
MPI_COMM_ACCEPT, MPI_COMM_CONNECT	i-s-c-f	c+f	nl	b	b-op	C	sq	W1	collective over comm

Procedure	Stages	Cpl	Loc	Blk	Op	Collective	Blocked resources	and remarks
						C sq S/W		
Chapter 12: One-Sided Communication								
Window Allocation and Destruction								
MPI_WIN_CREATE, MPI_WIN_ALLOCATE, MPI_WIN_ALLOCATE_SHARED, MPI_WIN_CREATE_DYNAMIC	i-----		nl			C sq W1	window buffer address	(36) (40)
MPI_WIN_SET_INFO	i-----		nl			C sq W1	window buffer address	(40)
MPI_WIN_FREE	-----f	f	nl			C sq W1	(36) (40)	
With Fence Synchronization in the Role of an Origin Process								
MPI_WIN_FENCE	i-----	-	nl			C sq W1	(30) (38)	
MPI_PUT, MPI_GET, MPI_ACCUMULATE, MPI_GET_ACCUMULATE, MPI_COMPARE_AND_SWAP, MPI_FETCH_AND_OP	i-s----	ic	l	nb	nb-op	-	argument buffer	(14) (30)
MPI_WIN_FENCE	---c-f	c+f	nl			C sq W1	(31) (38)	
With Fence Synchronization in the Role of a Target Process								
MPI_WIN_FENCE	--s----	ic	nl			C sq W1	window buffer address+content	(38)
MPI_WIN_FENCE	---c-	c	nl			C sq W1	window buffer address	(32) (38)
With General Active Target Synchronization in the Role of an Origin Process								
MPI_WIN_START	i-----		nl			-	(30) (34)	
MPI_PUT, MPI_GET, MPI_ACCUMULATE, MPI_GET_ACCUMULATE, MPI_COMPARE_AND_SWAP, MPI_FETCH_AND_OP	i-s----	ic	l	nb	nb-op	-	argument buffer	(14) (30) (34)
MPI_WIN_COMPLETE	---c-f	c+f	l			-	(31) (34)	
With General Active Target Synchronization in the Role of a Target Process								
MPI_WIN_POST	--s----	ic	l			-	window buffer address+content	
MPI_WIN_TEST returning flag=FALSE	-----		l			-		
MPI_WIN_TEST returning flag=TRUE	---c-	c	l			-	window buffer address	(32)
MPI_WIN_WAIT	---c-	c	nl			-	window buffer address	(32)
With Lock/Unlock Synchronization in the Role of an Origin Process								
MPI_WIN_LOCK, MPI_WIN_LOCK_ALL	i-----		nl			-	(30) (35)	
MPI_PUT, MPI_GET, MPI_ACCUMULATE, MPI_GET_ACCUMULATE, MPI_COMPARE_AND_SWAP, MPI_FETCH_AND_OP	i-s----	ic	l	nb	nb-op	-	argument buffer	(14) (30) (35)
MPI_RPUT, MPI_RGET, MPI_RACCUMULATE, MPI_RGET_ACCUMULATE	i-s----	ic	l	nb	nb-op	-	argument buffer	(14) (30) (35)
corresponding MPI_TEST returning flag=FALSE	-----		l		nb-op	-	argument buffer	(7)
corresponding MPI_TEST returning flag=TRUE	---c-f	c+f	l		nb-op	-	(30) (31)	
corresponding MPI_WAIT	---c-f	c+f	l*		nb-op	-	(30) (31) (35)	
MPI_WIN_FLUSH_LOCAL, MPI_WIN_FLUSH_LOCAL_ALL	i--c-f	c+f	l			-	(30) (31) (35) (37) (39)	
MPI_WIN_FLUSH, MPI_WIN_FLUSH_ALL	i--c-f	c+f	l			-	(30) (33) (35) (37) (39)	
MPI_WIN_UNLOCK, MPI_WIN_UNLOCK_ALL	---c-f	c+f	l			-	(30) (33) (35) (37)	

Procedure	Stages	Cpl	Loc	Blk	Op	Collective			Blocked resources and remarks
						C	sq	S/W	
Chapter 14: I/O									
MPI_FILE_READ/WRITE[_AT]_SHARED], MPI_FILE_DELETE/SEEK/GET_VIEW	i-s-c-f	c+f	l*	b	b-op	-			12)*
MPI_FILE_READ/WRITE_AT_ALL, MPI_FILE_READ/WRITE_ALL ORDERED, MPI_FILE_OPEN/CLOSE/SEEK_SHARED, MPI_FILE_PREALLOCATE/SYNC, MPI_FILE_SET_VIEW/SIZE/INFO/ATOMICITY	i-s-c-f	c+f	nl	b	b-op	C	sq	W1	
MPI_FILE_IREAD/IWRITE[_AT]_SHARED]	i-s----	ic	l	nb	nb-op	-			buffer 10a)
MPI_FILE_IREAD/IWRITE[_AT]_ALL	i-s----	ic	l	nb	nb-op	C	sq		buffer 10b)
MPI_FILE_READ/WRITE[_AT]_ALL_BEGIN MPI_FILE_READ/WRITE_ORDERED_BEGIN	i-s----	ic	nl*	b	b-op	C	sq	W1	buffer 8)*
MPI_FILE_READ/WRITE[_AT]_ALL_END MPI_FILE_READ/WRITE_ORDERED_END	---c-f	c+f	nl	b	b-op	C	sq	W1	

A.3 C Bindings

A.3.1 Point-to-Point Communication C Bindings

```
1  int MPI_Bsend(const void *buf, int count, MPI_Datatype datatype, int dest,
2              int tag, MPI_Comm comm)
3
4  int MPI_Bsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
5                  int dest, int tag, MPI_Comm comm)
6
7  int MPI_Bsend_init(const void *buf, int count, MPI_Datatype datatype, int dest,
8                    int tag, MPI_Comm comm, MPI_Request *request)
9
10 int MPI_Bsend_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
11                     int dest, int tag, MPI_Comm comm, MPI_Request *request)
12
13 int MPI_Buffer_attach(void *buffer, int size)
14
15 int MPI_Buffer_attach_c(void *buffer, MPI_Count size)
16
17 int MPI_Buffer_detach(void *buffer_addr, int *size)
18
19 int MPI_Buffer_detach_c(void *buffer_addr, MPI_Count *size)
20
21 int MPI_Buffer_flush(void)
22
23 int MPI_Buffer_iflush(MPI_Request *request)
24
25 int MPI_Cancel(MPI_Request *request)
26
27 int MPI_Comm_attach_buffer(MPI_Comm comm, void *buffer, int size)
28
29 int MPI_Comm_attach_buffer_c(MPI_Comm comm, void *buffer, MPI_Count size)
30
31 int MPI_Comm_detach_buffer(MPI_Comm comm, void *buffer_addr, int *size)
32
33 int MPI_Comm_detach_buffer_c(MPI_Comm comm, void *buffer_addr, MPI_Count *size)
34
35 int MPI_Comm_flush_buffer(MPI_Comm comm)
36
37 int MPI_Comm_iflush_buffer(MPI_Comm comm, MPI_Request *request)
38
39 int MPI_Get_count(const MPI_Status *status, MPI_Datatype datatype, int *count)
40
41 int MPI_Get_count_c(const MPI_Status *status, MPI_Datatype datatype,
42                    MPI_Count *count)
43
44 int MPI_Ibsend(const void *buf, int count, MPI_Datatype datatype, int dest,
45               int tag, MPI_Comm comm, MPI_Request *request)
46
47 int MPI_Ibsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
48                  int dest, int tag, MPI_Comm comm, MPI_Request *request)
49
50 int MPI_Improbe(int source, int tag, MPI_Comm comm, int *flag,
51                MPI_Message *message, MPI_Status *status)
52
53 int MPI_Imrecv(void *buf, int count, MPI_Datatype datatype,
54               MPI_Message *message, MPI_Request *request)
55
56 int MPI_Imrecv_c(void *buf, MPI_Count count, MPI_Datatype datatype,
57                 MPI_Message *message, MPI_Request *request)
```

```
int MPI_Iprobe(int source, int tag, MPI_Comm comm, int *flag,
               MPI_Status *status)
int MPI_Irecv(void *buf, int count, MPI_Datatype datatype, int source, int tag,
              MPI_Comm comm, MPI_Request *request)
int MPI_Irecv_c(void *buf, MPI_Count count, MPI_Datatype datatype, int source,
                int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Irsend(const void *buf, int count, MPI_Datatype datatype, int dest,
               int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Irsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
                 int dest, int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Isend(const void *buf, int count, MPI_Datatype datatype, int dest,
              int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Isend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
                int dest, int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Isendrecv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
                  int dest, int sendtag, void *recvbuf, int recvcount,
                  MPI_Datatype recvtype, int source, int recvtag, MPI_Comm comm,
                  MPI_Request *request)
int MPI_Isendrecv_c(const void *sendbuf, MPI_Count sendcount,
                    MPI_Datatype sendtype, int dest, int sendtag, void *recvbuf,
                    MPI_Count recvcount, MPI_Datatype recvtype, int source,
                    int recvtag, MPI_Comm comm, MPI_Request *request)
int MPI_Isendrecv_replace(void *buf, int count, MPI_Datatype datatype,
                           int dest, int sendtag, int source, int recvtag, MPI_Comm comm,
                           MPI_Request *request)
int MPI_Isendrecv_replace_c(void *buf, MPI_Count count, MPI_Datatype datatype,
                             int dest, int sendtag, int source, int recvtag, MPI_Comm comm,
                             MPI_Request *request)
int MPI_Issend(const void *buf, int count, MPI_Datatype datatype, int dest,
               int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Issend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
                 int dest, int tag, MPI_Comm comm, MPI_Request *request)
int MPI_Mprobe(int source, int tag, MPI_Comm comm, MPI_Message *message,
               MPI_Status *status)
int MPI_Mrecv(void *buf, int count, MPI_Datatype datatype,
              MPI_Message *message, MPI_Status *status)
int MPI_Mrecv_c(void *buf, MPI_Count count, MPI_Datatype datatype,
                MPI_Message *message, MPI_Status *status)
int MPI_Probe(int source, int tag, MPI_Comm comm, MPI_Status *status)
int MPI_Recv(void *buf, int count, MPI_Datatype datatype, int source, int tag,
             MPI_Comm comm, MPI_Status *status)
```

```
1 int MPI_Recv_c(void *buf, MPI_Count count, MPI_Datatype datatype, int source,
2               int tag, MPI_Comm comm, MPI_Status *status)
3
4 int MPI_Recv_init(void *buf, int count, MPI_Datatype datatype, int source,
5                  int tag, MPI_Comm comm, MPI_Request *request)
6
7 int MPI_Recv_init_c(void *buf, MPI_Count count, MPI_Datatype datatype,
8                    int source, int tag, MPI_Comm comm, MPI_Request *request)
9
10 int MPI_Request_free(MPI_Request *request)
11
12 int MPI_Request_get_status(MPI_Request request, int *flag, MPI_Status *status)
13
14 int MPI_Request_get_status_all(int count,
15                               const MPI_Request array_of_requests[], int *flag,
16                               MPI_Status array_of_statuses[])
17
18 int MPI_Request_get_status_any(int count,
19                               const MPI_Request array_of_requests[], int *index, int *flag,
20                               MPI_Status *status)
21
22 int MPI_Request_get_status_some(int incount,
23                               const MPI_Request array_of_requests[], int *outcount,
24                               int array_of_indices[], MPI_Status array_of_statuses[])
25
26 int MPI_Rsend(const void *buf, int count, MPI_Datatype datatype, int dest,
27              int tag, MPI_Comm comm)
28
29 int MPI_Rsend_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
30                int dest, int tag, MPI_Comm comm)
31
32 int MPI_Rsend_init(const void *buf, int count, MPI_Datatype datatype, int dest,
33                   int tag, MPI_Comm comm, MPI_Request *request)
34
35 int MPI_Rsend_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
36                     int dest, int tag, MPI_Comm comm, MPI_Request *request)
37
38 int MPI_Send(const void *buf, int count, MPI_Datatype datatype, int dest,
39              int tag, MPI_Comm comm)
40
41 int MPI_Send_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
42                int dest, int tag, MPI_Comm comm)
43
44 int MPI_Send_init(const void *buf, int count, MPI_Datatype datatype, int dest,
45                  int tag, MPI_Comm comm, MPI_Request *request)
46
47 int MPI_Send_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype,
48                    int dest, int tag, MPI_Comm comm, MPI_Request *request)
49
50 int MPI_Sendrecv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
51                 int dest, int sendtag, void *recvbuf, int recvcount,
52                 MPI_Datatype recvtype, int source, int recvtag, MPI_Comm comm,
53                 MPI_Status *status)
54
55 int MPI_Sendrecv_c(const void *sendbuf, MPI_Count sendcount,
56                   MPI_Datatype sendtype, int dest, int sendtag, void *recvbuf,
57                   MPI_Count recvcount, MPI_Datatype recvtype, int source,
58                   int recvtag, MPI_Comm comm, MPI_Status *status)
```



```
int MPI_Sendrecv_replace(void *buf, int count, MPI_Datatype datatype, int dest, 1
                          int sendtag, int source, int recvtag, MPI_Comm comm, 2
                          MPI_Status *status) 3
int MPI_Sendrecv_replace_c(void *buf, MPI_Count count, MPI_Datatype datatype, 4
                          int dest, int sendtag, int source, int recvtag, MPI_Comm comm, 5
                          MPI_Status *status) 6
int MPI_Session_attach_buffer(MPI_Session session, void *buffer, int size) 7
int MPI_Session_attach_buffer_c(MPI_Session session, void *buffer, 8
                              MPI_Count size) 9
int MPI_Session_detach_buffer(MPI_Session session, void *buffer_addr, 10
                              int *size) 11
int MPI_Session_detach_buffer_c(MPI_Session session, void *buffer_addr, 12
                              MPI_Count *size) 13
int MPI_Session_flush_buffer(MPI_Session session) 14
int MPI_Session_iflush_buffer(MPI_Session session, MPI_Request *request) 15
int MPI_Ssend(const void *buf, int count, MPI_Datatype datatype, int dest, 16
              int tag, MPI_Comm comm) 17
int MPI_Ssend_c(const void *buf, MPI_Count count, MPI_Datatype datatype, 18
               int dest, int tag, MPI_Comm comm) 19
int MPI_Ssend_init(const void *buf, int count, MPI_Datatype datatype, int dest, 20
                  int tag, MPI_Comm comm, MPI_Request *request) 21
int MPI_Ssend_init_c(const void *buf, MPI_Count count, MPI_Datatype datatype, 22
                    int dest, int tag, MPI_Comm comm, MPI_Request *request) 23
int MPI_Start(MPI_Request *request) 24
int MPI_Startall(int count, MPI_Request array_of_requests[]) 25
int MPI_Status_get_error(MPI_Status *status, int *err) 26
int MPI_Status_get_source(MPI_Status *status, int *source) 27
int MPI_Status_get_tag(MPI_Status *status, int *tag) 28
int MPI_Test(MPI_Request *request, int *flag, MPI_Status *status) 29
int MPI_Test_cancelled(const MPI_Status *status, int *flag) 30
int MPI_Testall(int count, MPI_Request array_of_requests[], int *flag, 31
               MPI_Status array_of_statuses[]) 32
int MPI_Testany(int count, MPI_Request array_of_requests[], int *index, 33
               int *flag, MPI_Status *status) 34
int MPI_Testsome(int incount, MPI_Request array_of_requests[], int *outcount, 35
                 int array_of_indices[], MPI_Status array_of_statuses[]) 36
int MPI_Wait(MPI_Request *request, MPI_Status *status) 37
int MPI_Waitall(int count, MPI_Request array_of_requests[], MPI_Status array_of_statuses[]) 38
int MPI_Waitany(int count, MPI_Request array_of_requests[], MPI_Status *status, MPI_Request *request) 39
int MPI_Waitsome(int incount, MPI_Request array_of_requests[], MPI_Status array_of_statuses[], MPI_Request *request) 40
int MPI_Win_create(void *buf, MPI_Count size, MPI_Datatype datatype, MPI_Comm comm, MPI_Win *win) 41
int MPI_Win_create_c(void *buf, MPI_Count count, MPI_Datatype datatype, MPI_Comm comm, MPI_Win *win) 42
int MPI_Win_detach(MPI_Win *win, MPI_Comm comm) 43
int MPI_Win_detach_c(MPI_Win *win, MPI_Comm comm) 44
int MPI_Win_flush_all(MPI_Win *win, MPI_Comm comm) 45
int MPI_Win_flush_all_c(MPI_Win *win, MPI_Comm comm) 46
int MPI_Win_flush_local(MPI_Win *win, MPI_Comm comm) 47
int MPI_Win_flush_local_c(MPI_Win *win, MPI_Comm comm) 48
```

```

1 int MPI_Waitall(int count, MPI_Request array_of_requests[],
2               MPI_Status array_of_statuses[])
3
4 int MPI_Waitany(int count, MPI_Request array_of_requests[], int *index,
5               MPI_Status *status)
6
7 int MPI_Waitsome(int incount, MPI_Request array_of_requests[], int *outcount,
8               int array_of_indices[], MPI_Status array_of_statuses[])

```

A.3.2 Partitioned Communication C Bindings

```

10
11 int MPI_Parrived(MPI_Request request, int partition, int *flag)
12
13 int MPI_Pready(int partition, MPI_Request request)
14
15 int MPI_Pready_list(int length, const int array_of_partitions[],
16                   MPI_Request request)
17
18 int MPI_Pready_range(int partition_low, int partition_high,
19                   MPI_Request request)
20
21 int MPI_Precv_init(void *buf, int partitions, MPI_Count count,
22                 MPI_Datatype datatype, int source, int tag, MPI_Comm comm,
23                 MPI_Info info, MPI_Request *request)
24
25 int MPI_Psend_init(const void *buf, int partitions, MPI_Count count,
26                 MPI_Datatype datatype, int dest, int tag, MPI_Comm comm,
27                 MPI_Info info, MPI_Request *request)

```

A.3.3 Datatypes C Bindings

```

28 MPI_Aint MPI_Aint_add(MPI_Aint base, MPI_Aint disp)
29
30 MPI_Aint MPI_Aint_diff(MPI_Aint addr1, MPI_Aint addr2)
31
32 int MPI_Get_address(const void *location, MPI_Aint *address)
33
34 int MPI_Get_elements(const MPI_Status *status, MPI_Datatype datatype,
35                   int *count)
36
37 int MPI_Get_elements_c(const MPI_Status *status, MPI_Datatype datatype,
38                   MPI_Count *count)
39
40 int MPI_Pack(const void *inbuf, int incount, MPI_Datatype datatype,
41            void *outbuf, int outsize, int *position, MPI_Comm comm)
42
43 int MPI_Pack_c(const void *inbuf, MPI_Count incount, MPI_Datatype datatype,
44            void *outbuf, MPI_Count outsize, MPI_Count *position,
45            MPI_Comm comm)
46
47 int MPI_Pack_external(const char datarep[], const void *inbuf, int incount,
48                   MPI_Datatype datatype, void *outbuf, MPI_Aint outsize,
49                   MPI_Aint *position)
50
51 int MPI_Pack_external_c(const char datarep[], const void *inbuf,
52                   MPI_Count incount, MPI_Datatype datatype, void *outbuf,
53                   MPI_Count outsize, MPI_Count *position)

```

```
int MPI_Pack_external_size(const char datarep[], int incount,      1
                          MPI_Datatype datatype, MPI_Aint *size)  2
int MPI_Pack_external_size_c(const char datarep[], MPI_Count incount,  3
                             MPI_Datatype datatype, MPI_Count *size)  4
int MPI_Pack_size(int incount, MPI_Datatype datatype, MPI_Comm comm, int *size)  5
int MPI_Pack_size_c(MPI_Count incount, MPI_Datatype datatype, MPI_Comm comm,  6
                   MPI_Count *size)  7
int MPI_Type_commit(MPI_Datatype *datatype)  8
int MPI_Type_contiguous(int count, MPI_Datatype oldtype, MPI_Datatype *newtype)  9
int MPI_Type_contiguous_c(MPI_Count count, MPI_Datatype oldtype, 10
                          MPI_Datatype *newtype) 11
int MPI_Type_create_darray(int size, int rank, int ndims, 12
                          const int array_of_gsizes[], const int array_of_distribs[], 13
                          const int array_of_dargs[], const int array_of_psize[], 14
                          int order, MPI_Datatype oldtype, MPI_Datatype *newtype) 15
int MPI_Type_create_darray_c(int size, int rank, int ndims, 16
                             const MPI_Count array_of_gsizes[], const int array_of_distribs[], 17
                             const int array_of_dargs[], const int array_of_psize[], 18
                             int order, MPI_Datatype oldtype, MPI_Datatype *newtype) 19
int MPI_Type_create_hindexed(int count, const int array_of_blocklengths[], 20
                             const MPI_Aint array_of_displacements[], MPI_Datatype oldtype, 21
                             MPI_Datatype *newtype) 22
int MPI_Type_create_hindexed_block(int count, int blocklength, 23
                                   const MPI_Aint array_of_displacements[], MPI_Datatype oldtype, 24
                                   MPI_Datatype *newtype) 25
int MPI_Type_create_hindexed_block_c(MPI_Count count, MPI_Count blocklength, 26
                                     const MPI_Count array_of_displacements[], MPI_Datatype oldtype, 27
                                     MPI_Datatype *newtype) 28
int MPI_Type_create_hindexed_c(MPI_Count count, 29
                               const MPI_Count array_of_blocklengths[], 30
                               const MPI_Count array_of_displacements[], MPI_Datatype oldtype, 31
                               MPI_Datatype *newtype) 32
int MPI_Type_create_hvector(int count, int blocklength, MPI_Aint stride, 33
                            MPI_Datatype oldtype, MPI_Datatype *newtype) 34
int MPI_Type_create_hvector_c(MPI_Count count, MPI_Count blocklength, 35
                              MPI_Count stride, MPI_Datatype oldtype, MPI_Datatype *newtype) 36
int MPI_Type_create_indexed_block(int count, int blocklength, 37
                                  const int array_of_displacements[], MPI_Datatype oldtype, 38
                                  MPI_Datatype *newtype) 39
int MPI_Type_create_indexed_block_c(MPI_Count count, MPI_Count blocklength, 40
                                    const MPI_Count array_of_displacements[], MPI_Datatype oldtype, 41
                                    MPI_Datatype *newtype) 42
int MPI_Type_create_indexed_block_c(MPI_Count count, MPI_Count blocklength, 43
                                    const MPI_Count array_of_displacements[], MPI_Datatype oldtype, 44
                                    MPI_Datatype *newtype) 45
int MPI_Type_create_indexed_block_c(MPI_Count count, MPI_Count blocklength, 46
                                    const MPI_Count array_of_displacements[], MPI_Datatype oldtype, 47
                                    MPI_Datatype *newtype) 48
```

```
1 int MPI_Type_create_resized(MPI_Datatype oldtype, MPI_Aint lb, MPI_Aint extent,
2     MPI_Datatype *newtype)
3
4 int MPI_Type_create_resized_c(MPI_Datatype oldtype, MPI_Count lb,
5     MPI_Count extent, MPI_Datatype *newtype)
6
7 int MPI_Type_create_struct(int count, const int array_of_blocklengths[],
8     const MPI_Aint array_of_displacements[],
9     const MPI_Datatype array_of_types[], MPI_Datatype *newtype)
10
11 int MPI_Type_create_struct_c(MPI_Count count,
12     const MPI_Count array_of_blocklengths[],
13     const MPI_Count array_of_displacements[],
14     const MPI_Datatype array_of_types[], MPI_Datatype *newtype)
15
16 int MPI_Type_create_subarray(int ndims, const int array_of_sizes[],
17     const int array_of_subsizes[], const int array_of_starts[],
18     int order, MPI_Datatype oldtype, MPI_Datatype *newtype)
19
20 int MPI_Type_create_subarray_c(int ndims, const MPI_Count array_of_sizes[],
21     const MPI_Count array_of_subsizes[],
22     const MPI_Count array_of_starts[], int order,
23     MPI_Datatype oldtype, MPI_Datatype *newtype)
24
25 int MPI_Type_dup(MPI_Datatype oldtype, MPI_Datatype *newtype)
26
27 int MPI_Type_free(MPI_Datatype *datatype)
28
29 int MPI_Type_get_contents(MPI_Datatype datatype, int max_integers,
30     int max_addresses, int max_datatypes, int array_of_integers[],
31     MPI_Aint array_of_addresses[], MPI_Datatype array_of_datatypes[])
32
33 int MPI_Type_get_contents_c(MPI_Datatype datatype, MPI_Count max_integers,
34     MPI_Count max_addresses, MPI_Count max_large_counts,
35     MPI_Count max_datatypes, int array_of_integers[],
36     MPI_Aint array_of_addresses[], MPI_Count array_of_large_counts[],
37     MPI_Datatype array_of_datatypes[])
38
39 int MPI_Type_get_envelope(MPI_Datatype datatype, int *num_integers,
40     int *num_addresses, int *num_datatypes, int *combiner)
41
42 int MPI_Type_get_envelope_c(MPI_Datatype datatype, MPI_Count *num_integers,
43     MPI_Count *num_addresses, MPI_Count *num_large_counts,
44     MPI_Count *num_datatypes, int *combiner)
45
46 int MPI_Type_get_extent(MPI_Datatype datatype, MPI_Aint *lb, MPI_Aint *extent)
47
48 int MPI_Type_get_extent_c(MPI_Datatype datatype, MPI_Count *lb,
49     MPI_Count *extent)
50
51 int MPI_Type_get_true_extent(MPI_Datatype datatype, MPI_Aint *true_lb,
52     MPI_Aint *true_extent)
53
54 int MPI_Type_get_true_extent_c(MPI_Datatype datatype, MPI_Count *true_lb,
55     MPI_Count *true_extent)
```

```

int MPI_Type_indexed(int count, const int array_of_blocklengths[],
                    const int array_of_displacements[], MPI_Datatype oldtype,
                    MPI_Datatype *newtype)
int MPI_Type_indexed_c(MPI_Count count,
                    const MPI_Count array_of_blocklengths[],
                    const MPI_Count array_of_displacements[], MPI_Datatype oldtype,
                    MPI_Datatype *newtype)
int MPI_Type_size(MPI_Datatype datatype, int *size)
int MPI_Type_size_c(MPI_Datatype datatype, MPI_Count *size)
int MPI_Type_vector(int count, int blocklength, int stride,
                    MPI_Datatype oldtype, MPI_Datatype *newtype)
int MPI_Type_vector_c(MPI_Count count, MPI_Count blocklength, MPI_Count stride,
                    MPI_Datatype oldtype, MPI_Datatype *newtype)
int MPI_Unpack(const void *inbuf, int insize, int *position, void *outbuf,
              int outcount, MPI_Datatype datatype, MPI_Comm comm)
int MPI_Unpack_c(const void *inbuf, MPI_Count insize, MPI_Count *position,
                void *outbuf, MPI_Count outcount, MPI_Datatype datatype,
                MPI_Comm comm)
int MPI_Unpack_external(const char datarep[], const void *inbuf,
                       MPI_Aint insize, MPI_Aint *position, void *outbuf, int outcount,
                       MPI_Datatype datatype)
int MPI_Unpack_external_c(const char datarep[], const void *inbuf,
                          MPI_Count insize, MPI_Count *position, void *outbuf,
                          MPI_Count outcount, MPI_Datatype datatype)

```

A.3.4 Collective Communication C Bindings

```

int MPI_Allgather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
                 void *recvbuf, int recvcount, MPI_Datatype recvtype,
                 MPI_Comm comm)
int MPI_Allgather_c(const void *sendbuf, MPI_Count sendcount,
                   MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                   MPI_Datatype recvtype, MPI_Comm comm)
int MPI_Allgather_init(const void *sendbuf, int sendcount,
                      MPI_Datatype sendtype, void *recvbuf, int recvcount,
                      MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                      MPI_Request *request)
int MPI_Allgather_init_c(const void *sendbuf, MPI_Count sendcount,
                        MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                        MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                        MPI_Request *request)

```

```
1 int MPI_Allgatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
2                   void *recvbuf, const int recvcnts[], const int displs[],
3                   MPI_Datatype recvtype, MPI_Comm comm)
4
5 int MPI_Allgatherv_c(const void *sendbuf, MPI_Count sendcount,
6                    MPI_Datatype sendtype, void *recvbuf,
7                    const MPI_Count recvcnts[], const MPI_Aint displs[],
8                    MPI_Datatype recvtype, MPI_Comm comm)
9
10 int MPI_Allgatherv_init(const void *sendbuf, int sendcount,
11                       MPI_Datatype sendtype, void *recvbuf, const int recvcnts[],
12                      const int displs[], MPI_Datatype recvtype, MPI_Comm comm,
13                      MPI_Info info, MPI_Request *request)
14
15 int MPI_Allgatherv_init_c(const void *sendbuf, MPI_Count sendcount,
16                          MPI_Datatype sendtype, void *recvbuf,
17                          const MPI_Count recvcnts[], const MPI_Aint displs[],
18                          MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
19                          MPI_Request *request)
20
21 int MPI_Allreduce(const void *sendbuf, void *recvbuf, int count,
22                 MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
23
24 int MPI_Allreduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,
25                   MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
26
27 int MPI_Allreduce_init(const void *sendbuf, void *recvbuf, int count,
28                      MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
29                      MPI_Request *request)
30
31 int MPI_Allreduce_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
32                         MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
33                         MPI_Request *request)
34
35 int MPI_Alltoall(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
36                void *recvbuf, int recvcnt, MPI_Datatype recvtype,
37                MPI_Comm comm)
38
39 int MPI_Alltoall_c(const void *sendbuf, MPI_Count sendcount,
40                  MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcnt,
41                  MPI_Datatype recvtype, MPI_Comm comm)
42
43 int MPI_Alltoall_init(const void *sendbuf, int sendcount,
44                    MPI_Datatype sendtype, void *recvbuf, int recvcnt,
45                    MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
46                    MPI_Request *request)
47
48 int MPI_Alltoall_init_c(const void *sendbuf, MPI_Count sendcount,
49                       MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcnt,
50                       MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
51                       MPI_Request *request)
52
53 int MPI_Alltoallv(const void *sendbuf, const int sendcounts[],
54                 const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
```

```

    const int recvcounts[], const int rdispls[],
    MPI_Datatype recvtype, MPI_Comm comm)
1
2
3
int MPI_Alltoallv_c(const void *sendbuf, const MPI_Count sendcounts[],
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
    const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf,
    const MPI_Count recvcounts[], const MPI_Aint rdispls[],
    MPI_Datatype recvtype, MPI_Comm comm)

int MPI_Alltoallv_init(const void *sendbuf, const int sendcounts[],
    const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
    const int recvcounts[], const int rdispls[],
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
    MPI_Request *request)

int MPI_Alltoallv_init_c(const void *sendbuf, const MPI_Count sendcounts[],
    const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf,
    const MPI_Count recvcounts[], const MPI_Aint rdispls[],
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
    MPI_Request *request)

int MPI_Alltoallw(const void *sendbuf, const int sendcounts[],
    const int sdispls[], const MPI_Datatype sendtypes[],
    void *recvbuf, const int recvcounts[], const int rdispls[],
    const MPI_Datatype recvtypes[], MPI_Comm comm)

int MPI_Alltoallw_c(const void *sendbuf, const MPI_Count sendcounts[],
    const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
    void *recvbuf, const MPI_Count recvcounts[],
    const MPI_Aint rdispls[], const MPI_Datatype recvtypes[],
    MPI_Comm comm)

int MPI_Alltoallw_init(const void *sendbuf, const int sendcounts[],
    const int sdispls[], const MPI_Datatype sendtypes[],
    void *recvbuf, const int recvcounts[], const int rdispls[],
    const MPI_Datatype recvtypes[], MPI_Comm comm, MPI_Info info,
    MPI_Request *request)

int MPI_Alltoallw_init_c(const void *sendbuf, const MPI_Count sendcounts[],
    const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
    void *recvbuf, const MPI_Count recvcounts[],
    const MPI_Aint rdispls[], const MPI_Datatype recvtypes[],
    MPI_Comm comm, MPI_Info info, MPI_Request *request)

int MPI_Barrier(MPI_Comm comm)

int MPI_Barrier_init(MPI_Comm comm, MPI_Info info, MPI_Request *request)

int MPI_Bcast(void *buffer, int count, MPI_Datatype datatype, int root,
    MPI_Comm comm)

int MPI_Bcast_c(void *buffer, MPI_Count count, MPI_Datatype datatype, int root,
    MPI_Comm comm)

int MPI_Bcast_init(void *buffer, int count, MPI_Datatype datatype, int root,
    MPI_Comm comm, MPI_Info info, MPI_Request *request)

```

```
1 int MPI_Bcast_init_c(void *buffer, MPI_Count count, MPI_Datatype datatype,
2     int root, MPI_Comm comm, MPI_Info info, MPI_Request *request)
3
4 int MPI_Exscan(const void *sendbuf, void *recvbuf, int count,
5     MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
6
7 int MPI_Exscan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
8     MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
9
10 int MPI_Exscan_init(const void *sendbuf, void *recvbuf, int count,
11     MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
12     MPI_Request *request)
13
14 int MPI_Exscan_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
15     MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
16     MPI_Request *request)
17
18 int MPI_Gather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
19     void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
20     MPI_Comm comm)
21
22 int MPI_Gather_c(const void *sendbuf, MPI_Count sendcount,
23     MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
24     MPI_Datatype recvtype, int root, MPI_Comm comm)
25
26 int MPI_Gather_init(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
27     void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
28     MPI_Comm comm, MPI_Info info, MPI_Request *request)
29
30 int MPI_Gather_init_c(const void *sendbuf, MPI_Count sendcount,
31     MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
32     MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info,
33     MPI_Request *request)
34
35 int MPI_Gatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
36     void *recvbuf, const int recvcounts[], const int displs[],
37     MPI_Datatype recvtype, int root, MPI_Comm comm)
38
39 int MPI_Gatherv_c(const void *sendbuf, MPI_Count sendcount,
40     MPI_Datatype sendtype, void *recvbuf,
41     const MPI_Count recvcounts[], const MPI_Aint displs[],
42     MPI_Datatype recvtype, int root, MPI_Comm comm)
43
44 int MPI_Gatherv_init(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
45     void *recvbuf, const int recvcounts[], const int displs[],
46     MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info,
47     MPI_Request *request)
48
49 int MPI_Gatherv_init_c(const void *sendbuf, MPI_Count sendcount,
50     MPI_Datatype sendtype, void *recvbuf,
51     const MPI_Count recvcounts[], const MPI_Aint displs[],
52     MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info,
53     MPI_Request *request)
```



```

int MPI_Iallgather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,      1
                  void *recvbuf, int recvcount, MPI_Datatype recvtype,          2
                  MPI_Comm comm, MPI_Request *request)                          3
                                                                              4
int MPI_Iallgather_c(const void *sendbuf, MPI_Count sendcount,                  5
                    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,  6
                    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request) 7
                                                                              8
int MPI_Iallgatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,  9
                   void *recvbuf, const int recvcounts[], const int displs[], 10
                   MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request) 11
                                                                              12
int MPI_Iallgatherv_c(const void *sendbuf, MPI_Count sendcount,                13
                    MPI_Datatype sendtype, void *recvbuf,                    14
                    const MPI_Count recvcounts[], const MPI_Aint displs[],    15
                    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)
                                                                              16
int MPI_Iallreduce(const void *sendbuf, void *recvbuf, int count,              17
                  MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,            18
                  MPI_Request *request)                                        19
                                                                              20
int MPI_Iallreduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,     21
                    MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,          22
                    MPI_Request *request)                                     23
                                                                              24
int MPI_Ialltoall(const void *sendbuf, int sendcount, MPI_Datatype sendtype,  25
                 void *recvbuf, int recvcount, MPI_Datatype recvtype,        26
                 MPI_Comm comm, MPI_Request *request)                        27
                                                                              28
int MPI_Ialltoall_c(const void *sendbuf, MPI_Count sendcount,                 29
                   MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount, 30
                   MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request) 31
                                                                              32
int MPI_Ialltoallv(const void *sendbuf, const int sendcounts[],               33
                  const int sdispls[], MPI_Datatype sendtype, void *recvbuf,  34
                  const int recvcounts[], const int rdispls[],                35
                  MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request) 36
                                                                              37
int MPI_Ialltoallv_c(const void *sendbuf, const MPI_Count sendcounts[],        38
                   const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf, 39
                   const MPI_Count recvcounts[], const MPI_Aint rdispls[],    40
                   MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request) 41
                                                                              42
int MPI_Ialltoallw(const void *sendbuf, const int sendcounts[],               43
                  const int sdispls[], const MPI_Datatype sendtypes[],        44
                  void *recvbuf, const int recvcounts[], const int rdispls[], 45
                  const MPI_Datatype recvtypes[], MPI_Comm comm,              46
                  MPI_Request *request)                                        47
                                                                              48
int MPI_Ialltoallw_c(const void *sendbuf, const MPI_Count sendcounts[],        49
                   const MPI_Aint sdispls[], const MPI_Datatype sendtypes[], 50
                   void *recvbuf, const MPI_Count recvcounts[],               51
                   const MPI_Aint rdispls[], const MPI_Datatype recvtypes[], 52
                   MPI_Comm comm, MPI_Request *request)                      53

```

```
1 int MPI_Ibarrier(MPI_Comm comm, MPI_Request *request)
2
3 int MPI_Ibcast(void *buffer, int count, MPI_Datatype datatype, int root,
4               MPI_Comm comm, MPI_Request *request)
5
6 int MPI_Ibcast_c(void *buffer, MPI_Count count, MPI_Datatype datatype,
7                 int root, MPI_Comm comm, MPI_Request *request)
8
9 int MPI_Iexscan(const void *sendbuf, void *recvbuf, int count,
10               MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
11               MPI_Request *request)
12
13 int MPI_Iexscan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
14                 MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
15                 MPI_Request *request)
16
17 int MPI_Igather(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
18               void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
19               MPI_Comm comm, MPI_Request *request)
20
21 int MPI_Igather_c(const void *sendbuf, MPI_Count sendcount,
22                 MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
23                 MPI_Datatype recvtype, int root, MPI_Comm comm,
24                 MPI_Request *request)
25
26 int MPI_Igatherv(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
27               void *recvbuf, const int recvcnts[], const int displs[],
28               MPI_Datatype recvtype, int root, MPI_Comm comm,
29               MPI_Request *request)
30
31 int MPI_Igatherv_c(const void *sendbuf, MPI_Count sendcount,
32                 MPI_Datatype sendtype, void *recvbuf,
33                 const MPI_Count recvcnts[], const MPI_Aint displs[],
34                 MPI_Datatype recvtype, int root, MPI_Comm comm,
35                 MPI_Request *request)
36
37 int MPI_Ireduce(const void *sendbuf, void *recvbuf, int count,
38               MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
39               MPI_Request *request)
40
41 int MPI_Ireduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,
42                 MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
43                 MPI_Request *request)
44
45 int MPI_Ireduce_scatter(const void *sendbuf, void *recvbuf,
46                       const int recvcnts[], MPI_Datatype datatype, MPI_Op op,
47                       MPI_Comm comm, MPI_Request *request)
48
49 int MPI_Ireduce_scatter_block(const void *sendbuf, void *recvbuf,
50                              int recvcount, MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
51                              MPI_Request *request)
52
53 int MPI_Ireduce_scatter_block_c(const void *sendbuf, void *recvbuf,
54                                MPI_Count recvcount, MPI_Datatype datatype, MPI_Op op,
55                                MPI_Comm comm, MPI_Request *request)
```

```

int MPI_Ireduce_scatter_c(const void *sendbuf, void *recvbuf,
    const MPI_Count recvcnts[], MPI_Datatype datatype, MPI_Op op,
    MPI_Comm comm, MPI_Request *request)
int MPI_Iscan(const void *sendbuf, void *recvbuf, int count,
    MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
    MPI_Request *request)
int MPI_Iscan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
    MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
    MPI_Request *request)
int MPI_Iscatter(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
    void *recvbuf, int recvcnt, MPI_Datatype recvtype, int root,
    MPI_Comm comm, MPI_Request *request)
int MPI_Iscatter_c(const void *sendbuf, MPI_Count sendcount,
    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcnt,
    MPI_Datatype recvtype, int root, MPI_Comm comm,
    MPI_Request *request)
int MPI_Iscatterv(const void *sendbuf, const int sendcounts[],
    const int displs[], MPI_Datatype sendtype, void *recvbuf,
    int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm,
    MPI_Request *request)
int MPI_Iscatterv_c(const void *sendbuf, const MPI_Count sendcounts[],
    const MPI_Aint displs[], MPI_Datatype sendtype, void *recvbuf,
    MPI_Count recvcnt, MPI_Datatype recvtype, int root,
    MPI_Comm comm, MPI_Request *request)
int MPI_Op_commutative(MPI_Op op, int *commute)
int MPI_Op_create(MPI_User_function *user_fn, int commute, MPI_Op *op)
int MPI_Op_create_c(MPI_User_function_c *user_fn, int commute, MPI_Op *op)
int MPI_Op_free(MPI_Op *op)
int MPI_Reduce(const void *sendbuf, void *recvbuf, int count,
    MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)
int MPI_Reduce_c(const void *sendbuf, void *recvbuf, MPI_Count count,
    MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)
int MPI_Reduce_init(const void *sendbuf, void *recvbuf, int count,
    MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
    MPI_Info info, MPI_Request *request)
int MPI_Reduce_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
    MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm,
    MPI_Info info, MPI_Request *request)
int MPI_Reduce_local(const void *inbuf, void *inoutbuf, int count,
    MPI_Datatype datatype, MPI_Op op)

```

```
1 int MPI_Reduce_local_c(const void *inbuf, void *inoutbuf, MPI_Count count,
2                       MPI_Datatype datatype, MPI_Op op)
3
4 int MPI_Reduce_scatter(const void *sendbuf, void *recvbuf,
5                       const int recvcnts[], MPI_Datatype datatype, MPI_Op op,
6                       MPI_Comm comm)
7
8 int MPI_Reduce_scatter_block(const void *sendbuf, void *recvbuf, int recvcount,
9                             MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
10
11 int MPI_Reduce_scatter_block_c(const void *sendbuf, void *recvbuf,
12                               MPI_Count recvcount, MPI_Datatype datatype, MPI_Op op,
13                               MPI_Comm comm)
14
15 int MPI_Reduce_scatter_block_init(const void *sendbuf, void *recvbuf,
16                                  int recvcount, MPI_Datatype datatype, MPI_Op op, MPI_Comm comm,
17                                  MPI_Info info, MPI_Request *request)
18
19 int MPI_Reduce_scatter_block_init_c(const void *sendbuf, void *recvbuf,
20                                    MPI_Count recvcount, MPI_Datatype datatype, MPI_Op op,
21                                    MPI_Comm comm, MPI_Info info, MPI_Request *request)
22
23 int MPI_Reduce_scatter_c(const void *sendbuf, void *recvbuf,
24                          const MPI_Count recvcnts[], MPI_Datatype datatype, MPI_Op op,
25                          MPI_Comm comm)
26
27 int MPI_Reduce_scatter_init(const void *sendbuf, void *recvbuf,
28                             const int recvcnts[], MPI_Datatype datatype, MPI_Op op,
29                             MPI_Comm comm, MPI_Info info, MPI_Request *request)
30
31 int MPI_Reduce_scatter_init_c(const void *sendbuf, void *recvbuf,
32                               const MPI_Count recvcnts[], MPI_Datatype datatype, MPI_Op op,
33                               MPI_Comm comm, MPI_Info info, MPI_Request *request)
34
35 int MPI_Scan(const void *sendbuf, void *recvbuf, int count,
36             MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
37
38 int MPI_Scan_c(const void *sendbuf, void *recvbuf, MPI_Count count,
39              MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
40
41 int MPI_Scan_init(const void *sendbuf, void *recvbuf, int count,
42                 MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
43                 MPI_Request *request)
44
45 int MPI_Scan_init_c(const void *sendbuf, void *recvbuf, MPI_Count count,
46                   MPI_Datatype datatype, MPI_Op op, MPI_Comm comm, MPI_Info info,
47                   MPI_Request *request)
48
49 int MPI_Scatter(const void *sendbuf, int sendcount, MPI_Datatype sendtype,
50               void *recvbuf, int recvcount, MPI_Datatype recvtype, int root,
51               MPI_Comm comm)
52
53 int MPI_Scatter_c(const void *sendbuf, MPI_Count sendcount,
54                 MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
55                 MPI_Datatype recvtype, int root, MPI_Comm comm)
```

```

int MPI_Scatter_init(const void *sendbuf, int sendcount, MPI_Datatype sendtype, 1
                    void *recvbuf, int recvcount, MPI_Datatype recvtype, int root, 2
                    MPI_Comm comm, MPI_Info info, MPI_Request *request) 3
4
int MPI_Scatter_init_c(const void *sendbuf, MPI_Count sendcount, 5
                      MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount, 6
                      MPI_Datatype recvtype, int root, MPI_Comm comm, MPI_Info info, 7
                      MPI_Request *request) 8
9
int MPI_Scatterv(const void *sendbuf, const int sendcounts[], 9
                const int displs[], MPI_Datatype sendtype, void *recvbuf, 10
                int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm) 11
12
int MPI_Scatterv_c(const void *sendbuf, const MPI_Count sendcounts[], 13
                  const MPI_Aint displs[], MPI_Datatype sendtype, void *recvbuf, 14
                  MPI_Count recvcount, MPI_Datatype recvtype, int root, 15
                  MPI_Comm comm) 16
17
int MPI_Scatterv_init(const void *sendbuf, const int sendcounts[], 17
                     const int displs[], MPI_Datatype sendtype, void *recvbuf, 18
                     int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm, 19
                     MPI_Info info, MPI_Request *request) 20
21
int MPI_Scatterv_init_c(const void *sendbuf, const MPI_Count sendcounts[], 22
                       const MPI_Aint displs[], MPI_Datatype sendtype, void *recvbuf, 23
                       MPI_Count recvcount, MPI_Datatype recvtype, int root, 24
                       MPI_Comm comm, MPI_Info info, MPI_Request *request) 25
26
int MPI_Type_get_value_index(MPI_Datatype value_type, MPI_Datatype index_type, 26
                             MPI_Datatype *pair_type) 27
28
A.3.5 Groups, Contexts, Communicators, and Caching C Bindings 29
int MPI_Comm_compare(MPI_Comm comm1, MPI_Comm comm2, int *result) 30
int MPI_Comm_create(MPI_Comm comm, MPI_Group group, MPI_Comm *newcomm) 32
int MPI_Comm_create_from_group(MPI_Group group, const char *stringtag, 34
                               MPI_Info info, MPI_Errhandler errhandler, MPI_Comm *newcomm) 35
int MPI_Comm_create_group(MPI_Comm comm, MPI_Group group, int tag, 36
                          MPI_Comm *newcomm) 37
int MPI_Comm_create_keyval(MPI_Comm_copy_attr_function *comm_copy_attr_fn, 38
                          MPI_Comm_delete_attr_function *comm_delete_attr_fn, 39
                          int *comm_keyval, void *extra_state) 40
int MPI_Comm_delete_attr(MPI_Comm comm, int comm_keyval) 42
int MPI_Comm_dup(MPI_Comm comm, MPI_Comm *newcomm) 43
int MPI_COMM_DUP_FN(MPI_Comm oldcomm, int comm_keyval, void *extra_state, 45
                   void *attribute_val_in, void *attribute_val_out, int *flag) 46
int MPI_Comm_dup_with_info(MPI_Comm comm, MPI_Info info, MPI_Comm *newcomm) 47

```

```
1 int MPI_Comm_free(MPI_Comm *comm)
2
3 int MPI_Comm_free_keyval(int *comm_keyval)
4
5 int MPI_Comm_get_attr(MPI_Comm comm, int comm_keyval, void *attribute_val,
6                       int *flag)
7
8 int MPI_Comm_get_info(MPI_Comm comm, MPI_Info *info_used)
9
10 int MPI_Comm_get_name(MPI_Comm comm, char *comm_name, int *resultlen)
11
12 int MPI_Comm_group(MPI_Comm comm, MPI_Group *group)
13
14 int MPI_Comm_idup(MPI_Comm comm, MPI_Comm *newcomm, MPI_Request *request)
15
16 int MPI_Comm_idup_with_info(MPI_Comm comm, MPI_Info info, MPI_Comm *newcomm,
17                             MPI_Request *request)
18
19 int MPI_COMM_NULL_COPY_FN(MPI_Comm oldcomm, int comm_keyval, void *extra_state,
20                             void *attribute_val_in, void *attribute_val_out, int *flag)
21
22 int MPI_COMM_NULL_DELETE_FN(MPI_Comm comm, int comm_keyval,
23                             void *attribute_val, void *extra_state)
24
25 int MPI_Comm_rank(MPI_Comm comm, int *rank)
26
27 int MPI_Comm_remote_group(MPI_Comm comm, MPI_Group *group)
28
29 int MPI_Comm_remote_size(MPI_Comm comm, int *size)
30
31 int MPI_Comm_set_attr(MPI_Comm comm, int comm_keyval, void *attribute_val)
32
33 int MPI_Comm_set_info(MPI_Comm comm, MPI_Info info)
34
35 int MPI_Comm_set_name(MPI_Comm comm, const char *comm_name)
36
37 int MPI_Comm_size(MPI_Comm comm, int *size)
38
39 int MPI_Comm_split(MPI_Comm comm, int color, int key, MPI_Comm *newcomm)
40
41 int MPI_Comm_split_type(MPI_Comm comm, int split_type, int key, MPI_Info info,
42                         MPI_Comm *newcomm)
43
44 int MPI_Comm_test_inter(MPI_Comm comm, int *flag)
45
46 int MPI_Group_compare(MPI_Group group1, MPI_Group group2, int *result)
47
48 int MPI_Group_difference(MPI_Group group1, MPI_Group group2,
49                         MPI_Group *newgroup)
50
51 int MPI_Group_excl(MPI_Group group, int n, const int ranks[],
52                   MPI_Group *newgroup)
53
54 int MPI_Group_free(MPI_Group *group)
55
56 int MPI_Group_from_session_pset(MPI_Session session, const char *pset_name,
57                                 MPI_Group *newgroup)
58
59 int MPI_Group_incl(MPI_Group group, int n, const int ranks[],
60                   MPI_Group *newgroup)
61
62 int MPI_Group_intersection(MPI_Group group1, MPI_Group group2,
63                            MPI_Group *newgroup)
```

```

int MPI_Group_range_excl(MPI_Group group, int n, int ranges[][3],
                        MPI_Group *newgroup)
int MPI_Group_range_incl(MPI_Group group, int n, int ranges[][3],
                        MPI_Group *newgroup)
int MPI_Group_rank(MPI_Group group, int *rank)
int MPI_Group_size(MPI_Group group, int *size)
int MPI_Group_translate_ranks(MPI_Group group1, int n, const int ranks1[],
                             MPI_Group group2, int ranks2[])
int MPI_Group_union(MPI_Group group1, MPI_Group group2, MPI_Group *newgroup)
int MPI_Intercomm_create(MPI_Comm local_comm, int local_leader,
                        MPI_Comm peer_comm, int remote_leader, int tag,
                        MPI_Comm *newintercomm)
int MPI_Intercomm_create_from_groups(MPI_Group local_group, int local_leader,
                                     MPI_Group remote_group, int remote_leader, const char *stringtag,
                                     MPI_Info info, MPI_Errhandler errhandler, MPI_Comm *newintercomm)
int MPI_Intercomm_merge(MPI_Comm intercomm, int high, MPI_Comm *newintracomm)
int MPI_Type_create_keyval(MPI_Type_copy_attr_function *type_copy_attr_fn,
                          MPI_Type_delete_attr_function *type_delete_attr_fn,
                          int *type_keyval, void *extra_state)
int MPI_Type_delete_attr(MPI_Datatype datatype, int type_keyval)
int MPI_TYPE_DUP_FN(MPI_Datatype oldtype, int type_keyval, void *extra_state,
                   void *attribute_val_in, void *attribute_val_out, int *flag)
int MPI_Type_free_keyval(int *type_keyval)
int MPI_Type_get_attr(MPI_Datatype datatype, int type_keyval,
                    void *attribute_val, int *flag)
int MPI_Type_get_name(MPI_Datatype datatype, char *type_name, int *resultlen)
int MPI_TYPE_NULL_COPY_FN(MPI_Datatype oldtype, int type_keyval,
                          void *extra_state, void *attribute_val_in,
                          void *attribute_val_out, int *flag)
int MPI_TYPE_NULL_DELETE_FN(MPI_Datatype datatype, int type_keyval,
                            void *attribute_val, void *extra_state)
int MPI_Type_set_attr(MPI_Datatype datatype, int type_keyval,
                    void *attribute_val)
int MPI_Type_set_name(MPI_Datatype datatype, const char *type_name)
int MPI_Win_create_keyval(MPI_Win_copy_attr_function *win_copy_attr_fn,
                         MPI_Win_delete_attr_function *win_delete_attr_fn,
                         int *win_keyval, void *extra_state)
int MPI_Win_delete_attr(MPI_Win win, int win_keyval)

```

```
1 int MPI_WIN_DUP_FN(MPI_Win oldwin, int win_keyval, void *extra_state,
2                   void *attribute_val_in, void *attribute_val_out, int *flag)
3
4 int MPI_Win_free_keyval(int *win_keyval)
5
6 int MPI_Win_get_attr(MPI_Win win, int win_keyval, void *attribute_val,
7                     int *flag)
8
9 int MPI_Win_get_name(MPI_Win win, char *win_name, int *resultlen)
10
11 int MPI_WIN_NULL_COPY_FN(MPI_Win oldwin, int win_keyval, void *extra_state,
12                          void *attribute_val_in, void *attribute_val_out, int *flag)
13
14 int MPI_WIN_NULL_DELETE_FN(MPI_Win win, int win_keyval, void *attribute_val,
15                            void *extra_state)
16
17 int MPI_Win_set_attr(MPI_Win win, int win_keyval, void *attribute_val)
18
19 int MPI_Win_set_name(MPI_Win win, const char *win_name)
```

A.3.6 Virtual Topologies for MPI Processes C Bindings

```
19 int MPI_Cart_coords(MPI_Comm comm, int rank, int maxdims, int coords[])
20
21 int MPI_Cart_create(MPI_Comm comm_old, int ndims, const int dims[],
22                   const int periods[], int reorder, MPI_Comm *comm_cart)
23
24 int MPI_Cart_get(MPI_Comm comm, int maxdims, int dims[], int periods[],
25                 int coords[])
26
27 int MPI_Cart_map(MPI_Comm comm, int ndims, const int dims[],
28                 const int periods[], int *newrank)
29
30 int MPI_Cart_rank(MPI_Comm comm, const int coords[], int *rank)
31
32 int MPI_Cart_shift(MPI_Comm comm, int direction, int disp, int *rank_source,
33                   int *rank_dest)
34
35 int MPI_Cart_sub(MPI_Comm comm, const int remain_dims[], MPI_Comm *newcomm)
36
37 int MPI_Cartdim_get(MPI_Comm comm, int *ndims)
38
39 int MPI_Dims_create(int nnodes, int ndims, int dims[])
40
41 int MPI_Dist_graph_create(MPI_Comm comm_old, int n, const int sources[],
42                          const int degrees[], const int destinations[],
43                          const int weights[], MPI_Info info, int reorder,
44                          MPI_Comm *comm_dist_graph)
45
46 int MPI_Dist_graph_create_adjacent(MPI_Comm comm_old, int indegree,
47                                    const int sources[], const int sourceweights[], int outdegree,
48                                    const int destinations[], const int destweights[], MPI_Info info,
49                                    int reorder, MPI_Comm *comm_dist_graph)
50
51 int MPI_Dist_graph_neighbors(MPI_Comm comm, int maxindegree, int sources[],
52                              int sourceweights[], int maxoutdegree, int destinations[],
53                              int destweights[])
```



```

int MPI_Dist_graph_neighbors_count(MPI_Comm comm, int *indegree,      1
    int *outdegree, int *weighted)                                   2
                                                                    3
int MPI_Graph_create(MPI_Comm comm_old, int nnodes, const int index[],  4
    const int edges[], int reorder, MPI_Comm *comm_graph)          5
                                                                    6
int MPI_Graph_get(MPI_Comm comm, int maxindex, int maxedges, int index[],  7
    int edges[])                                                  8
                                                                    9
int MPI_Graph_map(MPI_Comm comm, int nnodes, const int index[],      10
    const int edges[], int *newrank)                               11
                                                                    12
int MPI_Graph_neighbors(MPI_Comm comm, int rank, int maxneighbors,   13
    int neighbors[])                                              14
                                                                    15
int MPI_Graph_neighbors_count(MPI_Comm comm, int rank, int *nneighbors) 16
                                                                    17
int MPI_Graphdims_get(MPI_Comm comm, int *nnodes, int *nedges)      18
                                                                    19
int MPI_Ineighbor_allgather(const void *sendbuf, int sendcount,      20
    MPI_Datatype sendtype, void *recvbuf, int recvcount,          21
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)    22
                                                                    23
int MPI_Ineighbor_allgather_c(const void *sendbuf, MPI_Count sendcount, 24
    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,    25
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)    26
                                                                    27
int MPI_Ineighbor_allgather_v(const void *sendbuf, int sendcount,    28
    MPI_Datatype sendtype, void *recvbuf, const int recvcounts[],  29
    const int displs[], MPI_Datatype recvtype, MPI_Comm comm,      30
    MPI_Request *request)                                          31
                                                                    32
int MPI_Ineighbor_allgather_v_c(const void *sendbuf, MPI_Count sendcount, 33
    MPI_Datatype sendtype, void *recvbuf,                          34
    const MPI_Count recvcounts[], const MPI_Aint displs[],          35
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)    36
                                                                    37
int MPI_Ineighbor_alltoall(const void *sendbuf, int sendcount,      38
    MPI_Datatype sendtype, void *recvbuf, int recvcount,          39
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)    40
                                                                    41
int MPI_Ineighbor_alltoall_c(const void *sendbuf, MPI_Count sendcount, 42
    MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,    43
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)    44
                                                                    45
int MPI_Ineighbor_alltoall_v(const void *sendbuf, const int sendcounts[], 46
    const int sdispls[], MPI_Datatype sendtype, void *recvbuf,    47
    const int recvcounts[], const int rdispls[],                    48
    MPI_Datatype recvtype, MPI_Comm comm, MPI_Request *request)

```

```
1 int MPI_Ineighbor_alltoallw(const void *sendbuf, const int sendcounts[],
2     const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
3     void *recvbuf, const int recvcounts[], const MPI_Aint rdispls[],
4     const MPI_Datatype recvtypes[], MPI_Comm comm,
5     MPI_Request *request)
6
7 int MPI_Ineighbor_alltoallw_c(const void *sendbuf,
8     const MPI_Count sendcounts[], const MPI_Aint sdispls[],
9     const MPI_Datatype sendtypes[], void *recvbuf,
10    const MPI_Count recvcounts[], const MPI_Aint rdispls[],
11    const MPI_Datatype recvtypes[], MPI_Comm comm,
12    MPI_Request *request)
13
14 int MPI_Neighbor_allgather(const void *sendbuf, int sendcount,
15     MPI_Datatype sendtype, void *recvbuf, int recvcount,
16     MPI_Datatype recvttype, MPI_Comm comm)
17
18 int MPI_Neighbor_allgather_c(const void *sendbuf, MPI_Count sendcount,
19     MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
20     MPI_Datatype recvttype, MPI_Comm comm)
21
22 int MPI_Neighbor_allgather_init(const void *sendbuf, int sendcount,
23     MPI_Datatype sendtype, void *recvbuf, int recvcount,
24     MPI_Datatype recvttype, MPI_Comm comm, MPI_Info info,
25     MPI_Request *request)
26
27 int MPI_Neighbor_allgather_init_c(const void *sendbuf, MPI_Count sendcount,
28     MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
29     MPI_Datatype recvttype, MPI_Comm comm, MPI_Info info,
30     MPI_Request *request)
31
32 int MPI_Neighbor_allgatherv(const void *sendbuf, int sendcount,
33     MPI_Datatype sendtype, void *recvbuf, const int recvcounts[],
34     const int displs[], MPI_Datatype recvttype, MPI_Comm comm)
35
36 int MPI_Neighbor_allgatherv_c(const void *sendbuf, MPI_Count sendcount,
37     MPI_Datatype sendtype, void *recvbuf,
38     const MPI_Count recvcounts[], const MPI_Aint displs[],
39     MPI_Datatype recvttype, MPI_Comm comm)
40
41 int MPI_Neighbor_allgatherv_init(const void *sendbuf, int sendcount,
42     MPI_Datatype sendtype, void *recvbuf, const int recvcounts[],
43     const int displs[], MPI_Datatype recvttype, MPI_Comm comm,
44     MPI_Info info, MPI_Request *request)
45
46 int MPI_Neighbor_allgatherv_init_c(const void *sendbuf, MPI_Count sendcount,
47     MPI_Datatype sendtype, void *recvbuf,
48     const MPI_Count recvcounts[], const MPI_Aint displs[],
49     MPI_Datatype recvttype, MPI_Comm comm, MPI_Info info,
50     MPI_Request *request)
```

```

int MPI_Neighbor_alltoall_c(const void *sendbuf, MPI_Count sendcount,      1
                           MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                           MPI_Datatype recvtype, MPI_Comm comm)          2
                                                                              3
int MPI_Neighbor_alltoall_init(const void *sendbuf, int sendcount,        4
                               MPI_Datatype sendtype, void *recvbuf, int recvcount,
                               MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                               MPI_Request *request)                       5
                                                                              6
                                                                              7
                                                                              8
int MPI_Neighbor_alltoall_init_c(const void *sendbuf, MPI_Count sendcount, 9
                                 MPI_Datatype sendtype, void *recvbuf, MPI_Count recvcount,
                                 MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                                 MPI_Request *request)                     10
                                                                              11
                                                                              12
int MPI_Neighbor_alltoallv(const void *sendbuf, const int sendcounts[],    13
                           const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
                           const int rcvcounts[], const int rdispls[],
                           MPI_Datatype recvtype, MPI_Comm comm)          14
                                                                              15
                                                                              16
                                                                              17
int MPI_Neighbor_alltoallv_c(const void *sendbuf, const MPI_Count sendcounts[], 18
                             const MPI_Aint sdispls[], MPI_Datatype sendtype, void *recvbuf,
                             const MPI_Count rcvcounts[], const MPI_Aint rdispls[],
                             MPI_Datatype recvtype, MPI_Comm comm)        19
                                                                              20
                                                                              21
int MPI_Neighbor_alltoallv_init(const void *sendbuf, const int sendcounts[], 22
                                const int sdispls[], MPI_Datatype sendtype, void *recvbuf,
                                const int rcvcounts[], const int rdispls[],
                                MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                                MPI_Request *request)                       23
                                                                              24
                                                                              25
                                                                              26
                                                                              27
int MPI_Neighbor_alltoallv_init_c(const void *sendbuf,
                                  const MPI_Count sendcounts[], const MPI_Aint sdispls[],
                                  MPI_Datatype sendtype, void *recvbuf,
                                  const MPI_Count rcvcounts[], const MPI_Aint rdispls[],
                                  MPI_Datatype recvtype, MPI_Comm comm, MPI_Info info,
                                  MPI_Request *request)                     28
                                                                              29
                                                                              30
                                                                              31
                                                                              32
                                                                              33
int MPI_Neighbor_alltoallw(const void *sendbuf, const int sendcounts[],    34
                           const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
                           void *recvbuf, const int rcvcounts[], const MPI_Aint rdispls[],
                           const MPI_Datatype rcvtypes[], MPI_Comm comm)   35
                                                                              36
                                                                              37
                                                                              38
int MPI_Neighbor_alltoallw_c(const void *sendbuf, const MPI_Count sendcounts[], 39
                             const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
                             void *recvbuf, const MPI_Count rcvcounts[],
                             const MPI_Aint rdispls[], const MPI_Datatype rcvtypes[],
                             MPI_Comm comm)                                40
                                                                              41
                                                                              42
                                                                              43
int MPI_Neighbor_alltoallw_init(const void *sendbuf, const int sendcounts[], 44
                                const MPI_Aint sdispls[], const MPI_Datatype sendtypes[],
                                void *recvbuf, const int rcvcounts[], const MPI_Aint rdispls[],
                                const MPI_Datatype rcvtypes[], MPI_Comm comm, MPI_Info info,
                                MPI_Request *request)                       45
                                                                              46
                                                                              47
                                                                              48

```

```
1 int MPI_Neighbor_alltoallw_init_c(const void *sendbuf,
2     const MPI_Count sendcounts[], const MPI_Aint sdispls[],
3     const MPI_Datatype sendtypes[], void *recvbuf,
4     const MPI_Count recvcounts[], const MPI_Aint rdispls[],
5     const MPI_Datatype recvtypes[], MPI_Comm comm, MPI_Info info,
6     MPI_Request *request)
7
8 int MPI_Topo_test(MPI_Comm comm, int *status)
```

10 A.3.7 MPI Environmental Management C Bindings

```
11 int MPI_Add_error_class(int *errorclass)
12
13 int MPI_Add_error_code(int errorclass, int *errorcode)
14
15 int MPI_Add_error_string(int errorcode, const char *string)
16
17 int MPI_Alloc_mem(MPI_Aint size, MPI_Info info, void *baseptr)
18
19 int MPI_Comm_call_errhandler(MPI_Comm comm, int errorcode)
20
21 int MPI_Comm_create_errhandler(
22     MPI_Comm_errhandler_function *comm_errhandler_fn,
23     MPI_Errhandler *errhandler)
24
25 int MPI_Comm_get_errhandler(MPI_Comm comm, MPI_Errhandler *errhandler)
26
27 int MPI_Comm_set_errhandler(MPI_Comm comm, MPI_Errhandler errhandler)
28
29 int MPI_Errhandler_free(MPI_Errhandler *errhandler)
30
31 int MPI_Error_class(int errorcode, int *errorclass)
32
33 int MPI_Error_string(int errorcode, char *string, int *resultlen)
34
35 int MPI_File_call_errhandler(MPI_File fh, int errorcode)
36
37 int MPI_File_create_errhandler(
38     MPI_File_errhandler_function *file_errhandler_fn,
39     MPI_Errhandler *errhandler)
40
41 int MPI_File_get_errhandler(MPI_File file, MPI_Errhandler *errhandler)
42
43 int MPI_File_set_errhandler(MPI_File file, MPI_Errhandler errhandler)
44
45 int MPI_Free_mem(void *base)
46
47 int MPI_Get_hw_resource_info(MPI_Info *hw_info)
48
49 int MPI_Get_library_version(char *version, int *resultlen)
50
51 int MPI_Get_processor_name(char *name, int *resultlen)
52
53 int MPI_Get_version(int *version, int *subversion)
54
55 int MPI_Remove_error_class(int errorclass)
56
57 int MPI_Remove_error_code(int errorcode)
58
59 int MPI_Remove_error_string(int errorcode)
```

```

int MPI_Session_call_errhandler(MPI_Session session, int errorcode) 1
int MPI_Session_create_errhandler( 2
    MPI_Session_errhandler_function *session_errhandler_fn, 3
    MPI_Errhandler *errhandler) 4
int MPI_Session_get_errhandler(MPI_Session session, MPI_Errhandler *errhandler) 5
int MPI_Session_set_errhandler(MPI_Session session, MPI_Errhandler errhandler) 6
int MPI_Win_call_errhandler(MPI_Win win, int errorcode) 7
int MPI_Win_create_errhandler(MPI_Win_errhandler_function *win_errhandler_fn, 8
    MPI_Errhandler *errhandler) 9
int MPI_Win_get_errhandler(MPI_Win win, MPI_Errhandler *errhandler) 10
int MPI_Win_set_errhandler(MPI_Win win, MPI_Errhandler errhandler) 11
double MPI_Wtick(void) 12
double MPI_Wtime(void) 13

```

A.3.8 The Info Object C Bindings

```

int MPI_Info_create(MPI_Info *info) 14
int MPI_Info_create_env(int argc, char *argv[], MPI_Info *info) 15
int MPI_Info_delete(MPI_Info info, const char *key) 16
int MPI_Info_dup(MPI_Info info, MPI_Info *newinfo) 17
int MPI_Info_free(MPI_Info *info) 18
int MPI_Info_get_nkeys(MPI_Info info, int *nkeys) 19
int MPI_Info_get_nthkey(MPI_Info info, int n, char *key) 20
int MPI_Info_get_string(MPI_Info info, const char *key, int *buflen, 21
    char *value, int *flag) 22
int MPI_Info_set(MPI_Info info, const char *key, const char *value) 23

```

A.3.9 Process Creation and Management C Bindings

```

int MPI_Abort(MPI_Comm comm, int errorcode) 24
int MPI_Close_port(const char *port_name) 25
int MPI_Comm_accept(const char *port_name, MPI_Info info, int root, 26
    MPI_Comm comm, MPI_Comm *newcomm) 27
int MPI_Comm_connect(const char *port_name, MPI_Info info, int root, 28
    MPI_Comm comm, MPI_Comm *newcomm) 29
int MPI_Comm_disconnect(MPI_Comm *comm) 30
int MPI_Comm_get_parent(MPI_Comm *parent) 31
int MPI_Comm_join(int fd, MPI_Comm *intercomm) 32

```

```
1 int MPI_Comm_spawn(const char *command, char *argv[], int maxprocs,
2                   MPI_Info info, int root, MPI_Comm comm, MPI_Comm *intercomm,
3                   int array_of_errcodes[])
4
5 int MPI_Comm_spawn_multiple(int count, char *array_of_commands[],
6                             char **array_of_argv[], const int array_of_maxprocs[],
7                             const MPI_Info array_of_info[], int root, MPI_Comm comm,
8                             MPI_Comm *intercomm, int array_of_errcodes[])
9
10 int MPI_Finalize(void)
11
12 int MPI_Finalized(int *flag)
13
14 int MPI_Init(int *argc, char ***argv)
15
16 int MPI_Init_thread(int *argc, char ***argv, int required, int *provided)
17
18 int MPI_Initialized(int *flag)
19
20 int MPI_Is_thread_main(int *flag)
21
22 int MPI_Lookup_name(const char *service_name, MPI_Info info, char *port_name)
23
24 int MPI_Open_port(MPI_Info info, char *port_name)
25
26 int MPI_Publish_name(const char *service_name, MPI_Info info,
27                     const char *port_name)
28
29 int MPI_Query_thread(int *provided)
30
31 int MPI_Session_finalize(MPI_Session *session)
32
33 int MPI_Session_get_info(MPI_Session session, MPI_Info *info_used)
34
35 int MPI_Session_get_nth_pset(MPI_Session session, MPI_Info info, int n,
36                             int *pset_len, char *pset_name)
37
38 int MPI_Session_get_num_psets(MPI_Session session, MPI_Info info,
39                             int *npset_names)
40
41 int MPI_Session_get_pset_info(MPI_Session session, const char *pset_name,
42                             MPI_Info *info)
43
44 int MPI_Session_init(MPI_Info info, MPI_Errhandler errhandler,
45                    MPI_Session *session)
46
47 int MPI_Unpublish_name(const char *service_name, MPI_Info info,
48                      const char *port_name)
```

A.3.10 One-Sided Communications C Bindings

```
42 int MPI_Accumulate(const void *origin_addr, int origin_count,
43                  MPI_Datatype origin_datatype, int target_rank,
44                  MPI_Aint target_disp, int target_count,
45                  MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
46
47 int MPI_Accumulate_c(const void *origin_addr, MPI_Count origin_count,
48                    MPI_Datatype origin_datatype, int target_rank,
```

```

        MPI_Aint target_disp, MPI_Count target_count,
        MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
1
2
3
int MPI_Compare_and_swap(const void *origin_addr, const void *compare_addr,
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
        MPI_Aint target_disp, MPI_Count target_count,
        MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
int MPI_Fetch_and_op(const void *origin_addr, void *result_addr,
        MPI_Datatype datatype, int target_rank, MPI_Aint target_disp,
        MPI_Op op, MPI_Win win)
int MPI_Get(void *origin_addr, int origin_count, MPI_Datatype origin_datatype,
        int target_rank, MPI_Aint target_disp, int target_count,
        MPI_Datatype target_datatype, MPI_Win win)
int MPI_Get_accumulate(const void *origin_addr, int origin_count,
        MPI_Datatype origin_datatype, void *result_addr,
        int result_count, MPI_Datatype result_datatype, int target_rank,
        MPI_Aint target_disp, int target_count,
        MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
int MPI_Get_accumulate_c(const void *origin_addr, MPI_Count origin_count,
        MPI_Datatype origin_datatype, void *result_addr,
        MPI_Count result_count, MPI_Datatype result_datatype,
        int target_rank, MPI_Aint target_disp, MPI_Count target_count,
        MPI_Datatype target_datatype, MPI_Op op, MPI_Win win)
int MPI_Get_c(void *origin_addr, MPI_Count origin_count,
        MPI_Datatype origin_datatype, int target_rank,
        MPI_Aint target_disp, MPI_Count target_count,
        MPI_Datatype target_datatype, MPI_Win win)
int MPI_Put(const void *origin_addr, int origin_count,
        MPI_Datatype origin_datatype, int target_rank,
        MPI_Aint target_disp, int target_count,
        MPI_Datatype target_datatype, MPI_Win win)
int MPI_Put_c(const void *origin_addr, MPI_Count origin_count,
        MPI_Datatype origin_datatype, int target_rank,
        MPI_Aint target_disp, MPI_Count target_count,
        MPI_Datatype target_datatype, MPI_Win win)
int MPI_Raccumulate(const void *origin_addr, int origin_count,
        MPI_Datatype origin_datatype, int target_rank,
        MPI_Aint target_disp, int target_count,
        MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
        MPI_Request *request)
int MPI_Raccumulate_c(const void *origin_addr, MPI_Count origin_count,
        MPI_Datatype origin_datatype, int target_rank,
        MPI_Aint target_disp, MPI_Count target_count,
        MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
        MPI_Request *request)

```

```
1 int MPI_Rget(void *origin_addr, int origin_count, MPI_Datatype origin_datatype,
2             int target_rank, MPI_Aint target_disp, int target_count,
3             MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)
4
5 int MPI_Rget_accumulate(const void *origin_addr, int origin_count,
6                         MPI_Datatype origin_datatype, void *result_addr,
7                         int result_count, MPI_Datatype result_datatype, int target_rank,
8                         MPI_Aint target_disp, int target_count,
9                         MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
10                        MPI_Request *request)
11
12 int MPI_Rget_accumulate_c(const void *origin_addr, MPI_Count origin_count,
13                          MPI_Datatype origin_datatype, void *result_addr,
14                          MPI_Count result_count, MPI_Datatype result_datatype,
15                          int target_rank, MPI_Aint target_disp, MPI_Count target_count,
16                          MPI_Datatype target_datatype, MPI_Op op, MPI_Win win,
17                          MPI_Request *request)
18
19 int MPI_Rget_c(void *origin_addr, MPI_Count origin_count,
20               MPI_Datatype origin_datatype, int target_rank,
21               MPI_Aint target_disp, MPI_Count target_count,
22               MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)
23
24 int MPI_Rput(const void *origin_addr, int origin_count,
25             MPI_Datatype origin_datatype, int target_rank,
26             MPI_Aint target_disp, int target_count,
27             MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)
28
29 int MPI_Rput_c(const void *origin_addr, MPI_Count origin_count,
30               MPI_Datatype origin_datatype, int target_rank,
31               MPI_Aint target_disp, MPI_Count target_count,
32               MPI_Datatype target_datatype, MPI_Win win, MPI_Request *request)
33
34 int MPI_Win_allocate(MPI_Aint size, int disp_unit, MPI_Info info,
35                    MPI_Comm comm, void *baseptr, MPI_Win *win)
36
37 int MPI_Win_allocate_c(MPI_Aint size, MPI_Aint disp_unit, MPI_Info info,
38                      MPI_Comm comm, void *baseptr, MPI_Win *win)
39
40 int MPI_Win_allocate_shared(MPI_Aint size, int disp_unit, MPI_Info info,
41                            MPI_Comm comm, void *baseptr, MPI_Win *win)
42
43 int MPI_Win_allocate_shared_c(MPI_Aint size, MPI_Aint disp_unit, MPI_Info info,
44                              MPI_Comm comm, void *baseptr, MPI_Win *win)
45
46 int MPI_Win_attach(MPI_Win win, void *base, MPI_Aint size)
47
48 int MPI_Win_complete(MPI_Win win)
49
50 int MPI_Win_create(void *base, MPI_Aint size, int disp_unit, MPI_Info info,
51                  MPI_Comm comm, MPI_Win *win)
52
53 int MPI_Win_create_c(void *base, MPI_Aint size, MPI_Aint disp_unit,
54                    MPI_Info info, MPI_Comm comm, MPI_Win *win)
55
56 int MPI_Win_create_dynamic(MPI_Info info, MPI_Comm comm, MPI_Win *win)
```



```

int MPI_Win_detach(MPI_Win win, const void *base) 1
int MPI_Win_fence(int assert, MPI_Win win) 2
int MPI_Win_flush(int rank, MPI_Win win) 3
int MPI_Win_flush_all(MPI_Win win) 4
int MPI_Win_flush_local(int rank, MPI_Win win) 5
int MPI_Win_flush_local_all(MPI_Win win) 6
int MPI_Win_free(MPI_Win *win) 7
int MPI_Win_get_group(MPI_Win win, MPI_Group *group) 8
int MPI_Win_get_info(MPI_Win win, MPI_Info *info_used) 9
int MPI_Win_lock(int lock_type, int rank, int assert, MPI_Win win) 10
int MPI_Win_lock_all(int assert, MPI_Win win) 11
int MPI_Win_post(MPI_Group group, int assert, MPI_Win win) 12
int MPI_Win_set_info(MPI_Win win, MPI_Info info) 13
int MPI_Win_shared_query(MPI_Win win, int rank, MPI_Aint *size, int *disp_unit, 14
    void *baseptr) 15
int MPI_Win_shared_query_c(MPI_Win win, int rank, MPI_Aint *size, 16
    MPI_Aint *disp_unit, void *baseptr) 17
int MPI_Win_start(MPI_Group group, int assert, MPI_Win win) 18
int MPI_Win_sync(MPI_Win win) 19
int MPI_Win_test(MPI_Win win, int *flag) 20
int MPI_Win_unlock(int rank, MPI_Win win) 21
int MPI_Win_unlock_all(MPI_Win win) 22
int MPI_Win_wait(MPI_Win win) 23
int MPI_Win_wait(MPI_Win win) 24
int MPI_Win_wait(MPI_Win win) 25
int MPI_Win_wait(MPI_Win win) 26
int MPI_Win_wait(MPI_Win win) 27
int MPI_Win_wait(MPI_Win win) 28
int MPI_Win_wait(MPI_Win win) 29
int MPI_Win_wait(MPI_Win win) 30
int MPI_Win_wait(MPI_Win win) 31
A.3.11 External Interfaces C Bindings 32
int MPI_Grequest_complete(MPI_Request request) 33
int MPI_Grequest_start(MPI_Grequest_query_function *query_fn, 34
    MPI_Grequest_free_function *free_fn, 35
    MPI_Grequest_cancel_function *cancel_fn, void *extra_state, 36
    MPI_Request *request) 37
int MPI_Status_set_cancelled(MPI_Status *status, int flag) 38
int MPI_Status_set_elements(MPI_Status *status, MPI_Datatype datatype, 39
    int count) 40
int MPI_Status_set_elements_c(MPI_Status *status, MPI_Datatype datatype, 41
    MPI_Count count) 42
int MPI_Status_set_error(MPI_Status *status, int err) 43
int MPI_Status_set_source(MPI_Status *status, int source) 44
int MPI_Status_set_tag(MPI_Status *status, int tag) 45
int MPI_Status_set_tag(MPI_Status *status, int tag) 46
int MPI_Status_set_tag(MPI_Status *status, int tag) 47
int MPI_Status_set_tag(MPI_Status *status, int tag) 48

```

A.3.12 I/O C Bindings

```
1  int MPI_CONVERSION_FN_NULL(void *userbuf, MPI_Datatype datatype, int count,
2                               void *filebuf, MPI_Offset position, void *extra_state)
3
4  int MPI_CONVERSION_FN_NULL_C(void *userbuf, MPI_Datatype datatype,
5                               MPI_Count count, void *filebuf, MPI_Offset position,
6                               void *extra_state)
7
8  int MPI_File_close(MPI_File *fh)
9
10 int MPI_File_delete(const char *filename, MPI_Info info)
11
12 int MPI_File_get_amode(MPI_File fh, int *amode)
13
14 int MPI_File_get_atomicity(MPI_File fh, int *flag)
15
16 int MPI_File_get_byte_offset(MPI_File fh, MPI_Offset offset, MPI_Offset *disp)
17
18 int MPI_File_get_group(MPI_File fh, MPI_Group *group)
19
20 int MPI_File_get_info(MPI_File fh, MPI_Info *info_used)
21
22 int MPI_File_get_position(MPI_File fh, MPI_Offset *offset)
23
24 int MPI_File_get_position_shared(MPI_File fh, MPI_Offset *offset)
25
26 int MPI_File_get_size(MPI_File fh, MPI_Offset *size)
27
28 int MPI_File_get_type_extent(MPI_File fh, MPI_Datatype datatype,
29                               MPI_Aint *extent)
30
31 int MPI_File_get_type_extent_c(MPI_File fh, MPI_Datatype datatype,
32                               MPI_Count *extent)
33
34 int MPI_File_get_view(MPI_File fh, MPI_Offset *disp, MPI_Datatype *etype,
35                               MPI_Datatype *filetype, char *datarep)
36
37 int MPI_File_iread(MPI_File fh, void *buf, int count, MPI_Datatype datatype,
38                               MPI_Request *request)
39
40 int MPI_File_iread_all(MPI_File fh, void *buf, int count,
41                               MPI_Datatype datatype, MPI_Request *request)
42
43 int MPI_File_iread_all_c(MPI_File fh, void *buf, MPI_Count count,
44                               MPI_Datatype datatype, MPI_Request *request)
45
46 int MPI_File_iread_at(MPI_File fh, MPI_Offset offset, void *buf, int count,
47                               MPI_Datatype datatype, MPI_Request *request)
48
49 int MPI_File_iread_at_all(MPI_File fh, MPI_Offset offset, void *buf, int count,
50                               MPI_Datatype datatype, MPI_Request *request)
51
52 int MPI_File_iread_at_all_c(MPI_File fh, MPI_Offset offset, void *buf,
53                               MPI_Count count, MPI_Datatype datatype, MPI_Request *request)
54
55 int MPI_File_iread_at_c(MPI_File fh, MPI_Offset offset, void *buf,
56                               MPI_Count count, MPI_Datatype datatype, MPI_Request *request)
57
58 int MPI_File_iread_c(MPI_File fh, void *buf, MPI_Count count,
59                               MPI_Datatype datatype, MPI_Request *request)
```

```

int MPI_File_iread_shared(MPI_File fh, void *buf, int count,          1
                        MPI_Datatype datatype, MPI_Request *request)  2
int MPI_File_iread_shared_c(MPI_File fh, void *buf, MPI_Count count,  3
                        MPI_Datatype datatype, MPI_Request *request)  4
int MPI_File_iwrite(MPI_File fh, const void *buf, int count,        5
                   MPI_Datatype datatype, MPI_Request *request)    6
int MPI_File_iwrite_all(MPI_File fh, const void *buf, int count,    7
                       MPI_Datatype datatype, MPI_Request *request)  8
int MPI_File_iwrite_all_c(MPI_File fh, const void *buf, MPI_Count count,  9
                          MPI_Datatype datatype, MPI_Request *request) 10
int MPI_File_iwrite_at(MPI_File fh, MPI_Offset offset, const void *buf, 11
                      int count, MPI_Datatype datatype, MPI_Request *request) 12
int MPI_File_iwrite_at_all(MPI_File fh, MPI_Offset offset, const void *buf, 13
                          int count, MPI_Datatype datatype, MPI_Request *request) 14
int MPI_File_iwrite_at_all_c(MPI_File fh, MPI_Offset offset, const void *buf, 15
                             MPI_Count count, MPI_Datatype datatype, MPI_Request *request) 16
int MPI_File_iwrite_at_c(MPI_File fh, MPI_Offset offset, const void *buf, 17
                        MPI_Count count, MPI_Datatype datatype, MPI_Request *request) 18
int MPI_File_iwrite_c(MPI_File fh, const void *buf, MPI_Count count, 19
                     MPI_Datatype datatype, MPI_Request *request)    20
int MPI_File_iwrite_shared(MPI_File fh, const void *buf, int count, 21
                          MPI_Datatype datatype, MPI_Request *request) 22
int MPI_File_iwrite_shared_c(MPI_File fh, const void *buf, MPI_Count count, 23
                             MPI_Datatype datatype, MPI_Request *request) 24
int MPI_File_open(MPI_Comm comm, const char *filename, int amode,    25
                 MPI_Info info, MPI_File *fh)                       26
int MPI_File_preallocate(MPI_File fh, MPI_Offset size)              27
int MPI_File_read(MPI_File fh, void *buf, int count, MPI_Datatype datatype, 28
                 MPI_Status *status)                                29
int MPI_File_read_all(MPI_File fh, void *buf, int count, MPI_Datatype datatype, 30
                    MPI_Status *status)                             31
int MPI_File_read_all_begin(MPI_File fh, void *buf, int count,     32
                            MPI_Datatype datatype)                 33
int MPI_File_read_all_begin_c(MPI_File fh, void *buf, MPI_Count count, 34
                              MPI_Datatype datatype)               35
int MPI_File_read_all_c(MPI_File fh, void *buf, MPI_Count count,   36
                       MPI_Datatype datatype, MPI_Status *status) 37
int MPI_File_read_all_end(MPI_File fh, void *buf, MPI_Status *status) 38
int MPI_File_read_at(MPI_File fh, MPI_Offset offset, void *buf, int count, 39
                   MPI_Datatype datatype, MPI_Status *status)     40
int MPI_File_read_at_all(MPI_File fh, MPI_Offset offset, void *buf, int count, 41
                       MPI_Datatype datatype, MPI_Status *status) 42
int MPI_File_read_at_all_c(MPI_File fh, MPI_Offset offset, void *buf, int count, 43
                          MPI_Datatype datatype, MPI_Status *status) 44
int MPI_File_read_at_all_end(MPI_File fh, MPI_Offset offset, void *buf, int count, 45
                             MPI_Datatype datatype, MPI_Status *status) 46
int MPI_File_read_at_shared(MPI_File fh, MPI_Offset offset, void *buf, int count, 47
                           MPI_Datatype datatype, MPI_Status *status) 48

```

```
1 int MPI_File_read_at_all(MPI_File fh, MPI_Offset offset, void *buf, int count,
2     MPI_Datatype datatype, MPI_Status *status)
3
4 int MPI_File_read_at_all_begin(MPI_File fh, MPI_Offset offset, void *buf,
5     int count, MPI_Datatype datatype)
6
7 int MPI_File_read_at_all_begin_c(MPI_File fh, MPI_Offset offset, void *buf,
8     MPI_Count count, MPI_Datatype datatype)
9
10 int MPI_File_read_at_all_c(MPI_File fh, MPI_Offset offset, void *buf,
11     MPI_Count count, MPI_Datatype datatype, MPI_Status *status)
12
13 int MPI_File_read_at_all_end(MPI_File fh, void *buf, MPI_Status *status)
14
15 int MPI_File_read_at_c(MPI_File fh, MPI_Offset offset, void *buf,
16     MPI_Count count, MPI_Datatype datatype, MPI_Status *status)
17
18 int MPI_File_read_c(MPI_File fh, void *buf, MPI_Count count,
19     MPI_Datatype datatype, MPI_Status *status)
20
21 int MPI_File_read_ordered(MPI_File fh, void *buf, int count,
22     MPI_Datatype datatype, MPI_Status *status)
23
24 int MPI_File_read_ordered_begin(MPI_File fh, void *buf, int count,
25     MPI_Datatype datatype)
26
27 int MPI_File_read_ordered_begin_c(MPI_File fh, void *buf, MPI_Count count,
28     MPI_Datatype datatype)
29
30 int MPI_File_read_ordered_c(MPI_File fh, void *buf, MPI_Count count,
31     MPI_Datatype datatype, MPI_Status *status)
32
33 int MPI_File_read_ordered_end(MPI_File fh, void *buf, MPI_Status *status)
34
35 int MPI_File_read_shared(MPI_File fh, void *buf, int count,
36     MPI_Datatype datatype, MPI_Status *status)
37
38 int MPI_File_read_shared_c(MPI_File fh, void *buf, MPI_Count count,
39     MPI_Datatype datatype, MPI_Status *status)
40
41 int MPI_File_seek(MPI_File fh, MPI_Offset offset, int whence)
42
43 int MPI_File_seek_shared(MPI_File fh, MPI_Offset offset, int whence)
44
45 int MPI_File_set_atomicity(MPI_File fh, int flag)
46
47 int MPI_File_set_info(MPI_File fh, MPI_Info info)
48
49 int MPI_File_set_size(MPI_File fh, MPI_Offset size)
50
51 int MPI_File_set_view(MPI_File fh, MPI_Offset disp, MPI_Datatype etype,
52     MPI_Datatype filetype, const char *datarep, MPI_Info info)
53
54 int MPI_File_sync(MPI_File fh)
55
56 int MPI_File_write(MPI_File fh, const void *buf, int count,
57     MPI_Datatype datatype, MPI_Status *status)
58
59 int MPI_File_write_all(MPI_File fh, const void *buf, int count,
60     MPI_Datatype datatype, MPI_Status *status)
```

```
int MPI_File_write_all_begin(MPI_File fh, const void *buf, int count,      1
                             MPI_Datatype datatype)                        2
int MPI_File_write_all_begin_c(MPI_File fh, const void *buf, MPI_Count count, 3
                               MPI_Datatype datatype)                    4
int MPI_File_write_all_c(MPI_File fh, const void *buf, MPI_Count count,    5
                          MPI_Datatype datatype, MPI_Status *status)      6
int MPI_File_write_all_end(MPI_File fh, const void *buf, MPI_Status *status) 7
int MPI_File_write_at(MPI_File fh, MPI_Offset offset, const void *buf,     8
                      int count, MPI_Datatype datatype, MPI_Status *status) 9
int MPI_File_write_at_all(MPI_File fh, MPI_Offset offset, const void *buf, 10
                          int count, MPI_Datatype datatype, MPI_Status *status) 11
int MPI_File_write_at_all_begin(MPI_File fh, MPI_Offset offset,            12
                                const void *buf, int count, MPI_Datatype datatype) 13
int MPI_File_write_at_all_begin_c(MPI_File fh, MPI_Offset offset,          14
                                   const void *buf, MPI_Count count, MPI_Datatype datatype) 15
int MPI_File_write_at_all_c(MPI_File fh, MPI_Offset offset, const void *buf, 16
                             MPI_Count count, MPI_Datatype datatype, MPI_Status *status) 17
int MPI_File_write_at_all_end(MPI_File fh, const void *buf, MPI_Status *status) 18
int MPI_File_write_at_c(MPI_File fh, MPI_Offset offset, const void *buf,   19
                        MPI_Count count, MPI_Datatype datatype, MPI_Status *status) 20
int MPI_File_write_c(MPI_File fh, const void *buf, MPI_Count count,       21
                     MPI_Datatype datatype, MPI_Status *status)          22
int MPI_File_write_ordered(MPI_File fh, const void *buf, int count,       23
                            MPI_Datatype datatype, MPI_Status *status)    24
int MPI_File_write_ordered_begin(MPI_File fh, const void *buf, int count,  25
                                  MPI_Datatype datatype)                  26
int MPI_File_write_ordered_begin_c(MPI_File fh, const void *buf,          27
                                    MPI_Count count, MPI_Datatype datatype) 28
int MPI_File_write_ordered_c(MPI_File fh, const void *buf, MPI_Count count, 29
                              MPI_Datatype datatype, MPI_Status *status) 30
int MPI_File_write_ordered_end(MPI_File fh, const void *buf,              31
                               MPI_Status *status)                       32
int MPI_File_write_shared(MPI_File fh, const void *buf, int count,        33
                          MPI_Datatype datatype, MPI_Status *status)      34
int MPI_File_write_shared_c(MPI_File fh, const void *buf, MPI_Count count, 35
                             MPI_Datatype datatype, MPI_Status *status)    36
int MPI_Register_datarep(const char *datarep,                             37
                        MPI_Datarep_conversion_function *read_conversion_fn, 38
                        MPI_Datarep_conversion_function *write_conversion_fn, 39
                        MPI_Status *status)                                40
int MPI_File_write_shared_c(MPI_File fh, const void *buf, MPI_Count count, 41
                             MPI_Datatype datatype, MPI_Status *status)    42
int MPI_File_write_shared_c(MPI_File fh, const void *buf, MPI_Count count, 43
                             MPI_Datatype datatype, MPI_Status *status)    44
int MPI_Register_datarep(const char *datarep,                             45
                        MPI_Datarep_conversion_function *read_conversion_fn, 46
                        MPI_Datarep_conversion_function *write_conversion_fn, 47
                        MPI_Status *status)                                48
```

```

1         MPI_Datarep_extent_function *dtype_file_extent_fn,
2         void *extra_state)
3
4 int MPI_Register_datarep_c(const char *datarep,
5         MPI_Datarep_conversion_function_c *read_conversion_fn,
6         MPI_Datarep_conversion_function_c *write_conversion_fn,
7         MPI_Datarep_extent_function *dtype_file_extent_fn,
8         void *extra_state)
9

```

10 A.3.13 Language Bindings C Bindings

```

11 MPI_Fint MPI_Comm_c2f(MPI_Comm comm)
12 MPI_Comm MPI_Comm_f2c(MPI_Fint comm)
13 MPI_Fint MPI_Errhandler_c2f(MPI_Errhandler errhandler)
14 MPI_Errhandler MPI_Errhandler_f2c(MPI_Fint errhandler)
15 MPI_Fint MPI_File_c2f(MPI_File file)
16 MPI_File MPI_File_f2c(MPI_Fint file)
17 MPI_Fint MPI_Group_c2f(MPI_Group group)
18 MPI_Group MPI_Group_f2c(MPI_Fint group)
19 MPI_Fint MPI_Info_c2f(MPI_Info info)
20 MPI_Info MPI_Info_f2c(MPI_Fint info)
21 MPI_Fint MPI_Message_c2f(MPI_Message message)
22 MPI_Message MPI_Message_f2c(MPI_Fint message)
23 MPI_Fint MPI_Op_c2f(MPI_Op op)
24 MPI_Op MPI_Op_f2c(MPI_Fint op)
25 MPI_Fint MPI_Request_c2f(MPI_Request request)
26 MPI_Request MPI_Request_f2c(MPI_Fint request)
27 MPI_Fint MPI_Session_c2f(MPI_Session session)
28 MPI_Session MPI_Session_f2c(MPI_Fint session)
29
30 int MPI_Status_c2f(const MPI_Status *c_status, MPI_Fint *f_status)
31
32 int MPI_Status_c2f08(const MPI_Status *c_status, MPI_F08_status *f08_status)
33
34 int MPI_Status_f082c(const MPI_F08_status *f08_status, MPI_Status *c_status)
35
36 int MPI_Status_f082f(const MPI_F08_status *f08_status, MPI_Fint *f_status)
37
38 int MPI_Status_f2c(const MPI_Fint *f_status, MPI_Status *c_status)
39
40 int MPI_Status_f2f08(const MPI_Fint *f_status, MPI_F08_status *f08_status)
41
42 MPI_Fint MPI_Type_c2f(MPI_Datatype datatype)
43
44 int MPI_Type_create_f90_complex(int p, int r, MPI_Datatype *newtype)
45
46
47
48

```

```

int MPI_Type_create_f90_integer(int r, MPI_Datatype *newtype)           1
int MPI_Type_create_f90_real(int p, int r, MPI_Datatype *newtype)     2
MPI_Datatype MPI_Type_f2c(MPI_Fint datatype)                          4
int MPI_Type_match_size(int typeclass, int size, MPI_Datatype *datatype) 5
MPI_Fint MPI_Win_c2f(MPI_Win win)                                     7
MPI_Win MPI_Win_f2c(MPI_Fint win)                                     8

```

A.3.14 Tools / Profiling Interface C Bindings

```

int MPI_Pcontrol(const int level, . . . )                             12

```

A.3.15 Tools / MPI Tool Information Interface C Bindings

```

int MPI_T_category_changed(int *update_number)                       16
int MPI_T_category_get_categories(int cat_index, int len, int indices[]) 18
int MPI_T_category_get_cvars(int cat_index, int len, int indices[])    20
int MPI_T_category_get_events(int cat_index, int len, int indices[])    21
int MPI_T_category_get_index(const char *name, int *cat_index)         22
int MPI_T_category_get_info(int cat_index, char *name, int *name_len,    24
                           char *desc, int *desc_len, int *num_cvars, int *num_pvars,
                           int *num_categories)                       26
int MPI_T_category_get_num(int *num_cat)                               27
int MPI_T_category_get_num_events(int cat_index, int *num_events)       29
int MPI_T_category_get_pvars(int cat_index, int len, int indices[])     30
int MPI_T_cvar_get_index(const char *name, int *cvar_index)           32
int MPI_T_cvar_get_info(int cvar_index, char *name, int *name_len,      33
                       int *verbosity, MPI_Datatype *datatype, MPI_T_enum *enumtype,
                       char *desc, int *desc_len, int *bind, int *scope) 35
int MPI_T_cvar_get_num(int *num_cvar)                                  36
int MPI_T_cvar_handle_alloc(int cvar_index, void *obj_handle,          38
                           MPI_T_cvar_handle *handle, int *count)     39
int MPI_T_cvar_handle_free(MPI_T_cvar_handle *handle)                 40
int MPI_T_cvar_read(MPI_T_cvar_handle handle, void *buf)              42
int MPI_T_cvar_write(MPI_T_cvar_handle handle, const void *buf)       43
int MPI_T_enum_get_info(MPI_T_enum enumtype, int *num, char *name,     45
                       int *name_len)                                  46
int MPI_T_enum_get_item(MPI_T_enum enumtype, int index, int *value, char *name, 47
                       int *name_len)                                  48

```

```
1 int MPI_T_event_callback_get_info(MPI_T_event_registration event_registration,
2     MPI_T_cb_safety cb_safety, MPI_Info *info_used)
3
4 int MPI_T_event_callback_set_info(MPI_T_event_registration event_registration,
5     MPI_T_cb_safety cb_safety, MPI_Info info)
6
7 int MPI_T_event_copy(MPI_T_event_instance event_instance, void *buffer)
8
9 int MPI_T_event_get_index(const char *name, int *event_index)
10
11 int MPI_T_event_get_info(int event_index, char *name, int *name_len,
12     int *verbosity, MPI_Datatype array_of_datatypes[],
13     MPI_Aint array_of_displacements[], int *num_elements,
14     MPI_T_enum *enumtype, MPI_Info *info, char *desc, int *desc_len,
15     int *bind)
16
17 int MPI_T_event_get_num(int *num_events)
18
19 int MPI_T_event_get_source(MPI_T_event_instance event_instance,
20     int *source_index)
21
22 int MPI_T_event_get_timestamp(MPI_T_event_instance event_instance,
23     MPI_Count *event_timestamp)
24
25 int MPI_T_event_handle_alloc(int event_index, void *obj_handle, MPI_Info info,
26     MPI_T_event_registration *event_registration)
27
28 int MPI_T_event_handle_free(MPI_T_event_registration event_registration,
29     void *user_data, MPI_T_event_free_cb_function free_cb_function)
30
31 int MPI_T_event_handle_get_info(MPI_T_event_registration event_registration,
32     MPI_Info *info_used)
33
34 int MPI_T_event_handle_set_info(MPI_T_event_registration event_registration,
35     MPI_Info info)
36
37 int MPI_T_event_read(MPI_T_event_instance event_instance, int element_index,
38     void *buffer)
39
40 int MPI_T_event_register_callback(MPI_T_event_registration event_registration,
41     MPI_T_cb_safety cb_safety, MPI_Info info, void *user_data,
42     MPI_T_event_cb_function event_cb_function)
43
44 int MPI_T_event_set_dropped_handler(
45     MPI_T_event_registration event_registration,
46     MPI_T_event_dropped_cb_function dropped_cb_function)
47
48 int MPI_T_finalize(void)
49
50 int MPI_T_init_thread(int required, int *provided)
51
52 int MPI_T_pvar_get_index(const char *name, int var_class, int *pvar_index)
53
54 int MPI_T_pvar_get_info(int pvar_index, char *name, int *name_len,
55     int *verbosity, int *var_class, MPI_Datatype *datatype,
56     MPI_T_enum *enumtype, char *desc, int *desc_len, int *bind,
57     int *readonly, int *continuous, int *atomic)
58
59 int MPI_T_pvar_get_num(int *num_pvar)
```



```

int MPI_T_pvar_handle_alloc(MPI_T_pvar_session pe_session, int pvar_index,      1
                           void *obj_handle, MPI_T_pvar_handle *handle, int *count)      2
int MPI_T_pvar_handle_free(MPI_T_pvar_session pe_session,                      3
                           MPI_T_pvar_handle *handle)                            4
int MPI_T_pvar_read(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle,    5
                   void *buf)                                                  6
int MPI_T_pvar_readreset(MPI_T_pvar_session pe_session,                       7
                        MPI_T_pvar_handle handle, void *buf)                   8
int MPI_T_pvar_reset(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle)   9
int MPI_T_pvar_session_create(MPI_T_pvar_session *pe_session)                 10
int MPI_T_pvar_session_free(MPI_T_pvar_session *pe_session)                   11
int MPI_T_pvar_start(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle)  12
int MPI_T_pvar_stop(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle)  13
int MPI_T_pvar_write(MPI_T_pvar_session pe_session, MPI_T_pvar_handle handle,  14
                    const void *buf)                                          15
int MPI_T_source_get_info(int source_index, char *name, int *name_len,         16
                          char *desc, int *desc_len, MPI_T_source_order *ordering,
                          MPI_Count *ticks_per_second, MPI_Count *max_ticks,
                          MPI_Info *info)                                     17
int MPI_T_source_get_num(int *num_sources)                                     18
int MPI_T_source_get_timestamp(int source_index, MPI_Count *timestamp)         19
A.3.16 Deprecated C Bindings                                               20
int MPI_Attr_delete(MPI_Comm comm, int keyval)                               21
int MPI_Attr_get(MPI_Comm comm, int keyval, void *attribute_val, int *flag)   22
int MPI_Attr_put(MPI_Comm comm, int keyval, void *attribute_val)              23
int MPI_DUP_FN(MPI_Comm oldcomm, int keyval, void *extra_state,              24
               void *attribute_val_in, void *attribute_val_out, int *flag)    25
int MPI_Get_elements_x(const MPI_Status *status, MPI_Datatype datatype,       26
                      MPI_Count *count)                                       27
int MPI_Info_get(MPI_Info info, const char *key, int valuelen, char *value,   28
                 int *flag)                                                    29
int MPI_Info_get_valuelen(MPI_Info info, const char *key, int *valuelen,     30
                          int *flag)                                           31
int MPI_Keyval_create(MPI_Copy_function *copy_fn,                             32
                    MPI_Delete_function *delete_fn, int *keyval, void *extra_state) 33
int MPI_Keyval_free(int *keyval)                                              34

```

```
1 int MPI_NULL_COPY_FN(MPI_Comm oldcomm, int keyval, void *extra_state,  
2     void *attribute_val_in, void *attribute_val_out, int *flag)  
3  
4 int MPI_NULL_DELETE_FN(MPI_Comm comm, int keyval, void *attribute_val,  
5     void *extra_state)  
6  
7 int MPI_Status_set_elements_x(MPI_Status *status, MPI_Datatype datatype,  
8     MPI_Count count)  
9  
10 int MPI_Type_get_extent_x(MPI_Datatype datatype, MPI_Count *lb,  
11     MPI_Count *extent)  
12  
13 int MPI_Type_get_true_extent_x(MPI_Datatype datatype, MPI_Count *true_lb,  
14     MPI_Count *true_extent)  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48
```

A.4 Fortran 2008 Bindings with the mpi_f08 Module

A.4.1 Point-to-Point Communication Fortran 2008 Bindings

```

MPI_Bsend(buf, count, datatype, dest, tag, comm, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER, INTENT(IN) :: count, dest, tag
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Bsend(buf, count, datatype, dest, tag, comm, ierror) !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER, INTENT(IN) :: dest, tag
  TYPE(MPI_Comm), INTENT(IN) :: comm
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Bsend_init(buf, count, datatype, dest, tag, comm, request, ierror)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
  INTEGER, INTENT(IN) :: count, dest, tag
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Bsend_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
  TYPE(MPI_Datatype), INTENT(IN) :: datatype
  INTEGER, INTENT(IN) :: dest, tag
  TYPE(MPI_Comm), INTENT(IN) :: comm
  TYPE(MPI_Request), INTENT(OUT) :: request
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Buffer_attach(buffer, size, ierror)
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
  INTEGER, INTENT(IN) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Buffer_attach(buffer, size, ierror) !(_c)
  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Buffer_detach(buffer_addr, size, ierror)
  USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
  TYPE(C_PTR), INTENT(OUT) :: buffer_addr
  INTEGER, INTENT(OUT) :: size
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Buffer_detach(buffer_addr, size, ierror) !(_c)
2   USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
3   TYPE(C_PTR), INTENT(OUT) :: buffer_addr
4   INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
5   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
6
7 MPI_Buffer_flush(ierror)
8   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
9
10 MPI_Buffer_iflush(request, ierror)
11   TYPE(MPI_Request), INTENT(OUT) :: request
12   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13
14 MPI_Cancel(request, ierror)
15   TYPE(MPI_Request), INTENT(IN) :: request
16   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
17
18 MPI_Comm_attach_buffer(comm, buffer, size, ierror)
19   TYPE(MPI_Comm), INTENT(IN) :: comm
20   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
21   INTEGER, INTENT(IN) :: size
22   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24 MPI_Comm_attach_buffer(comm, buffer, size, ierror) !(_c)
25   TYPE(MPI_Comm), INTENT(IN) :: comm
26   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
27   INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: size
28   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
29
30 MPI_Comm_detach_buffer(comm, buffer_addr, size, ierror)
31   USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
32   TYPE(MPI_Comm), INTENT(IN) :: comm
33   TYPE(C_PTR), INTENT(OUT) :: buffer_addr
34   INTEGER, INTENT(OUT) :: size
35   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
37 MPI_Comm_detach_buffer(comm, buffer_addr, size, ierror) !(_c)
38   USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
39   TYPE(MPI_Comm), INTENT(IN) :: comm
40   TYPE(C_PTR), INTENT(OUT) :: buffer_addr
41   INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
42   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
43
44 MPI_Comm_flush_buffer(comm, ierror)
45   TYPE(MPI_Comm), INTENT(IN) :: comm
46   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
47
48 MPI_Comm_iflush_buffer(comm, request, ierror)
49   TYPE(MPI_Comm), INTENT(IN) :: comm
50   TYPE(MPI_Request), INTENT(OUT) :: request
51   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
52
53 MPI_Get_count(status, datatype, count, ierror)

```

```

TYPE(MPI_Status), INTENT(IN) :: status
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(OUT) :: count
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Get_count(status, datatype, count, ierror) !(_c)
TYPE(MPI_Status), INTENT(IN) :: status
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: count
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ibsend(buf, count, datatype, dest, tag, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, dest, tag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ibsend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: dest, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Improbe(source, tag, comm, flag, message, status, ierror)
INTEGER, INTENT(IN) :: source, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
LOGICAL, INTENT(OUT) :: flag
TYPE(MPI_Message), INTENT(OUT) :: message
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Irecv(buf, count, datatype, message, request, ierror)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Message), INTENT(INOUT) :: message
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Irecv(buf, count, datatype, message, request, ierror) !(_c)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Message), INTENT(INOUT) :: message
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```
1 MPI_Iprobe(source, tag, comm, flag, status, ierror)
2   INTEGER, INTENT(IN) :: source, tag
3   TYPE(MPI_Comm), INTENT(IN) :: comm
4   LOGICAL, INTENT(OUT) :: flag
5   TYPE(MPI_Status) :: status
6   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
7
8 MPI_Irecv(buf, count, datatype, source, tag, comm, request, ierror)
9   TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
10  INTEGER, INTENT(IN) :: count, source, tag
11  TYPE(MPI_Datatype), INTENT(IN) :: datatype
12  TYPE(MPI_Comm), INTENT(IN) :: comm
13  TYPE(MPI_Request), INTENT(OUT) :: request
14  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15
16 MPI_Irecv(buf, count, datatype, source, tag, comm, request, ierror) !(_c)
17  TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
18  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
19  TYPE(MPI_Datatype), INTENT(IN) :: datatype
20  INTEGER, INTENT(IN) :: source, tag
21  TYPE(MPI_Comm), INTENT(IN) :: comm
22  TYPE(MPI_Request), INTENT(OUT) :: request
23  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
24
25 MPI_Irsend(buf, count, datatype, dest, tag, comm, request, ierror)
26  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
27  INTEGER, INTENT(IN) :: count, dest, tag
28  TYPE(MPI_Datatype), INTENT(IN) :: datatype
29  TYPE(MPI_Comm), INTENT(IN) :: comm
30  TYPE(MPI_Request), INTENT(OUT) :: request
31  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Irsend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
34  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
35  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
36  TYPE(MPI_Datatype), INTENT(IN) :: datatype
37  INTEGER, INTENT(IN) :: dest, tag
38  TYPE(MPI_Comm), INTENT(IN) :: comm
39  TYPE(MPI_Request), INTENT(OUT) :: request
40  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
41
42 MPI_Isend(buf, count, datatype, dest, tag, comm, request, ierror)
43  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
44  INTEGER, INTENT(IN) :: count, dest, tag
45  TYPE(MPI_Datatype), INTENT(IN) :: datatype
46  TYPE(MPI_Comm), INTENT(IN) :: comm
47  TYPE(MPI_Request), INTENT(OUT) :: request
48  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
49
50 MPI_Isend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
51  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
```

```

INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: dest, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Isendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
              recvtype, source, recvtag, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, dest, sendtag, recvcount, source, recvtag
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Isendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
              recvtype, source, recvtag, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Isendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
                     comm, request, ierror)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, dest, sendtag, source, recvtag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Isendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
                     comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Issend(buf, count, datatype, dest, tag, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, dest, tag

```

```
1     TYPE(MPI_Datatype), INTENT(IN) :: datatype
2     TYPE(MPI_Comm), INTENT(IN) :: comm
3     TYPE(MPI_Request), INTENT(OUT) :: request
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6 MPI_Issend(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
7     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
8     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
9     TYPE(MPI_Datatype), INTENT(IN) :: datatype
10    INTEGER, INTENT(IN) :: dest, tag
11    TYPE(MPI_Comm), INTENT(IN) :: comm
12    TYPE(MPI_Request), INTENT(OUT) :: request
13    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14
15 MPI_Mprobe(source, tag, comm, message, status, ierror)
16    INTEGER, INTENT(IN) :: source, tag
17    TYPE(MPI_Comm), INTENT(IN) :: comm
18    TYPE(MPI_Message), INTENT(OUT) :: message
19    TYPE(MPI_Status) :: status
20    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Mrecv(buf, count, datatype, message, status, ierror)
23    TYPE(*), DIMENSION(..) :: buf
24    INTEGER, INTENT(IN) :: count
25    TYPE(MPI_Datatype), INTENT(IN) :: datatype
26    TYPE(MPI_Message), INTENT(INOUT) :: message
27    TYPE(MPI_Status) :: status
28    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
29
30 MPI_Mrecv(buf, count, datatype, message, status, ierror) !(_c)
31    TYPE(*), DIMENSION(..) :: buf
32    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
33    TYPE(MPI_Datatype), INTENT(IN) :: datatype
34    TYPE(MPI_Message), INTENT(INOUT) :: message
35    TYPE(MPI_Status) :: status
36    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Probe(source, tag, comm, status, ierror)
39    INTEGER, INTENT(IN) :: source, tag
40    TYPE(MPI_Comm), INTENT(IN) :: comm
41    TYPE(MPI_Status) :: status
42    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
43
44 MPI_Recv(buf, count, datatype, source, tag, comm, status, ierror)
45    TYPE(*), DIMENSION(..) :: buf
46    INTEGER, INTENT(IN) :: count, source, tag
47    TYPE(MPI_Datatype), INTENT(IN) :: datatype
48    TYPE(MPI_Comm), INTENT(IN) :: comm
49    TYPE(MPI_Status) :: status
50    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
51
52 MPI_Recv(buf, count, datatype, source, tag, comm, status, ierror) !(_c)
```



```

TYPE(*), DIMENSION(..) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: source, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Recv_init(buf, count, datatype, source, tag, comm, request, ierror)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, source, tag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Recv_init(buf, count, datatype, source, tag, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: source, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Request_free(request, ierror)
TYPE(MPI_Request), INTENT(INOUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Request_get_status(request, flag, status, ierror)
TYPE(MPI_Request), INTENT(IN) :: request
LOGICAL, INTENT(OUT) :: flag
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Request_get_status_all(count, array_of_requests, flag, array_of_statuses,
                           ierror)
INTEGER, INTENT(IN) :: count
TYPE(MPI_Request), INTENT(IN) :: array_of_requests(count)
LOGICAL, INTENT(OUT) :: flag
TYPE(MPI_Status) :: array_of_statuses(*)
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Request_get_status_any(count, array_of_requests, index, flag, status,
                           ierror)
INTEGER, INTENT(IN) :: count
TYPE(MPI_Request), INTENT(IN) :: array_of_requests(count)
INTEGER, INTENT(OUT) :: index
LOGICAL, INTENT(OUT) :: flag
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Request_get_status_some(incount, array_of_requests, outcount,
2     array_of_indices, array_of_statuses, ierror)
3     INTEGER, INTENT(IN) :: incount
4     TYPE(MPI_Request), INTENT(IN) :: array_of_requests(incount)
5     INTEGER, INTENT(OUT) :: outcount, array_of_indices(*)
6     TYPE(MPI_Status) :: array_of_statuses(*)
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8
9 MPI_Rsend(buf, count, datatype, dest, tag, comm, ierror)
10    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
11    INTEGER, INTENT(IN) :: count, dest, tag
12    TYPE(MPI_Datatype), INTENT(IN) :: datatype
13    TYPE(MPI_Comm), INTENT(IN) :: comm
14    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15
16 MPI_Rsend(buf, count, datatype, dest, tag, comm, ierror) !(_c)
17    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
18    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
19    TYPE(MPI_Datatype), INTENT(IN) :: datatype
20    INTEGER, INTENT(IN) :: dest, tag
21    TYPE(MPI_Comm), INTENT(IN) :: comm
22    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24 MPI_Rsend_init(buf, count, datatype, dest, tag, comm, request, ierror)
25    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
26    INTEGER, INTENT(IN) :: count, dest, tag
27    TYPE(MPI_Datatype), INTENT(IN) :: datatype
28    TYPE(MPI_Comm), INTENT(IN) :: comm
29    TYPE(MPI_Request), INTENT(OUT) :: request
30    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Rsend_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
33    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
34    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
35    TYPE(MPI_Datatype), INTENT(IN) :: datatype
36    INTEGER, INTENT(IN) :: dest, tag
37    TYPE(MPI_Comm), INTENT(IN) :: comm
38    TYPE(MPI_Request), INTENT(OUT) :: request
39    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41 MPI_Send(buf, count, datatype, dest, tag, comm, ierror)
42    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
43    INTEGER, INTENT(IN) :: count, dest, tag
44    TYPE(MPI_Datatype), INTENT(IN) :: datatype
45    TYPE(MPI_Comm), INTENT(IN) :: comm
46    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
47
48 MPI_Send(buf, count, datatype, dest, tag, comm, ierror) !(_c)
49    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
50    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
51    TYPE(MPI_Datatype), INTENT(IN) :: datatype

```

```

INTEGER, INTENT(IN) :: dest, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Send_init(buf, count, datatype, dest, tag, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count, dest, tag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Send_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: dest, tag
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Sendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
             recvtype, source, recvtag, comm, status, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcount, dest, sendtag, recvcount, source, recvtag
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Sendrecv(sendbuf, sendcount, sendtype, dest, sendtag, recvbuf, recvcount,
             recvtype, source, recvtag, comm, status, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Sendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
                    comm, status, ierror)
TYPE(*), DIMENSION(..) :: buf
INTEGER, INTENT(IN) :: count, dest, sendtag, source, recvtag
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Sendrecv_replace(buf, count, datatype, dest, sendtag, source, recvtag,
2     comm, status, ierror) !(_c)
3     TYPE(*), DIMENSION(..) :: buf
4     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
5     TYPE(MPI_Datatype), INTENT(IN) :: datatype
6     INTEGER, INTENT(IN) :: dest, sendtag, source, recvtag
7     TYPE(MPI_Comm), INTENT(IN) :: comm
8     TYPE(MPI_Status) :: status
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10
11 MPI_Session_attach_buffer(session, buffer, size, ierror)
12     TYPE(MPI_Session), INTENT(IN) :: session
13     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
14     INTEGER, INTENT(IN) :: size
15     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
16
17 MPI_Session_attach_buffer(session, buffer, size, ierror) !(_c)
18     TYPE(MPI_Session), INTENT(IN) :: session
19     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
20     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: size
21     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22
23 MPI_Session_detach_buffer(session, buffer_addr, size, ierror)
24     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
25     TYPE(MPI_Session), INTENT(IN) :: session
26     TYPE(C_PTR), INTENT(OUT) :: buffer_addr
27     INTEGER, INTENT(OUT) :: size
28     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
29
30 MPI_Session_detach_buffer(session, buffer_addr, size, ierror) !(_c)
31     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
32     TYPE(MPI_Session), INTENT(IN) :: session
33     TYPE(C_PTR), INTENT(OUT) :: buffer_addr
34     INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
35     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
37 MPI_Session_flush_buffer(session, ierror)
38     TYPE(MPI_Session), INTENT(IN) :: session
39     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41 MPI_Session_iflush_buffer(session, request, ierror)
42     TYPE(MPI_Session), INTENT(IN) :: session
43     TYPE(MPI_Request), INTENT(OUT) :: request
44     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
45
46 MPI_Ssend(buf, count, datatype, dest, tag, comm, ierror)
47     TYPE(*), DIMENSION(..), INTENT(IN) :: buf
48     INTEGER, INTENT(IN) :: count, dest, tag
49     TYPE(MPI_Datatype), INTENT(IN) :: datatype
50     TYPE(MPI_Comm), INTENT(IN) :: comm
51     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Ssend(buf, count, datatype, dest, tag, comm, ierror) !(_c)           1
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf                             2
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count                     3
  TYPE(MPI_Datatype), INTENT(IN) :: datatype                            4
  INTEGER, INTENT(IN) :: dest, tag                                     5
  TYPE(MPI_Comm), INTENT(IN) :: comm                                   6
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             7
                                                                           8
MPI_Ssend_init(buf, count, datatype, dest, tag, comm, request, ierror)  9
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf             10
  INTEGER, INTENT(IN) :: count, dest, tag                             11
  TYPE(MPI_Datatype), INTENT(IN) :: datatype                          12
  TYPE(MPI_Comm), INTENT(IN) :: comm                                  13
  TYPE(MPI_Request), INTENT(OUT) :: request                           14
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             15
                                                                           16
MPI_Ssend_init(buf, count, datatype, dest, tag, comm, request, ierror) !(_c) 16
  TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf             17
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count                     18
  TYPE(MPI_Datatype), INTENT(IN) :: datatype                          19
  INTEGER, INTENT(IN) :: dest, tag                                     20
  TYPE(MPI_Comm), INTENT(IN) :: comm                                   21
  TYPE(MPI_Request), INTENT(OUT) :: request                           22
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             23
                                                                           24
MPI_Start(request, ierror)                                              25
  TYPE(MPI_Request), INTENT(INOUT) :: request                          26
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             27
                                                                           28
MPI_Startall(count, array_of_requests, ierror)                          28
  INTEGER, INTENT(IN) :: count                                         29
  TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)         30
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             31
                                                                           32
MPI_Status_get_error(status, err, ierror)                                32
  TYPE(MPI_Status), INTENT(IN) :: status                               33
  INTEGER, INTENT(OUT) :: err                                          34
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             35
                                                                           36
MPI_Status_get_source(status, source, ierror)                            36
  TYPE(MPI_Status), INTENT(IN) :: status                               37
  INTEGER, INTENT(OUT) :: source                                       38
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             39
                                                                           40
MPI_Status_get_tag(status, tag, ierror)                                  40
  TYPE(MPI_Status), INTENT(IN) :: status                               41
  INTEGER, INTENT(OUT) :: tag                                          42
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                             43
                                                                           44
MPI_Test(request, flag, status, ierror)                                  44
  TYPE(MPI_Request), INTENT(INOUT) :: request                          45
  LOGICAL, INTENT(OUT) :: flag                                         46
                                                                           47
                                                                           48

```

```
1     TYPE(MPI_Status) :: status
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
3
4 MPI_Test_cancelled(status, flag, ierror)
5     TYPE(MPI_Status), INTENT(IN) :: status
6     LOGICAL, INTENT(OUT) :: flag
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8
9 MPI_Testall(count, array_of_requests, flag, array_of_statuses, ierror)
10    INTEGER, INTENT(IN) :: count
11    TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
12    LOGICAL, INTENT(OUT) :: flag
13    TYPE(MPI_Status) :: array_of_statuses(*)
14    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15
16 MPI_Testany(count, array_of_requests, index, flag, status, ierror)
17    INTEGER, INTENT(IN) :: count
18    TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
19    INTEGER, INTENT(OUT) :: index
20    LOGICAL, INTENT(OUT) :: flag
21    TYPE(MPI_Status) :: status
22    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24 MPI_Testsome(incount, array_of_requests, outcount, array_of_indices,
25             array_of_statuses, ierror)
26    INTEGER, INTENT(IN) :: incount
27    TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(incount)
28    INTEGER, INTENT(OUT) :: outcount, array_of_indices(*)
29    TYPE(MPI_Status) :: array_of_statuses(*)
30    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Wait(request, status, ierror)
33    TYPE(MPI_Request), INTENT(INOUT) :: request
34    TYPE(MPI_Status) :: status
35    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
37 MPI_Waitall(count, array_of_requests, array_of_statuses, ierror)
38    INTEGER, INTENT(IN) :: count
39    TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
40    TYPE(MPI_Status) :: array_of_statuses(*)
41    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Waitany(count, array_of_requests, index, status, ierror)
44    INTEGER, INTENT(IN) :: count
45    TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(count)
46    INTEGER, INTENT(OUT) :: index
47    TYPE(MPI_Status) :: status
48    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
49
50 MPI_Waitsome(incount, array_of_requests, outcount, array_of_indices,
51             array_of_statuses, ierror)
52    INTEGER, INTENT(IN) :: incount
```

```

TYPE(MPI_Request), INTENT(INOUT) :: array_of_requests(incount)      1
INTEGER, INTENT(OUT) :: outcount, array_of_indices(*)              2
TYPE(MPI_Status) :: array_of_statuses(*)                            3
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                            4
                                                                    5
                                                                    6

```

A.4.2 Partitioned Communication Fortran 2008 Bindings

```

MPI_Parrived(request, partition, flag, ierror)                      8
  TYPE(MPI_Request), INTENT(IN) :: request                          9
  INTEGER, INTENT(IN) :: partition                                 10
  LOGICAL, INTENT(OUT) :: flag                                    11
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        12
                                                                    13
MPI_Pready(partition, request, ierror)                              14
  INTEGER, INTENT(IN) :: partition                                 15
  TYPE(MPI_Request), INTENT(IN) :: request                        16
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        17
                                                                    18
MPI_Pready_list(length, array_of_partitions, request, ierror)      19
  INTEGER, INTENT(IN) :: length, array_of_partitions(length)     20
  TYPE(MPI_Request), INTENT(IN) :: request                        21
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        22
                                                                    23
MPI_Pready_range(partition_low, partition_high, request, ierror)  24
  INTEGER, INTENT(IN) :: partition_low, partition_high           25
  TYPE(MPI_Request), INTENT(IN) :: request                        26
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        27
                                                                    28
MPI_Precv_init(buf, partitions, count, datatype, source, tag, comm, info,
               request, ierror)                                     29
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf                      30
  INTEGER, INTENT(IN) :: partitions, source, tag                  31
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count              32
  TYPE(MPI_Datatype), INTENT(IN) :: datatype                     33
  TYPE(MPI_Comm), INTENT(IN) :: comm                              34
  TYPE(MPI_Info), INTENT(IN) :: info                             35
  TYPE(MPI_Request), INTENT(OUT) :: request                       36
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        37
                                                                    38
MPI_Psend_init(buf, partitions, count, datatype, dest, tag, comm, info,
               request, ierror)                                     39
  TYPE(*), DIMENSION(..), INTENT(IN) :: buf                      40
  INTEGER, INTENT(IN) :: partitions, dest, tag                    41
  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count              42
  TYPE(MPI_Datatype), INTENT(IN) :: datatype                     43
  TYPE(MPI_Comm), INTENT(IN) :: comm                              44
  TYPE(MPI_Info), INTENT(IN) :: info                             45
  TYPE(MPI_Request), INTENT(OUT) :: request                       46
  INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        47
                                                                    48

```

```

1  A.4.3 Datatypes Fortran 2008 Bindings
2
3  INTEGER(KIND=MPI_ADDRESS_KIND) MPI_Aint_add(base, disp)
4      INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: base, disp
5
6  INTEGER(KIND=MPI_ADDRESS_KIND) MPI_Aint_diff(addr1, addr2)
7      INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: addr1, addr2
8
9  MPI_Get_address(location, address, ierror)
10     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: location
11     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: address
12     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13
14 MPI_Get_elements(status, datatype, count, ierror)
15     TYPE(MPI_Status), INTENT(IN) :: status
16     TYPE(MPI_Datatype), INTENT(IN) :: datatype
17     INTEGER, INTENT(OUT) :: count
18     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19
20 MPI_Get_elements(status, datatype, count, ierror) !(_c)
21     TYPE(MPI_Status), INTENT(IN) :: status
22     TYPE(MPI_Datatype), INTENT(IN) :: datatype
23     INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: count
24     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
25
26 MPI_Pack(inbuf, incount, datatype, outbuf, outsize, position, comm, ierror)
27     TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
28     INTEGER, INTENT(IN) :: incount, outsize
29     TYPE(MPI_Datatype), INTENT(IN) :: datatype
30     TYPE(*), DIMENSION(..) :: outbuf
31     INTEGER, INTENT(INOUT) :: position
32     TYPE(MPI_Comm), INTENT(IN) :: comm
33     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
34
35 MPI_Pack(inbuf, incount, datatype, outbuf, outsize, position, comm, ierror)
36     !(_c)
37     TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
38     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incount, outsize
39     TYPE(MPI_Datatype), INTENT(IN) :: datatype
40     TYPE(*), DIMENSION(..) :: outbuf
41     INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position
42     TYPE(MPI_Comm), INTENT(IN) :: comm
43     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45 MPI_Pack_external(datarep, inbuf, incount, datatype, outbuf, outsize, position,
46     ierror)
47     CHARACTER(LEN=*), INTENT(IN) :: datarep
48     TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
49     INTEGER, INTENT(IN) :: incount
50     TYPE(MPI_Datatype), INTENT(IN) :: datatype
51     TYPE(*), DIMENSION(..) :: outbuf
52     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: outsize

```



```

INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(INOUT) :: position
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Pack_external(datarep, inbuf, incount, datatype, outbuf, outsize, position,
                 ierror) !(_c)
CHARACTER(LEN=*), INTENT(IN) :: datarep
TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incount, outsize
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(*), DIMENSION(..) :: outbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Pack_external_size(datarep, incount, datatype, size, ierror)
CHARACTER(LEN=*), INTENT(IN) :: datarep
INTEGER, INTENT(IN) :: incount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: size
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Pack_external_size(datarep, incount, datatype, size, ierror) !(_c)
CHARACTER(LEN=*), INTENT(IN) :: datarep
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Pack_size(incount, datatype, comm, size, ierror)
INTEGER, INTENT(IN) :: incount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(OUT) :: size
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Pack_size(incount, datatype, comm, size, ierror) !(_c)
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: incount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_commit(datatype, ierror)
TYPE(MPI_Datatype), INTENT(INOUT) :: datatype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_contiguous(count, oldtype, newtype, ierror)
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_contiguous(count, oldtype, newtype, ierror) !(_c)

```

```

1     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
2     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
3     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6 MPI_Type_create_darray(size, rank, ndims, array_of_gsizes, array_of_distribs,
7     array_of_dargs, array_of_psize, order, oldtype, newtype, ierror)
8     INTEGER, INTENT(IN) :: size, rank, ndims, array_of_gsizes(ndims),
9     array_of_distribs(ndims), array_of_dargs(ndims),
10    array_of_psize(ndims), order
11    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
12    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
13    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14
15 MPI_Type_create_darray(size, rank, ndims, array_of_gsizes, array_of_distribs,
16    array_of_dargs, array_of_psize, order, oldtype, newtype, ierror)
17    !(_c)
18    INTEGER, INTENT(IN) :: size, rank, ndims, array_of_distribs(ndims),
19    array_of_dargs(ndims), array_of_psize(ndims), order
20    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: array_of_gsizes(ndims)
21    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
22    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
23    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
24
25 MPI_Type_create_hindexed(count, array_of_blocklengths, array_of_displacements,
26    oldtype, newtype, ierror)
27    INTEGER, INTENT(IN) :: count, array_of_blocklengths(count)
28    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: array_of_displacements(count)
29    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
30    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
31    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Type_create_hindexed(count, array_of_blocklengths, array_of_displacements,
34    oldtype, newtype, ierror) !(_c)
35    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count,
36    array_of_blocklengths(count), array_of_displacements(count)
37    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
38    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
39    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41 MPI_Type_create_hindexed_block(count, blocklength, array_of_displacements,
42    oldtype, newtype, ierror)
43    INTEGER, INTENT(IN) :: count, blocklength
44    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: array_of_displacements(count)
45    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
46    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
47    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48
49 MPI_Type_create_hindexed_block(count, blocklength, array_of_displacements,
50    oldtype, newtype, ierror) !(_c)

```

```

INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength,
    array_of_displacements(count)
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_hvector(count, blocklength, stride, oldtype, newtype, ierror)
INTEGER, INTENT(IN) :: count, blocklength
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: stride
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_hvector(count, blocklength, stride, oldtype, newtype, ierror)
    !(_c)
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength, stride
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_indexed_block(count, blocklength, array_of_displacements,
    oldtype, newtype, ierror)
INTEGER, INTENT(IN) :: count, blocklength, array_of_displacements(count)
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_indexed_block(count, blocklength, array_of_displacements,
    oldtype, newtype, ierror) !(_c)
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength,
    array_of_displacements(count)
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_resized(oldtype, lb, extent, newtype, ierror)
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: lb, extent
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_resized(oldtype, lb, extent, newtype, ierror) !(_c)
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: lb, extent
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_create_struct(count, array_of_blocklengths, array_of_displacements,
    array_of_types, newtype, ierror)
INTEGER, INTENT(IN) :: count, array_of_blocklengths(count)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: array_of_displacements(count)
TYPE(MPI_Datatype), INTENT(IN) :: array_of_types(count)

```

```

1     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
3
4     MPI_Type_create_struct(count, array_of_blocklengths, array_of_displacements,
5         array_of_types, newtype, ierror) !(_c)
6     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count,
7         array_of_blocklengths(count), array_of_displacements(count)
8     TYPE(MPI_Datatype), INTENT(IN) :: array_of_types(count)
9     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12    MPI_Type_create_subarray(ndims, array_of_sizes, array_of_subsizes,
13        array_of_starts, order, oldtype, newtype, ierror)
14    INTEGER, INTENT(IN) :: ndims, array_of_sizes(ndims),
15        array_of_subsizes(ndims), array_of_starts(ndims), order
16    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
17    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
18    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19
20    MPI_Type_create_subarray(ndims, array_of_sizes, array_of_subsizes,
21        array_of_starts, order, oldtype, newtype, ierror) !(_c)
22    INTEGER, INTENT(IN) :: ndims, order
23    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: array_of_sizes(ndims),
24        array_of_subsizes(ndims), array_of_starts(ndims)
25    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
26    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
27    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
28
29    MPI_Type_dup(oldtype, newtype, ierror)
30    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
31    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
32    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
33
34    MPI_Type_free(datatype, ierror)
35    TYPE(MPI_Datatype), INTENT(INOUT) :: datatype
36    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38    MPI_Type_get_contents(datatype, max_integers, max_addresses, max_datatypes,
39        array_of_integers, array_of_addresses, array_of_datatypes,
40        ierror)
41    TYPE(MPI_Datatype), INTENT(IN) :: datatype
42    INTEGER, INTENT(IN) :: max_integers, max_addresses, max_datatypes
43    INTEGER, INTENT(OUT) :: array_of_integers(max_integers)
44    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) ::
45        array_of_addresses(max_addresses)
46    TYPE(MPI_Datatype), INTENT(OUT) :: array_of_datatypes(max_datatypes)
47    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48
49    MPI_Type_get_contents(datatype, max_integers, max_addresses, max_large_counts,
50        max_datatypes, array_of_integers, array_of_addresses,
51        array_of_large_counts, array_of_datatypes, ierror) !(_c)
52    TYPE(MPI_Datatype), INTENT(IN) :: datatype

```

```

INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: max_integers, max_addresses,
    max_large_counts, max_datatypes
INTEGER, INTENT(OUT) :: array_of_integers(max_integers)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) ::
    array_of_addresses(max_addresses)
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) ::
    array_of_large_counts(max_large_counts)
TYPE(MPI_Datatype), INTENT(OUT) :: array_of_datatypes(max_datatypes)
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_get_envelope(datatype, num_integers, num_addresses, num_datatypes,
    combiner, ierror)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(OUT) :: num_integers, num_addresses, num_datatypes,
    combiner
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_get_envelope(datatype, num_integers, num_addresses, num_large_counts,
    num_datatypes, combiner, ierror) !(_c)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: num_integers, num_addresses,
    num_large_counts, num_datatypes
INTEGER, INTENT(OUT) :: combiner
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_get_extent(datatype, lb, extent, ierror)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: lb, extent
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_get_extent(datatype, lb, extent, ierror) !(_c)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: lb, extent
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_get_true_extent(datatype, true_lb, true_extent, ierror)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: true_lb, true_extent
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_get_true_extent(datatype, true_lb, true_extent, ierror) !(_c)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: true_lb, true_extent
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Type_indexed(count, array_of_blocklengths, array_of_displacements, oldtype,
    newtype, ierror)
INTEGER, INTENT(IN) :: count, array_of_blocklengths(count),
    array_of_displacements(count)
TYPE(MPI_Datatype), INTENT(IN) :: oldtype
TYPE(MPI_Datatype), INTENT(OUT) :: newtype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Type_indexed(count, array_of_blocklengths, array_of_displacements, oldtype,
2     newtype, ierror) !(_c)
3     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count,
4     array_of_blocklengths(count), array_of_displacements(count)
5     TYPE(MPI_Datatype), INTENT(IN) :: oldtype
6     TYPE(MPI_Datatype), INTENT(OUT) :: newtype
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8
9 MPI_Type_size(datatype, size, ierror)
10    TYPE(MPI_Datatype), INTENT(IN) :: datatype
11    INTEGER, INTENT(OUT) :: size
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13
14 MPI_Type_size(datatype, size, ierror) !(_c)
15    TYPE(MPI_Datatype), INTENT(IN) :: datatype
16    INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
17    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
18
19 MPI_Type_vector(count, blocklength, stride, oldtype, newtype, ierror)
20    INTEGER, INTENT(IN) :: count, blocklength, stride
21    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
22    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
23    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
24
25 MPI_Type_vector(count, blocklength, stride, oldtype, newtype, ierror) !(_c)
26    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count, blocklength, stride
27    TYPE(MPI_Datatype), INTENT(IN) :: oldtype
28    TYPE(MPI_Datatype), INTENT(OUT) :: newtype
29    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
30
31 MPI_Unpack(inbuf, insize, position, outbuf, outcount, datatype, comm, ierror)
32    TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
33    INTEGER, INTENT(IN) :: insize, outcount
34    INTEGER, INTENT(INOUT) :: position
35    TYPE(*), DIMENSION(..) :: outbuf
36    TYPE(MPI_Datatype), INTENT(IN) :: datatype
37    TYPE(MPI_Comm), INTENT(IN) :: comm
38    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
39
40 MPI_Unpack(inbuf, insize, position, outbuf, outcount, datatype, comm, ierror)
41    !(_c)
42    TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
43    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: insize, outcount
44    INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position
45    TYPE(*), DIMENSION(..) :: outbuf
46    TYPE(MPI_Datatype), INTENT(IN) :: datatype
47    TYPE(MPI_Comm), INTENT(IN) :: comm
48    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
49
50 MPI_Unpack_external(datarep, inbuf, insize, position, outbuf, outcount,
51     datatype, ierror)
52    CHARACTER(LEN=*), INTENT(IN) :: datarep

```

```

TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf           1
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: insize  2
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(INOUT) :: position  3
TYPE(*), DIMENSION(..) :: outbuf                    4
INTEGER, INTENT(IN) :: outcount                      5
TYPE(MPI_Datatype), INTENT(IN) :: datatype           6
INTEGER, OPTIONAL, INTENT(OUT) :: ierror            7
MPI_Unpack_external(datatype, inbuf, insize, position, outbuf, outcount, 8
                    datatype, ierror) !(_c)          9
CHARACTER(LEN=*), INTENT(IN) :: datatype            10
TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf        11
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: insize, outcount  12
INTEGER(KIND=MPI_COUNT_KIND), INTENT(INOUT) :: position  13
TYPE(*), DIMENSION(..) :: outbuf                  14
TYPE(MPI_Datatype), INTENT(IN) :: datatype         15
INTEGER, OPTIONAL, INTENT(OUT) :: ierror           16
MPI_Allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm, 17
              ierror)                               18
A.4.4 Collective Communication Fortran 2008 Bindings 19
MPI_Allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm, 20
              ierror)                               21
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf      22
INTEGER, INTENT(IN) :: sendcount, recvcount        23
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype  24
TYPE(*), DIMENSION(..) :: recvbuf                 25
TYPE(MPI_Comm), INTENT(IN) :: comm                 26
INTEGER, OPTIONAL, INTENT(OUT) :: ierror           27
MPI_Allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm, 28
              ierror) !(_c)                         29
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf      30
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount  31
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype  32
TYPE(*), DIMENSION(..) :: recvbuf                 33
TYPE(MPI_Comm), INTENT(IN) :: comm                 34
INTEGER, OPTIONAL, INTENT(OUT) :: ierror           35
MPI_Allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, 36
                  comm, info, request, ierror)      37
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf  38
INTEGER, INTENT(IN) :: sendcount, recvcount        39
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype  40
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf    41
TYPE(MPI_Comm), INTENT(IN) :: comm                 42
TYPE(MPI_Info), INTENT(IN) :: info                 43
TYPE(MPI_Request), INTENT(OUT) :: request          44
INTEGER, OPTIONAL, INTENT(OUT) :: ierror           45

```

```

1 MPI_Allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
2                   comm, info, request, ierror) !(_c)
3     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
4     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
5     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
6     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
7     TYPE(MPI_Comm), INTENT(IN) :: comm
8     TYPE(MPI_Info), INTENT(IN) :: info
9     TYPE(MPI_Request), INTENT(OUT) :: request
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12 MPI_Allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
13              recvtype, comm, ierror)
14    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
15    INTEGER, INTENT(IN) :: sendcount, recvcounts(*), displs(*)
16    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
17    TYPE(*), DIMENSION(..) :: recvbuf
18    TYPE(MPI_Comm), INTENT(IN) :: comm
19    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21 MPI_Allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
22              recvtype, comm, ierror) !(_c)
23    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
24    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcounts(*)
25    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
26    TYPE(*), DIMENSION(..) :: recvbuf
27    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
28    TYPE(MPI_Comm), INTENT(IN) :: comm
29    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
30
31 MPI_Allgatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
32                   recvtype, comm, info, request, ierror)
33    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
34    INTEGER, INTENT(IN) :: sendcount
35    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
36    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
37    INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*), displs(*)
38    TYPE(MPI_Comm), INTENT(IN) :: comm
39    TYPE(MPI_Info), INTENT(IN) :: info
40    TYPE(MPI_Request), INTENT(OUT) :: request
41    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Allgatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
44                   recvtype, comm, info, request, ierror) !(_c)
45    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
46    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
47    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
48    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
49    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
50    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)

```



```

TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Allreduce(sendbuf, recvbuf, count, datatype, op, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
TYPE(*), DIMENSION(..) :: recvbuf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Allreduce(sendbuf, recvbuf, count, datatype, op, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
TYPE(*), DIMENSION(..) :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Allreduce_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Allreduce_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf

```

```

1     INTEGER, INTENT(IN) :: sendcount, recvcount
2     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
3     TYPE(*), DIMENSION(..) :: recvbuf
4     TYPE(MPI_Comm), INTENT(IN) :: comm
5     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
6
7 MPI_Alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm,
8             ierror) !(_c)
9     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
10    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
11    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
12    TYPE(*), DIMENSION(..) :: recvbuf
13    TYPE(MPI_Comm), INTENT(IN) :: comm
14    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15
16 MPI_Alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
17                  comm, info, request, ierror)
18    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
19    INTEGER, INTENT(IN) :: sendcount, recvcount
20    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
21    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
22    TYPE(MPI_Comm), INTENT(IN) :: comm
23    TYPE(MPI_Info), INTENT(IN) :: info
24    TYPE(MPI_Request), INTENT(OUT) :: request
25    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27 MPI_Alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
28                  comm, info, request, ierror) !(_c)
29    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
30    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
31    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
32    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
33    TYPE(MPI_Comm), INTENT(IN) :: comm
34    TYPE(MPI_Info), INTENT(IN) :: info
35    TYPE(MPI_Request), INTENT(OUT) :: request
36    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts,
39              rdispls, recvtype, comm, ierror)
40    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
41    INTEGER, INTENT(IN) :: sendcounts(*), sdispls(*), recvcounts(*), rdispls(*)
42    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
43    TYPE(*), DIMENSION(..) :: recvbuf
44    TYPE(MPI_Comm), INTENT(IN) :: comm
45    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
46
47 MPI_Alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts,
48              rdispls, recvtype, comm, ierror) !(_c)
49    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
50    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcounts(*)
51    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)

```

```

TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcoun
    rdispls, recvtype, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
    recvcoun
    rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcoun
    rdispls, recvtype, comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
    recvcoun
    rdispls(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
    rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcoun
    rdispls, recvtypes, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcounts(*), sdispls(*), recvcoun
    rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcoun
    rdispls, recvtypes, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcoun
    rdispls(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
2     recvcounts, rdispls, recvtypes, comm, info, request, ierror)
3     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
4     INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
5     recvcounts(*), rdispls(*)
6     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
7     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
8     TYPE(MPI_Comm), INTENT(IN) :: comm
9     TYPE(MPI_Info), INTENT(IN) :: info
10    TYPE(MPI_Request), INTENT(OUT) :: request
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_Alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
14     recvcounts, rdispls, recvtypes, comm, info, request, ierror)
15     !(_c)
16     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
17     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
18     recvcounts(*)
19     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
20     rdispls(*)
21     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
22     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
23     TYPE(MPI_Comm), INTENT(IN) :: comm
24     TYPE(MPI_Info), INTENT(IN) :: info
25     TYPE(MPI_Request), INTENT(OUT) :: request
26     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28 MPI_Barrier(comm, ierror)
29     TYPE(MPI_Comm), INTENT(IN) :: comm
30     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Barrier_init(comm, info, request, ierror)
33     TYPE(MPI_Comm), INTENT(IN) :: comm
34     TYPE(MPI_Info), INTENT(IN) :: info
35     TYPE(MPI_Request), INTENT(OUT) :: request
36     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Bcast(buffer, count, datatype, root, comm, ierror)
39     TYPE(*), DIMENSION(..) :: buffer
40     INTEGER, INTENT(IN) :: count, root
41     TYPE(MPI_Datatype), INTENT(IN) :: datatype
42     TYPE(MPI_Comm), INTENT(IN) :: comm
43     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45 MPI_Bcast(buffer, count, datatype, root, comm, ierror) !(_c)
46     TYPE(*), DIMENSION(..) :: buffer
47     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
48     TYPE(MPI_Datatype), INTENT(IN) :: datatype
49     INTEGER, INTENT(IN) :: root
50     TYPE(MPI_Comm), INTENT(IN) :: comm
51     INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Bcast_init(buffer, count, datatype, root, comm, info, request, ierror)      1
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer                          2
    INTEGER, INTENT(IN) :: count, root                                       3
    TYPE(MPI_Datatype), INTENT(IN) :: datatype                               4
    TYPE(MPI_Comm), INTENT(IN) :: comm                                        5
    TYPE(MPI_Info), INTENT(IN) :: info                                       6
    TYPE(MPI_Request), INTENT(OUT) :: request                                 7
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  8
                                                                              9
MPI_Bcast_init(buffer, count, datatype, root, comm, info, request, ierror)    10
    !(_c)                                                                      11
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer                          12
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count                       13
    TYPE(MPI_Datatype), INTENT(IN) :: datatype                               14
    INTEGER, INTENT(IN) :: root                                              15
    TYPE(MPI_Comm), INTENT(IN) :: comm                                        16
    TYPE(MPI_Info), INTENT(IN) :: info                                       17
    TYPE(MPI_Request), INTENT(OUT) :: request                                 18
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  19
                                                                              20
MPI_Exscan(sendbuf, recvbuf, count, datatype, op, comm, ierror)              21
    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf                          22
    TYPE(*), DIMENSION(..) :: recvbuf                                        23
    INTEGER, INTENT(IN) :: count                                             24
    TYPE(MPI_Datatype), INTENT(IN) :: datatype                               25
    TYPE(MPI_Op), INTENT(IN) :: op                                           26
    TYPE(MPI_Comm), INTENT(IN) :: comm                                        27
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  28
                                                                              29
MPI_Exscan(sendbuf, recvbuf, count, datatype, op, comm, ierror) !(_c)       30
    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf                          31
    TYPE(*), DIMENSION(..) :: recvbuf                                        32
    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count                       33
    TYPE(MPI_Datatype), INTENT(IN) :: datatype                               34
    TYPE(MPI_Op), INTENT(IN) :: op                                           35
    TYPE(MPI_Comm), INTENT(IN) :: comm                                        36
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  37
                                                                              38
MPI_Exscan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,  39
    ierror)                                                                      40
    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf            41
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf                        42
    INTEGER, INTENT(IN) :: count                                             43
    TYPE(MPI_Datatype), INTENT(IN) :: datatype                               44
    TYPE(MPI_Op), INTENT(IN) :: op                                           45
    TYPE(MPI_Comm), INTENT(IN) :: comm                                        46
    TYPE(MPI_Info), INTENT(IN) :: info                                       47
    TYPE(MPI_Request), INTENT(OUT) :: request                                 48
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Exscan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
2     ierror) !(_c)
3     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
4     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
5     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
6     TYPE(MPI_Datatype), INTENT(IN) :: datatype
7     TYPE(MPI_Op), INTENT(IN) :: op
8     TYPE(MPI_Comm), INTENT(IN) :: comm
9     TYPE(MPI_Info), INTENT(IN) :: info
10    TYPE(MPI_Request), INTENT(OUT) :: request
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_Gather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
14     comm, ierror)
15     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
16     INTEGER, INTENT(IN) :: sendcount, recvcount, root
17     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
18     TYPE(*), DIMENSION(..) :: recvbuf
19     TYPE(MPI_Comm), INTENT(IN) :: comm
20     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Gather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
23     comm, ierror) !(_c)
24     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
25     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
26     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
27     TYPE(*), DIMENSION(..) :: recvbuf
28     INTEGER, INTENT(IN) :: root
29     TYPE(MPI_Comm), INTENT(IN) :: comm
30     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Gather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
33     root, comm, info, request, ierror)
34     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
35     INTEGER, INTENT(IN) :: sendcount, recvcount, root
36     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
37     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
38     TYPE(MPI_Comm), INTENT(IN) :: comm
39     TYPE(MPI_Info), INTENT(IN) :: info
40     TYPE(MPI_Request), INTENT(OUT) :: request
41     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Gather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
44     root, comm, info, request, ierror) !(_c)
45     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
46     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
47     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
48     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
49     INTEGER, INTENT(IN) :: root
50     TYPE(MPI_Comm), INTENT(IN) :: comm

```

```

TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Gatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
            recvtype, root, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcnts(*), displs(*), root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Gatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
            recvtype, root, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcnts(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Gatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
                recvtype, root, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcnts(*), displs(*)
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Gatherv_init(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
                recvtype, root, comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcnts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Iallgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
2               comm, request, ierror)
3     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
4     INTEGER, INTENT(IN) :: sendcount, recvcount
5     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
6     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
7     TYPE(MPI_Comm), INTENT(IN) :: comm
8     TYPE(MPI_Request), INTENT(OUT) :: request
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10
11 MPI_Iallgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
12               comm, request, ierror) !(_c)
13     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
14     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
15     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
16     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
17     TYPE(MPI_Comm), INTENT(IN) :: comm
18     TYPE(MPI_Request), INTENT(OUT) :: request
19     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21 MPI_Iallgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
22                recvtype, comm, request, ierror)
23     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
24     INTEGER, INTENT(IN) :: sendcount
25     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
26     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
27     INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*), displs(*)
28     TYPE(MPI_Comm), INTENT(IN) :: comm
29     TYPE(MPI_Request), INTENT(OUT) :: request
30     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Iallgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
33                recvtype, comm, request, ierror) !(_c)
34     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
35     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
36     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
37     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
38     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
39     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
40     TYPE(MPI_Comm), INTENT(IN) :: comm
41     TYPE(MPI_Request), INTENT(OUT) :: request
42     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
43
44 MPI_Iallreduce(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
45     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
46     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
47     INTEGER, INTENT(IN) :: count
48     TYPE(MPI_Datatype), INTENT(IN) :: datatype
49     TYPE(MPI_Op), INTENT(IN) :: op
50     TYPE(MPI_Comm), INTENT(IN) :: comm

```



```

TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Iallreduce(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
    !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ialltoall(sendbuf, sendcount, sendtype, recvbuf, recvcnt, recvtype, comm,
    request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcnt
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ialltoall(sendbuf, sendcount, sendtype, recvbuf, recvcnt, recvtype, comm,
    request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcnt
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ialltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcnts,
    rdispls, recvtype, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
    recvcnts(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Ialltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcnts,
    rdispls, recvtype, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
    recvcnts(*)

```

```

1     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
2         rdispls(*)
3     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
4     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
5     TYPE(MPI_Comm), INTENT(IN) :: comm
6     TYPE(MPI_Request), INTENT(OUT) :: request
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8
9 MPI_Ialltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcoun
10         rdispls, recvtypes, comm, request, ierror)
11     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
12     INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
13         recvcoun
14     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
15     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
16     TYPE(MPI_Comm), INTENT(IN) :: comm
17     TYPE(MPI_Request), INTENT(OUT) :: request
18     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19 MPI_Ialltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, recvcoun
20         rdispls, recvtypes, comm, request, ierror) !(_c)
21     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
22     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
23         recvcoun
24     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
25         rdispls(*)
26     TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
27     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
28     TYPE(MPI_Comm), INTENT(IN) :: comm
29     TYPE(MPI_Request), INTENT(OUT) :: request
30     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31 MPI_Ibarrier(comm, request, ierror)
32     TYPE(MPI_Comm), INTENT(IN) :: comm
33     TYPE(MPI_Request), INTENT(OUT) :: request
34     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
35 MPI_Ibcast(buffer, count, datatype, root, comm, request, ierror)
36     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
37     INTEGER, INTENT(IN) :: count, root
38     TYPE(MPI_Datatype), INTENT(IN) :: datatype
39     TYPE(MPI_Comm), INTENT(IN) :: comm
40     TYPE(MPI_Request), INTENT(OUT) :: request
41     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42 MPI_Ibcast(buffer, count, datatype, root, comm, request, ierror) !(_c)
43     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buffer
44     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
45     TYPE(MPI_Datatype), INTENT(IN) :: datatype
46     INTEGER, INTENT(IN) :: root
47     TYPE(MPI_Comm), INTENT(IN) :: comm

```

```

TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Iexscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Iexscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Igather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
            comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount, root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Igather(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
            comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Igatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts, displs,
            recvtype, root, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf

```

```

1     INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*), displs(*)
2     TYPE(MPI_Comm), INTENT(IN) :: comm
3     TYPE(MPI_Request), INTENT(OUT) :: request
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6 MPI_Igatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, displs,
7             recvtype, root, comm, request, ierror) !(_c)
8     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
9     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
10    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
11    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
12    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
13    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
14    INTEGER, INTENT(IN) :: root
15    TYPE(MPI_Comm), INTENT(IN) :: comm
16    TYPE(MPI_Request), INTENT(OUT) :: request
17    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
18
19 MPI_Ireduce(sendbuf, recvbuf, count, datatype, op, root, comm, request, ierror)
20    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
21    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
22    INTEGER, INTENT(IN) :: count, root
23    TYPE(MPI_Datatype), INTENT(IN) :: datatype
24    TYPE(MPI_Op), INTENT(IN) :: op
25    TYPE(MPI_Comm), INTENT(IN) :: comm
26    TYPE(MPI_Request), INTENT(OUT) :: request
27    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
28
29 MPI_Ireduce(sendbuf, recvbuf, count, datatype, op, root, comm, request, ierror)
30    !(_c)
31    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
32    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
33    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
34    TYPE(MPI_Datatype), INTENT(IN) :: datatype
35    TYPE(MPI_Op), INTENT(IN) :: op
36    INTEGER, INTENT(IN) :: root
37    TYPE(MPI_Comm), INTENT(IN) :: comm
38    TYPE(MPI_Request), INTENT(OUT) :: request
39    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41 MPI_Ireduce_scatter(sendbuf, recvbuf, recvcounts, datatype, op, comm, request,
42                    ierror)
43    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
44    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
45    INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
46    TYPE(MPI_Datatype), INTENT(IN) :: datatype
47    TYPE(MPI_Op), INTENT(IN) :: op
48    TYPE(MPI_Comm), INTENT(IN) :: comm
49    TYPE(MPI_Request), INTENT(OUT) :: request
50    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

MPI_Ireduce_scatter(sendbuf, recvbuf, recvcnt, datatype, op, comm, request,
                    ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcnt(*)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Ireduce_scatter_block(sendbuf, recvbuf, recvcnt, datatype, op, comm,
                          request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: recvcnt
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Ireduce_scatter_block(sendbuf, recvbuf, recvcnt, datatype, op, comm,
                          request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcnt
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Iscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

MPI_Iscan(sendbuf, recvbuf, count, datatype, op, comm, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm

```

```

1     TYPE(MPI_Request), INTENT(OUT) :: request
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
3
4 MPI_Isscatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
5               comm, request, ierror)
6     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
7     INTEGER, INTENT(IN) :: sendcount, recvcount, root
8     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
9     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
10    TYPE(MPI_Comm), INTENT(IN) :: comm
11    TYPE(MPI_Request), INTENT(OUT) :: request
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13
14 MPI_Isscatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
15               comm, request, ierror) !(_c)
16    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
17    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
18    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
19    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
20    INTEGER, INTENT(IN) :: root
21    TYPE(MPI_Comm), INTENT(IN) :: comm
22    TYPE(MPI_Request), INTENT(OUT) :: request
23    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
24
25 MPI_Isscatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
26               recvtype, root, comm, request, ierror)
27    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
28    INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), displs(*)
29    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
30    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
31    INTEGER, INTENT(IN) :: recvcount, root
32    TYPE(MPI_Comm), INTENT(IN) :: comm
33    TYPE(MPI_Request), INTENT(OUT) :: request
34    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
35
36 MPI_Isscatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
37               recvtype, root, comm, request, ierror) !(_c)
38    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
39    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*)
40    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
41    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
42    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
43    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcount
44    INTEGER, INTENT(IN) :: root
45    TYPE(MPI_Comm), INTENT(IN) :: comm
46    TYPE(MPI_Request), INTENT(OUT) :: request
47    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48
49 MPI_Op_commutative(op, commute, ierror)
50    TYPE(MPI_Op), INTENT(IN) :: op

```

```

LOGICAL, INTENT(OUT) :: commute 1
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 2
MPI_Op_create(user_fn, commute, op, ierror) 3
PROCEDURE(MPI_User_function) :: user_fn 4
LOGICAL, INTENT(IN) :: commute 5
TYPE(MPI_Op), INTENT(OUT) :: op 6
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 7
MPI_Op_create_c(user_fn, commute, op, ierror) !(_c) 8
PROCEDURE(MPI_User_function_c) :: user_fn 9
LOGICAL, INTENT(IN) :: commute 10
TYPE(MPI_Op), INTENT(OUT) :: op 11
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 12
MPI_Op_free(op, ierror) 13
TYPE(MPI_Op), INTENT(INOUT) :: op 14
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 15
MPI_Reduce(sendbuf, recvbuf, count, datatype, op, root, comm, ierror) 16
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf 17
TYPE(*), DIMENSION(..) :: recvbuf 18
INTEGER, INTENT(IN) :: count, root 19
TYPE(MPI_Datatype), INTENT(IN) :: datatype 20
TYPE(MPI_Op), INTENT(IN) :: op 21
TYPE(MPI_Comm), INTENT(IN) :: comm 22
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 23
MPI_Reduce(sendbuf, recvbuf, count, datatype, op, root, comm, ierror) !(_c) 24
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf 25
TYPE(*), DIMENSION(..) :: recvbuf 26
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count 27
TYPE(MPI_Datatype), INTENT(IN) :: datatype 28
TYPE(MPI_Op), INTENT(IN) :: op 29
INTEGER, INTENT(IN) :: root 30
TYPE(MPI_Comm), INTENT(IN) :: comm 31
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 32
MPI_Reduce_init(sendbuf, recvbuf, count, datatype, op, root, comm, info, 33
request, ierror) 34
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 35
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 36
INTEGER, INTENT(IN) :: count, root 37
TYPE(MPI_Datatype), INTENT(IN) :: datatype 38
TYPE(MPI_Op), INTENT(IN) :: op 39
TYPE(MPI_Comm), INTENT(IN) :: comm 40
TYPE(MPI_Info), INTENT(IN) :: info 41
TYPE(MPI_Request), INTENT(OUT) :: request 42
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 43
MPI_Reduce_init(sendbuf, recvbuf, count, datatype, op, root, comm, info, 44
request, ierror) !(_c) 45

```

```

1     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
2     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
3     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
4     TYPE(MPI_Datatype), INTENT(IN) :: datatype
5     TYPE(MPI_Op), INTENT(IN) :: op
6     INTEGER, INTENT(IN) :: root
7     TYPE(MPI_Comm), INTENT(IN) :: comm
8     TYPE(MPI_Info), INTENT(IN) :: info
9     TYPE(MPI_Request), INTENT(OUT) :: request
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12    MPI_Reduce_local(inbuf, inoutbuf, count, datatype, op, ierror)
13    TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
14    TYPE(*), DIMENSION(..) :: inoutbuf
15    INTEGER, INTENT(IN) :: count
16    TYPE(MPI_Datatype), INTENT(IN) :: datatype
17    TYPE(MPI_Op), INTENT(IN) :: op
18    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19
20    MPI_Reduce_local(inbuf, inoutbuf, count, datatype, op, ierror) !(_c)
21    TYPE(*), DIMENSION(..), INTENT(IN) :: inbuf
22    TYPE(*), DIMENSION(..) :: inoutbuf
23    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
24    TYPE(MPI_Datatype), INTENT(IN) :: datatype
25    TYPE(MPI_Op), INTENT(IN) :: op
26    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28    MPI_Reduce_scatter(sendbuf, recvbuf, recvcnt, datatype, op, comm, ierror)
29    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
30    TYPE(*), DIMENSION(..) :: recvbuf
31    INTEGER, INTENT(IN) :: recvcnt(*)
32    TYPE(MPI_Datatype), INTENT(IN) :: datatype
33    TYPE(MPI_Op), INTENT(IN) :: op
34    TYPE(MPI_Comm), INTENT(IN) :: comm
35    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
37    MPI_Reduce_scatter(sendbuf, recvbuf, recvcnt, datatype, op, comm, ierror)
38    !(_c)
39    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
40    TYPE(*), DIMENSION(..) :: recvbuf
41    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcnt(*)
42    TYPE(MPI_Datatype), INTENT(IN) :: datatype
43    TYPE(MPI_Op), INTENT(IN) :: op
44    TYPE(MPI_Comm), INTENT(IN) :: comm
45    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
46
47    MPI_Reduce_scatter_block(sendbuf, recvbuf, recvcnt, datatype, op, comm,
48    ierror)
49    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
50    TYPE(*), DIMENSION(..) :: recvbuf

```



```

INTEGER, INTENT(IN) :: recvcount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Reduce_scatter_block(sendbuf, recvbuf, recvcount, datatype, op, comm,
    ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
TYPE(*), DIMENSION(..) :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Reduce_scatter_block_init(sendbuf, recvbuf, recvcount, datatype, op, comm,
    info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: recvcount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Reduce_scatter_block_init(sendbuf, recvbuf, recvcount, datatype, op, comm,
    info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcount
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Reduce_scatter_init(sendbuf, recvbuf, recvcounts, datatype, op, comm, info,
    request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request

```

```

1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2
3 MPI_Reduce_scatter_init(sendbuf, recvbuf, recvcounts, datatype, op, comm, info,
4     request, ierror) !(_c)
5     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
6     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
7     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*)
8     TYPE(MPI_Datatype), INTENT(IN) :: datatype
9     TYPE(MPI_Op), INTENT(IN) :: op
10    TYPE(MPI_Comm), INTENT(IN) :: comm
11    TYPE(MPI_Info), INTENT(IN) :: info
12    TYPE(MPI_Request), INTENT(OUT) :: request
13    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14
15 MPI_Scan(sendbuf, recvbuf, count, datatype, op, comm, ierror)
16    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
17    TYPE(*), DIMENSION(..) :: recvbuf
18    INTEGER, INTENT(IN) :: count
19    TYPE(MPI_Datatype), INTENT(IN) :: datatype
20    TYPE(MPI_Op), INTENT(IN) :: op
21    TYPE(MPI_Comm), INTENT(IN) :: comm
22    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24 MPI_Scan(sendbuf, recvbuf, count, datatype, op, comm, ierror) !(_c)
25    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
26    TYPE(*), DIMENSION(..) :: recvbuf
27    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
28    TYPE(MPI_Datatype), INTENT(IN) :: datatype
29    TYPE(MPI_Op), INTENT(IN) :: op
30    TYPE(MPI_Comm), INTENT(IN) :: comm
31    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Scan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
34    ierror)
35    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
36    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
37    INTEGER, INTENT(IN) :: count
38    TYPE(MPI_Datatype), INTENT(IN) :: datatype
39    TYPE(MPI_Op), INTENT(IN) :: op
40    TYPE(MPI_Comm), INTENT(IN) :: comm
41    TYPE(MPI_Info), INTENT(IN) :: info
42    TYPE(MPI_Request), INTENT(OUT) :: request
43    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45 MPI_Scan_init(sendbuf, recvbuf, count, datatype, op, comm, info, request,
46    ierror) !(_c)
47    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
48    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
49    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
50    TYPE(MPI_Datatype), INTENT(IN) :: datatype

```

```

TYPE(MPI_Op), INTENT(IN) :: op
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Scatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
            comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount, root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Scatter(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root,
            comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Scatter_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
                root, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount, root
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Scatter_init(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype,
                root, comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
INTEGER, INTENT(IN) :: root
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Scatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
             recvtype, root, comm, ierror)

```

```

1     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
2     INTEGER, INTENT(IN) :: sendcounts(*), displs(*), recvcount, root
3     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
4     TYPE(*), DIMENSION(..) :: recvbuf
5     TYPE(MPI_Comm), INTENT(IN) :: comm
6     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
7
8     MPI_Scatterv(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
9                 recvtype, root, comm, ierror) !(_c)
10    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
11    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcount
12    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
13    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
14    TYPE(*), DIMENSION(..) :: recvbuf
15    INTEGER, INTENT(IN) :: root
16    TYPE(MPI_Comm), INTENT(IN) :: comm
17    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
18
19    MPI_Scatterv_init(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
20                    recvtype, root, comm, info, request, ierror)
21    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
22    INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), displs(*)
23    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
24    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
25    INTEGER, INTENT(IN) :: recvcount, root
26    TYPE(MPI_Comm), INTENT(IN) :: comm
27    TYPE(MPI_Info), INTENT(IN) :: info
28    TYPE(MPI_Request), INTENT(OUT) :: request
29    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
30
31    MPI_Scatterv_init(sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount,
32                    recvtype, root, comm, info, request, ierror) !(_c)
33    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
34    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*)
35    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*)
36    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
37    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
38    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: recvcount
39    INTEGER, INTENT(IN) :: root
40    TYPE(MPI_Comm), INTENT(IN) :: comm
41    TYPE(MPI_Info), INTENT(IN) :: info
42    TYPE(MPI_Request), INTENT(OUT) :: request
43    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45    MPI_Type_get_value_index(value_type, index_type, pair_type, ierror)
46    TYPE(MPI_Datatype), INTENT(IN) :: value_type, index_type
47    TYPE(MPI_Datatype), INTENT(OUT) :: pair_type
48    INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

A.4.5 Groups, Contexts, Communicators, and Caching Fortran 2008 Bindings

	1
	2
MPI_Comm_compare(comm1, comm2, result, ierror)	3
TYPE(MPI_Comm), INTENT(IN) :: comm1, comm2	4
INTEGER, INTENT(OUT) :: result	5
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	6
	7
MPI_Comm_create(comm, group, newcomm, ierror)	8
TYPE(MPI_Comm), INTENT(IN) :: comm	9
TYPE(MPI_Group), INTENT(IN) :: group	10
TYPE(MPI_Comm), INTENT(OUT) :: newcomm	11
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	12
	13
MPI_Comm_create_from_group(group, stringtag, info, errhandler, newcomm, ierror)	14
TYPE(MPI_Group), INTENT(IN) :: group	15
CHARACTER(LEN=*), INTENT(IN) :: stringtag	16
TYPE(MPI_Info), INTENT(IN) :: info	17
TYPE(MPI_Errhandler), INTENT(IN) :: errhandler	18
TYPE(MPI_Comm), INTENT(OUT) :: newcomm	19
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	20
	21
MPI_Comm_create_group(comm, group, tag, newcomm, ierror)	22
TYPE(MPI_Comm), INTENT(IN) :: comm	23
TYPE(MPI_Group), INTENT(IN) :: group	24
INTEGER, INTENT(IN) :: tag	25
TYPE(MPI_Comm), INTENT(OUT) :: newcomm	26
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	27
	28
MPI_Comm_create_keyval(comm_copy_attr_fn, comm_delete_attr_fn, comm_keyval,	29
extra_state, ierror)	30
PROCEDURE(MPI_Comm_copy_attr_function) :: comm_copy_attr_fn	31
PROCEDURE(MPI_Comm_delete_attr_function) :: comm_delete_attr_fn	32
INTEGER, INTENT(OUT) :: comm_keyval	33
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state	34
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	35
	36
MPI_Comm_delete_attr(comm, comm_keyval, ierror)	37
TYPE(MPI_Comm), INTENT(IN) :: comm	38
INTEGER, INTENT(IN) :: comm_keyval	39
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	40
	41
MPI_Comm_dup(comm, newcomm, ierror)	42
TYPE(MPI_Comm), INTENT(IN) :: comm	43
TYPE(MPI_Comm), INTENT(OUT) :: newcomm	44
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	45
	46
MPI_COMM_DUP_FN(oldcomm, comm_keyval, extra_state, attribute_val_in,	47
attribute_val_out, flag, ierror)	48
TYPE(MPI_Comm) :: oldcomm	
INTEGER :: comm_keyval, ierror	
INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,	
attribute_val_out	

```
1     LOGICAL :: flag
2
3 MPI_Comm_dup_with_info(comm, info, newcomm, ierror)
4     TYPE(MPI_Comm), INTENT(IN) :: comm
5     TYPE(MPI_Info), INTENT(IN) :: info
6     TYPE(MPI_Comm), INTENT(OUT) :: newcomm
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8
9 MPI_Comm_free(comm, ierror)
10    TYPE(MPI_Comm), INTENT(INOUT) :: comm
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_Comm_free_keyval(comm_keyval, ierror)
14    INTEGER, INTENT(INOUT) :: comm_keyval
15    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
16
17 MPI_Comm_get_attr(comm, comm_keyval, attribute_val, flag, ierror)
18    TYPE(MPI_Comm), INTENT(IN) :: comm
19    INTEGER, INTENT(IN) :: comm_keyval
20    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: attribute_val
21    LOGICAL, INTENT(OUT) :: flag
22    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24 MPI_Comm_get_info(comm, info_used, ierror)
25    TYPE(MPI_Comm), INTENT(IN) :: comm
26    TYPE(MPI_Info), INTENT(OUT) :: info_used
27    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
28
29 MPI_Comm_get_name(comm, comm_name, resultlen, ierror)
30    TYPE(MPI_Comm), INTENT(IN) :: comm
31    CHARACTER(LEN=MPI_MAX_OBJECT_NAME), INTENT(OUT) :: comm_name
32    INTEGER, INTENT(OUT) :: resultlen
33    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
34
35 MPI_Comm_group(comm, group, ierror)
36    TYPE(MPI_Comm), INTENT(IN) :: comm
37    TYPE(MPI_Group), INTENT(OUT) :: group
38    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
39
40 MPI_Comm_idup(comm, newcomm, request, ierror)
41    TYPE(MPI_Comm), INTENT(IN) :: comm
42    TYPE(MPI_Comm), INTENT(OUT), ASYNCHRONOUS :: newcomm
43    TYPE(MPI_Request), INTENT(OUT) :: request
44    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
45
46 MPI_Comm_idup_with_info(comm, info, newcomm, request, ierror)
47    TYPE(MPI_Comm), INTENT(IN) :: comm
48    TYPE(MPI_Info), INTENT(IN) :: info
49    TYPE(MPI_Comm), INTENT(OUT), ASYNCHRONOUS :: newcomm
50    TYPE(MPI_Request), INTENT(OUT) :: request
51    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
52
53 MPI_COMM_NULL_COPY_FN(oldcomm, comm_keyval, extra_state, attribute_val_in,
54    attribute_val_out, flag, ierror)
```

```

TYPE(MPI_Comm) :: oldcomm
INTEGER :: comm_keyval, ierror
INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
    attribute_val_out
LOGICAL :: flag
MPI_COMM_NULL_DELETE_FN(comm, comm_keyval, attribute_val, extra_state, ierror)
TYPE(MPI_Comm) :: comm
INTEGER :: comm_keyval, ierror
INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
MPI_Comm_rank(comm, rank, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(OUT) :: rank
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_remote_group(comm, group, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Group), INTENT(OUT) :: group
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_remote_size(comm, size, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(OUT) :: size
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_set_attr(comm, comm_keyval, attribute_val, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(IN) :: comm_keyval
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: attribute_val
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_set_info(comm, info, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_set_name(comm, comm_name, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
CHARACTER(LEN=*), INTENT(IN) :: comm_name
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_size(comm, size, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(OUT) :: size
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_split(comm, color, key, newcomm, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(IN) :: color, key
TYPE(MPI_Comm), INTENT(OUT) :: newcomm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Comm_split_type(comm, split_type, key, info, newcomm, ierror)

```

```
1     TYPE(MPI_Comm), INTENT(IN) :: comm
2     INTEGER, INTENT(IN) :: split_type, key
3     TYPE(MPI_Info), INTENT(IN) :: info
4     TYPE(MPI_Comm), INTENT(OUT) :: newcomm
5     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
6
7 MPI_Comm_test_inter(comm, flag, ierror)
8     TYPE(MPI_Comm), INTENT(IN) :: comm
9     LOGICAL, INTENT(OUT) :: flag
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12 MPI_Group_compare(group1, group2, result, ierror)
13    TYPE(MPI_Group), INTENT(IN) :: group1, group2
14    INTEGER, INTENT(OUT) :: result
15    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
16
17 MPI_Group_difference(group1, group2, newgroup, ierror)
18    TYPE(MPI_Group), INTENT(IN) :: group1, group2
19    TYPE(MPI_Group), INTENT(OUT) :: newgroup
20    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Group_excl(group, n, ranks, newgroup, ierror)
23    TYPE(MPI_Group), INTENT(IN) :: group
24    INTEGER, INTENT(IN) :: n, ranks(n)
25    TYPE(MPI_Group), INTENT(OUT) :: newgroup
26    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28 MPI_Group_free(group, ierror)
29    TYPE(MPI_Group), INTENT(INOUT) :: group
30    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Group_from_session_pset(session, pset_name, newgroup, ierror)
33    TYPE(MPI_Session), INTENT(IN) :: session
34    CHARACTER(LEN=*), INTENT(IN) :: pset_name
35    TYPE(MPI_Group), INTENT(OUT) :: newgroup
36    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Group_incl(group, n, ranks, newgroup, ierror)
39    TYPE(MPI_Group), INTENT(IN) :: group
40    INTEGER, INTENT(IN) :: n, ranks(n)
41    TYPE(MPI_Group), INTENT(OUT) :: newgroup
42    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
43
44 MPI_Group_intersection(group1, group2, newgroup, ierror)
45    TYPE(MPI_Group), INTENT(IN) :: group1, group2
46    TYPE(MPI_Group), INTENT(OUT) :: newgroup
47    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48
49 MPI_Group_range_excl(group, n, ranges, newgroup, ierror)
50    TYPE(MPI_Group), INTENT(IN) :: group
51    INTEGER, INTENT(IN) :: n, ranges(3, n)
52    TYPE(MPI_Group), INTENT(OUT) :: newgroup
53    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```


MPI_Group_range_incl(group, n, ranges, newgroup, ierror)	1
TYPE(MPI_Group), INTENT(IN) :: group	2
INTEGER, INTENT(IN) :: n, ranges(3, n)	3
TYPE(MPI_Group), INTENT(OUT) :: newgroup	4
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	5
MPI_Group_rank(group, rank, ierror)	6
TYPE(MPI_Group), INTENT(IN) :: group	7
INTEGER, INTENT(OUT) :: rank	8
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	9
MPI_Group_size(group, size, ierror)	10
TYPE(MPI_Group), INTENT(IN) :: group	11
INTEGER, INTENT(OUT) :: size	12
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	13
MPI_Group_translate_ranks(group1, n, ranks1, group2, ranks2, ierror)	14
TYPE(MPI_Group), INTENT(IN) :: group1, group2	15
INTEGER, INTENT(IN) :: n, ranks1(n)	16
INTEGER, INTENT(OUT) :: ranks2(n)	17
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	18
MPI_Group_union(group1, group2, newgroup, ierror)	19
TYPE(MPI_Group), INTENT(IN) :: group1, group2	20
TYPE(MPI_Group), INTENT(OUT) :: newgroup	21
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	22
MPI_Intercomm_create(local_comm, local_leader, peer_comm, remote_leader, tag, newintercomm, ierror)	23
TYPE(MPI_Comm), INTENT(IN) :: local_comm, peer_comm	24
INTEGER, INTENT(IN) :: local_leader, remote_leader, tag	25
TYPE(MPI_Comm), INTENT(OUT) :: newintercomm	26
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	27
MPI_Intercomm_create_from_groups(local_group, local_leader, remote_group, remote_leader, stringtag, info, errhandler, newintercomm, ierror)	28
TYPE(MPI_Group), INTENT(IN) :: local_group, remote_group	29
INTEGER, INTENT(IN) :: local_leader, remote_leader	30
CHARACTER(LEN=*), INTENT(IN) :: stringtag	31
TYPE(MPI_Info), INTENT(IN) :: info	32
TYPE(MPI_Errhandler), INTENT(IN) :: errhandler	33
TYPE(MPI_Comm), INTENT(OUT) :: newintercomm	34
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	35
MPI_Intercomm_merge(intercomm, high, newintracomm, ierror)	36
TYPE(MPI_Comm), INTENT(IN) :: intercomm	37
LOGICAL, INTENT(IN) :: high	38
TYPE(MPI_Comm), INTENT(OUT) :: newintracomm	39
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	40
MPI_Type_create_keyval(type_copy_attr_fn, type_delete_attr_fn, type_keyval, extra_state, ierror)	41
TYPE(MPI_Type), INTENT(OUT) :: type_keyval	42
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	43
	44
	45
	46
	47
	48

```

1     PROCEDURE(MPI_Type_copy_attr_function) :: type_copy_attr_fn
2     PROCEDURE(MPI_Type_delete_attr_function) :: type_delete_attr_fn
3     INTEGER, INTENT(OUT) :: type_keyval
4     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
5     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
6
7     MPI_Type_delete_attr(datatype, type_keyval, ierror)
8     TYPE(MPI_Datatype), INTENT(IN) :: datatype
9     INTEGER, INTENT(IN) :: type_keyval
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12    MPI_TYPE_DUP_FN(oldtype, type_keyval, extra_state, attribute_val_in,
13    attribute_val_out, flag, ierror)
14    TYPE(MPI_Datatype) :: oldtype
15    INTEGER :: type_keyval, ierror
16    INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
17    attribute_val_out
18    LOGICAL :: flag
19
20    MPI_Type_free_keyval(type_keyval, ierror)
21    INTEGER, INTENT(INOUT) :: type_keyval
22    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24    MPI_Type_get_attr(datatype, type_keyval, attribute_val, flag, ierror)
25    TYPE(MPI_Datatype), INTENT(IN) :: datatype
26    INTEGER, INTENT(IN) :: type_keyval
27    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: attribute_val
28    LOGICAL, INTENT(OUT) :: flag
29    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
30
31    MPI_Type_get_name(datatype, type_name, resultlen, ierror)
32    TYPE(MPI_Datatype), INTENT(IN) :: datatype
33    CHARACTER(LEN=MPI_MAX_OBJECT_NAME), INTENT(OUT) :: type_name
34    INTEGER, INTENT(OUT) :: resultlen
35    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
37    MPI_TYPE_NULL_COPY_FN(oldtype, type_keyval, extra_state, attribute_val_in,
38    attribute_val_out, flag, ierror)
39    TYPE(MPI_Datatype) :: oldtype
40    INTEGER :: type_keyval, ierror
41    INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
42    attribute_val_out
43    LOGICAL :: flag
44
45    MPI_TYPE_NULL_DELETE_FN(datatype, type_keyval, attribute_val, extra_state,
46    ierror)
47    TYPE(MPI_Datatype) :: datatype
48    INTEGER :: type_keyval
49    INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
50    INTEGER, INTENT(OUT) :: ierror
51
52    MPI_Type_set_attr(datatype, type_keyval, attribute_val, ierror)

```

```

TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, INTENT(IN) :: type_keyval
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: attribute_val
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Type_set_name(datatype, type_name, ierror)
TYPE(MPI_Datatype), INTENT(IN) :: datatype
CHARACTER(LEN=*), INTENT(IN) :: type_name
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_create_keyval(win_copy_attr_fn, win_delete_attr_fn, win_keyval,
                     extra_state, ierror)
PROCEDURE(MPI_Win_copy_attr_function) :: win_copy_attr_fn
PROCEDURE(MPI_Win_delete_attr_function) :: win_delete_attr_fn
INTEGER, INTENT(OUT) :: win_keyval
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_delete_attr(win, win_keyval, ierror)
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, INTENT(IN) :: win_keyval
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_WIN_DUP_FN(oldwin, win_keyval, extra_state, attribute_val_in,
               attribute_val_out, flag, ierror)
TYPE(MPI_Win) :: oldwin
INTEGER :: win_keyval, ierror
INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
attribute_val_out
LOGICAL :: flag
MPI_Win_free_keyval(win_keyval, ierror)
INTEGER, INTENT(INOUT) :: win_keyval
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_get_attr(win, win_keyval, attribute_val, flag, ierror)
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, INTENT(IN) :: win_keyval
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: attribute_val
LOGICAL, INTENT(OUT) :: flag
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_get_name(win, win_name, resultlen, ierror)
TYPE(MPI_Win), INTENT(IN) :: win
CHARACTER(LEN=MPI_MAX_OBJECT_NAME), INTENT(OUT) :: win_name
INTEGER, INTENT(OUT) :: resultlen
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_WIN_NULL_COPY_FN(oldwin, win_keyval, extra_state, attribute_val_in,
                     attribute_val_out, flag, ierror)
TYPE(MPI_Win) :: oldwin
INTEGER :: win_keyval, ierror

```

```

1      INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state, attribute_val_in,
2          attribute_val_out
3      LOGICAL :: flag
4
5      MPI_WIN_NULL_DELETE_FN(win, win_keyval, attribute_val, extra_state, ierror)
6          TYPE(MPI_Win) :: win
7          INTEGER :: win_keyval, ierror
8          INTEGER(KIND=MPI_ADDRESS_KIND) :: attribute_val, extra_state
9
10     MPI_Win_set_attr(win, win_keyval, attribute_val, ierror)
11     TYPE(MPI_Win), INTENT(IN) :: win
12     INTEGER, INTENT(IN) :: win_keyval
13     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: attribute_val
14     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15
16     MPI_Win_set_name(win, win_name, ierror)
17     TYPE(MPI_Win), INTENT(IN) :: win
18     CHARACTER(LEN=*), INTENT(IN) :: win_name
19     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21     A.4.6 Virtual Topologies for MPI Processes Fortran 2008 Bindings
22
23     MPI_Cart_coords(comm, rank, maxdims, coords, ierror)
24     TYPE(MPI_Comm), INTENT(IN) :: comm
25     INTEGER, INTENT(IN) :: rank, maxdims
26     INTEGER, INTENT(OUT) :: coords(maxdims)
27     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
28
29     MPI_Cart_create(comm_old, ndims, dims, periods, reorder, comm_cart, ierror)
30     TYPE(MPI_Comm), INTENT(IN) :: comm_old
31     INTEGER, INTENT(IN) :: ndims, dims(ndims)
32     LOGICAL, INTENT(IN) :: periods(ndims), reorder
33     TYPE(MPI_Comm), INTENT(OUT) :: comm_cart
34     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
35
36     MPI_Cart_get(comm, maxdims, dims, periods, coords, ierror)
37     TYPE(MPI_Comm), INTENT(IN) :: comm
38     INTEGER, INTENT(IN) :: maxdims
39     INTEGER, INTENT(OUT) :: dims(maxdims), coords(maxdims)
40     LOGICAL, INTENT(OUT) :: periods(maxdims)
41     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43     MPI_Cart_map(comm, ndims, dims, periods, newrank, ierror)
44     TYPE(MPI_Comm), INTENT(IN) :: comm
45     INTEGER, INTENT(IN) :: ndims, dims(ndims)
46     LOGICAL, INTENT(IN) :: periods(ndims)
47     INTEGER, INTENT(OUT) :: newrank
48     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
49
50     MPI_Cart_rank(comm, coords, rank, ierror)
51     TYPE(MPI_Comm), INTENT(IN) :: comm
52     INTEGER, INTENT(IN) :: coords(*)

```

```

INTEGER, INTENT(OUT) :: rank
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Cart_shift(comm, direction, disp, rank_source, rank_dest, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(IN) :: direction, disp
INTEGER, INTENT(OUT) :: rank_source, rank_dest
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Cart_sub(comm, remain_dims, newcomm, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
LOGICAL, INTENT(IN) :: remain_dims(*)
TYPE(MPI_Comm), INTENT(OUT) :: newcomm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Cartdim_get(comm, ndims, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(OUT) :: ndims
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Dims_create(nnodes, ndims, dims, ierror)
INTEGER, INTENT(IN) :: nnodes, ndims
INTEGER, INTENT(INOUT) :: dims(ndims)
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Dist_graph_create(comm_old, n, sources, degrees, destinations, weights,
info, reorder, comm_dist_graph, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm_old
INTEGER, INTENT(IN) :: n, sources(n), degrees(n), destinations(*),
weights(*)
TYPE(MPI_Info), INTENT(IN) :: info
LOGICAL, INTENT(IN) :: reorder
TYPE(MPI_Comm), INTENT(OUT) :: comm_dist_graph
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Dist_graph_create_adjacent(comm_old, indegree, sources, sourceweights,
outdegree, destinations, destweights, info, reorder,
comm_dist_graph, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm_old
INTEGER, INTENT(IN) :: indegree, sources(indegree), sourceweights(*),
outdegree, destinations(outdegree), destweights(*)
TYPE(MPI_Info), INTENT(IN) :: info
LOGICAL, INTENT(IN) :: reorder
TYPE(MPI_Comm), INTENT(OUT) :: comm_dist_graph
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Dist_graph_neighbors(comm, maxindegree, sources, sourceweights,
maxoutdegree, destinations, destweights, ierror)
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, INTENT(IN) :: maxindegree, maxoutdegree
INTEGER, INTENT(OUT) :: sources(maxindegree), destinations(maxoutdegree)
INTEGER :: sourceweights(*), destweights(*)

```

```

1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2
3 MPI_Dist_graph_neighbors_count(comm, indegree, outdegree, weighted, ierror)
4     TYPE(MPI_Comm), INTENT(IN) :: comm
5     INTEGER, INTENT(OUT) :: indegree, outdegree
6     LOGICAL, INTENT(OUT) :: weighted
7     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
8
9 MPI_Graph_create(comm_old, nnodes, index, edges, reorder, comm_graph, ierror)
10    TYPE(MPI_Comm), INTENT(IN) :: comm_old
11    INTEGER, INTENT(IN) :: nnodes, index(nnodes), edges(*)
12    LOGICAL, INTENT(IN) :: reorder
13    TYPE(MPI_Comm), INTENT(OUT) :: comm_graph
14    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15
16 MPI_Graph_get(comm, maxindex, maxedges, index, edges, ierror)
17    TYPE(MPI_Comm), INTENT(IN) :: comm
18    INTEGER, INTENT(IN) :: maxindex, maxedges
19    INTEGER, INTENT(OUT) :: index(maxindex), edges(maxedges)
20    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Graph_map(comm, nnodes, index, edges, newrank, ierror)
23    TYPE(MPI_Comm), INTENT(IN) :: comm
24    INTEGER, INTENT(IN) :: nnodes, index(nnodes), edges(*)
25    INTEGER, INTENT(OUT) :: newrank
26    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28 MPI_Graph_neighbors(comm, rank, maxneighbors, neighbors, ierror)
29    TYPE(MPI_Comm), INTENT(IN) :: comm
30    INTEGER, INTENT(IN) :: rank, maxneighbors
31    INTEGER, INTENT(OUT) :: neighbors(maxneighbors)
32    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
33
34 MPI_Graph_neighbors_count(comm, rank, nneighbors, ierror)
35    TYPE(MPI_Comm), INTENT(IN) :: comm
36    INTEGER, INTENT(IN) :: rank
37    INTEGER, INTENT(OUT) :: nneighbors
38    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
39
40 MPI_Graphdims_get(comm, nnodes, nedges, ierror)
41    TYPE(MPI_Comm), INTENT(IN) :: comm
42    INTEGER, INTENT(OUT) :: nnodes, nedges
43    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45 MPI_Ineighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
46    recvtype, comm, request, ierror)
47    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
48    INTEGER, INTENT(IN) :: sendcount, recvcount
49    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
50    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
51    TYPE(MPI_Comm), INTENT(IN) :: comm
52    TYPE(MPI_Request), INTENT(OUT) :: request

```

```

INTEGER, OPTIONAL, INTENT(OUT) :: ierror 1
MPI_Ineighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount, 2
    recvtype, comm, request, ierror) !(_c) 3
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 4
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount 5
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 6
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 7
TYPE(MPI_Comm), INTENT(IN) :: comm 8
TYPE(MPI_Request), INTENT(OUT) :: request 9
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 10
MPI_Ineighbor_allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, 11
    displs, recvtype, comm, request, ierror) 12
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 13
INTEGER, INTENT(IN) :: sendcount 14
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 15
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 16
INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcounts(*), displs(*) 17
TYPE(MPI_Comm), INTENT(IN) :: comm 18
TYPE(MPI_Request), INTENT(OUT) :: request 19
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 20
MPI_Ineighbor_allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcounts, 21
    displs, recvtype, comm, request, ierror) !(_c) 22
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 23
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount 24
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 25
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 26
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcounts(*) 27
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: displs(*) 28
TYPE(MPI_Comm), INTENT(IN) :: comm 29
TYPE(MPI_Request), INTENT(OUT) :: request 30
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 31
MPI_Ineighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, 32
    recvtype, comm, request, ierror) 33
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 34
INTEGER, INTENT(IN) :: sendcount, recvcount 35
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 36
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 37
TYPE(MPI_Comm), INTENT(IN) :: comm 38
TYPE(MPI_Request), INTENT(OUT) :: request 39
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 40
MPI_Ineighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcount, 41
    recvtype, comm, request, ierror) !(_c) 42
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 43
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount 44
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype 45
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 46

```

```

1     TYPE(MPI_Comm), INTENT(IN) :: comm
2     TYPE(MPI_Request), INTENT(OUT) :: request
3     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
4
5     MPI_Ineighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
6         recvcounts, rdispls, recvtype, comm, request, ierror)
7     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
8     INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
9         recvcounts(*), rdispls(*)
10    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
11    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
12    TYPE(MPI_Comm), INTENT(IN) :: comm
13    TYPE(MPI_Request), INTENT(OUT) :: request
14    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
15
16    MPI_Ineighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
17        recvcounts, rdispls, recvtype, comm, request, ierror) !(_c)
18    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
19    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
20        recvcounts(*)
21    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
22        rdispls(*)
23    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
24    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
25    TYPE(MPI_Comm), INTENT(IN) :: comm
26    TYPE(MPI_Request), INTENT(OUT) :: request
27    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
28
29    MPI_Ineighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
30        recvcounts, rdispls, recvtypes, comm, request, ierror)
31    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
32    INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), recvcounts(*)
33    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
34        rdispls(*)
35    TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
36    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
37    TYPE(MPI_Comm), INTENT(IN) :: comm
38    TYPE(MPI_Request), INTENT(OUT) :: request
39    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41    MPI_Ineighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
42        recvcounts, rdispls, recvtypes, comm, request, ierror) !(_c)
43    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
44    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
45        recvcounts(*)
46    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
47        rdispls(*)
48    TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
49    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
50    TYPE(MPI_Comm), INTENT(IN) :: comm

```



```

TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_allgatherv(sendbuf, sendcount, sendtype, recvbuf, recvcnts,
    displs, recvtype, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcnts(*), displs(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

```

1 MPI_Neighbor_allgather(sendbuf, sendcount, sendtype, recvbuf, recvcoun
2     displs, recvtype, comm, ierror) !(_c)
3     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
4     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcoun
5     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
6     TYPE(*), DIMENSION(..) :: recvbuf
7     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
8     TYPE(MPI_Comm), INTENT(IN) :: comm
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10
11 MPI_Neighbor_allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcoun
12     displs, recvtype, comm, info, request, ierror)
13     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
14     INTEGER, INTENT(IN) :: sendcount, displs(*)
15     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
16     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
17     INTEGER, INTENT(IN), ASYNCHRONOUS :: recvcoun
18     TYPE(MPI_Comm), INTENT(IN) :: comm
19     TYPE(MPI_Info), INTENT(IN) :: info
20     TYPE(MPI_Request), INTENT(OUT) :: request
21     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22
23 MPI_Neighbor_allgather_init(sendbuf, sendcount, sendtype, recvbuf, recvcoun
24     displs, recvtype, comm, info, request, ierror) !(_c)
25     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
26     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount
27     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
28     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
29     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: recvcoun
30     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: displs(*)
31     TYPE(MPI_Comm), INTENT(IN) :: comm
32     TYPE(MPI_Info), INTENT(IN) :: info
33     TYPE(MPI_Request), INTENT(OUT) :: request
34     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
35
36 MPI_Neighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcoun
37     recvtype, comm, ierror)
38     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
39     INTEGER, INTENT(IN) :: sendcount, recvcoun
40     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
41     TYPE(*), DIMENSION(..) :: recvbuf
42     TYPE(MPI_Comm), INTENT(IN) :: comm
43     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45 MPI_Neighbor_alltoall(sendbuf, sendcount, sendtype, recvbuf, recvcoun
46     recvtype, comm, ierror) !(_c)
47     TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
48     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcoun
49     TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
50     TYPE(*), DIMENSION(..) :: recvbuf

```

```

TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_alltoall_init(sendbuf, sendcount, sendtype, recvbuf, recvcount,
    recvtype, comm, info, request, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcount, recvcount
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Request), INTENT(OUT) :: request
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
    recvcounts, rdispls, recvtype, comm, ierror)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER, INTENT(IN) :: sendcounts(*), sdispls(*), recvcounts(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_alltoallv(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
    recvcounts, rdispls, recvtype, comm, ierror) !(_c)
TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcounts(*)
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..) :: recvbuf
TYPE(MPI_Comm), INTENT(IN) :: comm
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Neighbor_alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
    recvcounts, rdispls, recvtype, comm, info, request, ierror)
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), sdispls(*),
    recvcounts(*), rdispls(*)
TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf

```

```

1     TYPE(MPI_Comm), INTENT(IN) :: comm
2     TYPE(MPI_Info), INTENT(IN) :: info
3     TYPE(MPI_Request), INTENT(OUT) :: request
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6 MPI_Neighbor_alltoallv_init(sendbuf, sendcounts, sdispls, sendtype, recvbuf,
7     recvcounts, rdispls, recvtype, comm, info, request, ierror) !(_c)
8     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
9     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*),
10    recvcounts(*)
11    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
12    rdispls(*)
13    TYPE(MPI_Datatype), INTENT(IN) :: sendtype, recvtype
14    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
15    TYPE(MPI_Comm), INTENT(IN) :: comm
16    TYPE(MPI_Info), INTENT(IN) :: info
17    TYPE(MPI_Request), INTENT(OUT) :: request
18    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
19
20 MPI_Neighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
21    recvcounts, rdispls, recvtypes, comm, ierror)
22    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
23    INTEGER, INTENT(IN) :: sendcounts(*), recvcounts(*)
24    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
25    TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
26    TYPE(*), DIMENSION(..) :: recvbuf
27    TYPE(MPI_Comm), INTENT(IN) :: comm
28    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
29
30 MPI_Neighbor_alltoallw(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
31    recvcounts, rdispls, recvtypes, comm, ierror) !(_c)
32    TYPE(*), DIMENSION(..), INTENT(IN) :: sendbuf
33    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: sendcounts(*), recvcounts(*)
34    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: sdispls(*), rdispls(*)
35    TYPE(MPI_Datatype), INTENT(IN) :: sendtypes(*), recvtypes(*)
36    TYPE(*), DIMENSION(..) :: recvbuf
37    TYPE(MPI_Comm), INTENT(IN) :: comm
38    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
39
40 MPI_Neighbor_alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf,
41    recvcounts, rdispls, recvtypes, comm, info, request, ierror)
42    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf
43    INTEGER, INTENT(IN), ASYNCHRONOUS :: sendcounts(*), recvcounts(*)
44    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*),
45    rdispls(*)
46    TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*)
47    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf
48    TYPE(MPI_Comm), INTENT(IN) :: comm
49    TYPE(MPI_Info), INTENT(IN) :: info
50    TYPE(MPI_Request), INTENT(OUT) :: request

```

```

INTEGER, OPTIONAL, INTENT(OUT) :: ierror 1
MPI_Neighbor_alltoallw_init(sendbuf, sendcounts, sdispls, sendtypes, recvbuf, 2
    recvcounts, rdispls, recvtypes, comm, info, request, ierror) 3
    !(_c) 4
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: sendbuf 5
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN), ASYNCHRONOUS :: sendcounts(*), 6
    recvcounts(*) 7
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN), ASYNCHRONOUS :: sdispls(*), 8
    rdispls(*) 9
TYPE(MPI_Datatype), INTENT(IN), ASYNCHRONOUS :: sendtypes(*), recvtypes(*) 10
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: recvbuf 11
TYPE(MPI_Comm), INTENT(IN) :: comm 12
TYPE(MPI_Info), INTENT(IN) :: info 13
TYPE(MPI_Request), INTENT(OUT) :: request 14
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 15
MPI_Topo_test(comm, status, ierror) 16
    TYPE(MPI_Comm), INTENT(IN) :: comm 17
    INTEGER, INTENT(OUT) :: status 18
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror 19
20
21
22
A.4.7 MPI Environmental Management Fortran 2008 Bindings 23
DOUBLE PRECISION MPI_Wtick() 24
DOUBLE PRECISION MPI_Wtime() 25
MPI_Add_error_class(errorclass, ierror) 26
    INTEGER, INTENT(OUT) :: errorclass 27
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror 28
MPI_Add_error_code(errorclass, errorcode, ierror) 29
    INTEGER, INTENT(IN) :: errorclass 30
    INTEGER, INTENT(OUT) :: errorcode 31
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror 32
MPI_Add_error_string(errorcode, string, ierror) 33
    INTEGER, INTENT(IN) :: errorcode 34
    CHARACTER(LEN=*), INTENT(IN) :: string 35
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror 36
MPI_Alloc_mem(size, info, baseptr, ierror) 37
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR 38
    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size 39
    TYPE(MPI_Info), INTENT(IN) :: info 40
    TYPE(C_PTR), INTENT(OUT) :: baseptr 41
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror 42
MPI_Comm_call_errhandler(comm, errorcode, ierror) 43
    TYPE(MPI_Comm), INTENT(IN) :: comm 44
    INTEGER, INTENT(IN) :: errorcode 45
46
47
48

```

```
1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2
3 MPI_Comm_create_errhandler(comm_errhandler_fn, errhandler, ierror)
4     PROCEDURE(MPI_Comm_errhandler_function) :: comm_errhandler_fn
5     TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
6     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
7
8 MPI_Comm_get_errhandler(comm, errhandler, ierror)
9     TYPE(MPI_Comm), INTENT(IN) :: comm
10    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_Comm_set_errhandler(comm, errhandler, ierror)
14    TYPE(MPI_Comm), INTENT(IN) :: comm
15    TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
16    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
17
18 MPI_Errhandler_free(errhandler, ierror)
19    TYPE(MPI_Errhandler), INTENT(INOUT) :: errhandler
20    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Error_class(errorcode, errorclass, ierror)
23    INTEGER, INTENT(IN) :: errorcode
24    INTEGER, INTENT(OUT) :: errorclass
25    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27 MPI_Error_string(errorcode, string, resultlen, ierror)
28    INTEGER, INTENT(IN) :: errorcode
29    CHARACTER(LEN=MPI_MAX_ERROR_STRING), INTENT(OUT) :: string
30    INTEGER, INTENT(OUT) :: resultlen
31    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_File_call_errhandler(fh, errorcode, ierror)
34    TYPE(MPI_File), INTENT(IN) :: fh
35    INTEGER, INTENT(IN) :: errorcode
36    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_File_create_errhandler(file_errhandler_fn, errhandler, ierror)
39    PROCEDURE(MPI_File_errhandler_function) :: file_errhandler_fn
40    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
41    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_File_get_errhandler(file, errhandler, ierror)
44    TYPE(MPI_File), INTENT(IN) :: file
45    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
46    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
47
48 MPI_File_set_errhandler(file, errhandler, ierror)
49    TYPE(MPI_File), INTENT(IN) :: file
50    TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
51    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
52
53 MPI_Free_mem(base, ierror)
54    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: base
```

INTEGER, OPTIONAL, INTENT(OUT) :: ierror	1
MPI_Get_hw_resource_info(hw_info, ierror)	2
TYPE(MPI_Info), INTENT(OUT) :: hw_info	3
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	4
MPI_Get_library_version(version, resultlen, ierror)	5
CHARACTER(LEN=MPI_MAX_LIBRARY_VERSION_STRING), INTENT(OUT) :: version	6
INTEGER, INTENT(OUT) :: resultlen	7
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	8
MPI_Get_processor_name(name, resultlen, ierror)	9
CHARACTER(LEN=MPI_MAX_PROCESSOR_NAME), INTENT(OUT) :: name	10
INTEGER, INTENT(OUT) :: resultlen	11
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	12
MPI_Get_version(version, subversion, ierror)	13
INTEGER, INTENT(OUT) :: version, subversion	14
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	15
MPI_Remove_error_class(errorclass, ierror)	16
INTEGER, INTENT(IN) :: errorclass	17
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	18
MPI_Remove_error_code(errorcode, ierror)	19
INTEGER, INTENT(IN) :: errorcode	20
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	21
MPI_Remove_error_string(errorcode, ierror)	22
INTEGER, INTENT(IN) :: errorcode	23
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	24
MPI_Session_call_errhandler(session, errorcode, ierror)	25
TYPE(MPI_Session), INTENT(IN) :: session	26
INTEGER, INTENT(IN) :: errorcode	27
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	28
MPI_Session_create_errhandler(session_errhandler_fn, errhandler, ierror)	29
PROCEDURE(MPI_Session_errhandler_function) :: session_errhandler_fn	30
TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler	31
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	32
MPI_Session_get_errhandler(session, errhandler, ierror)	33
TYPE(MPI_Session), INTENT(IN) :: session	34
TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler	35
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	36
MPI_Session_set_errhandler(session, errhandler, ierror)	37
TYPE(MPI_Session), INTENT(IN) :: session	38
TYPE(MPI_Errhandler), INTENT(IN) :: errhandler	39
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	40
MPI_Win_call_errhandler(win, errorcode, ierror)	41
TYPE(MPI_Win), INTENT(IN) :: win	42
INTEGER, INTENT(IN) :: errorcode	43
	44
	45
	46
	47
	48

```
1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2
3 MPI_Win_create_errhandler(win_errhandler_fn, errhandler, ierror)
4     PROCEDURE(MPI_Win_errhandler_function) :: win_errhandler_fn
5     TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
6     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
7
8 MPI_Win_get_errhandler(win, errhandler, ierror)
9     TYPE(MPI_Win), INTENT(IN) :: win
10    TYPE(MPI_Errhandler), INTENT(OUT) :: errhandler
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_Win_set_errhandler(win, errhandler, ierror)
14    TYPE(MPI_Win), INTENT(IN) :: win
15    TYPE(MPI_Errhandler), INTENT(IN) :: errhandler
16    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
17
18 A.4.8 The Info Object Fortran 2008 Bindings
19
20 MPI_Info_create(info, ierror)
21    TYPE(MPI_Info), INTENT(OUT) :: info
22    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24 MPI_Info_create_env(info, ierror)
25    TYPE(MPI_Info), INTENT(OUT) :: info
26    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28 MPI_Info_delete(info, key, ierror)
29    TYPE(MPI_Info), INTENT(IN) :: info
30    CHARACTER(LEN=*), INTENT(IN) :: key
31    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Info_dup(info, newinfo, ierror)
34    TYPE(MPI_Info), INTENT(IN) :: info
35    TYPE(MPI_Info), INTENT(OUT) :: newinfo
36    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Info_free(info, ierror)
39    TYPE(MPI_Info), INTENT(INOUT) :: info
40    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
41
42 MPI_Info_get_nkeys(info, nkeys, ierror)
43    TYPE(MPI_Info), INTENT(IN) :: info
44    INTEGER, INTENT(OUT) :: nkeys
45    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
46
47 MPI_Info_get_nthkey(info, n, key, ierror)
48    TYPE(MPI_Info), INTENT(IN) :: info
49    INTEGER, INTENT(IN) :: n
50    CHARACTER(LEN=*), INTENT(OUT) :: key
51    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
52
53 MPI_Info_get_string(info, key, buflen, value, flag, ierror)
```


TYPE(MPI_Info), INTENT(IN) :: info	1
CHARACTER(LEN=*), INTENT(IN) :: key	2
INTEGER, INTENT(INOUT) :: buflen	3
CHARACTER(LEN=*), INTENT(OUT) :: value	4
LOGICAL, INTENT(OUT) :: flag	5
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	6
MPI_Info_set(info, key, value, ierror)	7
TYPE(MPI_Info), INTENT(IN) :: info	8
CHARACTER(LEN=*), INTENT(IN) :: key, value	9
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	10
	11
	12
A.4.9 Process Creation and Management Fortran 2008 Bindings	13
MPI_Abort(comm, errorcode, ierror)	14
TYPE(MPI_Comm), INTENT(IN) :: comm	15
INTEGER, INTENT(IN) :: errorcode	16
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	17
	18
MPI_Close_port(port_name, ierror)	19
CHARACTER(LEN=*), INTENT(IN) :: port_name	20
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	21
	22
MPI_Comm_accept(port_name, info, root, comm, newcomm, ierror)	23
CHARACTER(LEN=*), INTENT(IN) :: port_name	24
TYPE(MPI_Info), INTENT(IN) :: info	25
INTEGER, INTENT(IN) :: root	26
TYPE(MPI_Comm), INTENT(IN) :: comm	27
TYPE(MPI_Comm), INTENT(OUT) :: newcomm	28
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	29
	30
MPI_Comm_connect(port_name, info, root, comm, newcomm, ierror)	31
CHARACTER(LEN=*), INTENT(IN) :: port_name	32
TYPE(MPI_Info), INTENT(IN) :: info	33
INTEGER, INTENT(IN) :: root	34
TYPE(MPI_Comm), INTENT(IN) :: comm	35
TYPE(MPI_Comm), INTENT(OUT) :: newcomm	36
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	37
	38
MPI_Comm_disconnect(comm, ierror)	39
TYPE(MPI_Comm), INTENT(INOUT) :: comm	40
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	41
	42
MPI_Comm_get_parent(parent, ierror)	43
TYPE(MPI_Comm), INTENT(OUT) :: parent	44
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	45
	46
MPI_Comm_join(fd, intercomm, ierror)	47
INTEGER, INTENT(IN) :: fd	48
TYPE(MPI_Comm), INTENT(OUT) :: intercomm	
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	

```

1 MPI_Comm_spawn(command, argv, maxprocs, info, root, comm, intercomm,
2     array_of_errcodes, ierror)
3     CHARACTER(LEN=*), INTENT(IN) :: command, argv(*)
4     INTEGER, INTENT(IN) :: maxprocs, root
5     TYPE(MPI_Info), INTENT(IN) :: info
6     TYPE(MPI_Comm), INTENT(IN) :: comm
7     TYPE(MPI_Comm), INTENT(OUT) :: intercomm
8     INTEGER :: array_of_errcodes(*)
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10
11 MPI_Comm_spawn_multiple(count, array_of_commands, array_of_argv,
12     array_of_maxprocs, array_of_info, root, comm, intercomm,
13     array_of_errcodes, ierror)
14     INTEGER, INTENT(IN) :: count, array_of_maxprocs(*), root
15     CHARACTER(LEN=*), INTENT(IN) :: array_of_commands(*),
16     array_of_argv(count, *)
17     TYPE(MPI_Info), INTENT(IN) :: array_of_info(*)
18     TYPE(MPI_Comm), INTENT(IN) :: comm
19     TYPE(MPI_Comm), INTENT(OUT) :: intercomm
20     INTEGER :: array_of_errcodes(*)
21     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22
23 MPI_Finalize(ierror)
24     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
25
26 MPI_Finalized(flag, ierror)
27     LOGICAL, INTENT(OUT) :: flag
28     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
29
30 MPI_Init(ierror)
31     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Init_thread(required, provided, ierror)
34     INTEGER, INTENT(IN) :: required
35     INTEGER, INTENT(OUT) :: provided
36     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Initialized(flag, ierror)
39     LOGICAL, INTENT(OUT) :: flag
40     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
41
42 MPI_Is_thread_main(flag, ierror)
43     LOGICAL, INTENT(OUT) :: flag
44     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
45
46 MPI_Lookup_name(service_name, info, port_name, ierror)
47     CHARACTER(LEN=*), INTENT(IN) :: service_name
48     TYPE(MPI_Info), INTENT(IN) :: info
49     CHARACTER(LEN=MPI_MAX_PORT_NAME), INTENT(OUT) :: port_name
50     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
51
52 MPI_Open_port(info, port_name, ierror)
53     TYPE(MPI_Info), INTENT(IN) :: info

```

```

CHARACTER(LEN=MPI_MAX_PORT_NAME), INTENT(OUT) :: port_name      1
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        2
                                                                3
MPI_Publish_name(service_name, info, port_name, ierror)        4
CHARACTER(LEN=*), INTENT(IN) :: service_name, port_name        5
TYPE(MPI_Info), INTENT(IN) :: info                              6
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                        7
                                                                8
MPI_Query_thread(provided, ierror)                               9
INTEGER, INTENT(OUT) :: provided                               10
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       11
                                                                12
MPI_Session_finalize(session, ierror)                           13
TYPE(MPI_Session), INTENT(INOUT) :: session                   14
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       15
                                                                16
MPI_Session_get_info(session, info_used, ierror)                17
TYPE(MPI_Session), INTENT(IN) :: session                       18
TYPE(MPI_Info), INTENT(OUT) :: info_used                       19
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       20
                                                                21
MPI_Session_get_nth_pset(session, info, n, pset_len, pset_name, ierror)  22
TYPE(MPI_Session), INTENT(IN) :: session                       23
TYPE(MPI_Info), INTENT(IN) :: info                             24
INTEGER, INTENT(IN) :: n                                       25
INTEGER, INTENT(INOUT) :: pset_len                             26
CHARACTER(LEN=*), INTENT(OUT) :: pset_name                     27
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       28
                                                                29
MPI_Session_get_num_psets(session, info, npset_names, ierror)  30
TYPE(MPI_Session), INTENT(IN) :: session                       31
TYPE(MPI_Info), INTENT(IN) :: info                             32
INTEGER, INTENT(OUT) :: npset_names                            33
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       34
                                                                35
MPI_Session_get_pset_info(session, pset_name, info, ierror)    36
TYPE(MPI_Session), INTENT(IN) :: session                       37
CHARACTER(LEN=*), INTENT(IN) :: pset_name                     38
TYPE(MPI_Info), INTENT(OUT) :: info                             39
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       40
                                                                41
MPI_Session_init(info, errhandler, session, ierror)             42
TYPE(MPI_Info), INTENT(IN) :: info                             43
TYPE(MPI_Errhandler), INTENT(IN) :: errhandler                44
TYPE(MPI_Session), INTENT(OUT) :: session                     45
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       46
                                                                47
MPI_Unpublish_name(service_name, info, port_name, ierror)      48
CHARACTER(LEN=*), INTENT(IN) :: service_name, port_name
TYPE(MPI_Info), INTENT(IN) :: info
INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

A.4.10 One-Sided Communications Fortran 2008 Bindings

```

1 MPI_Accumulate(origin_addr, origin_count, origin_datatype, target_rank,
2               target_disp, target_count, target_datatype, op, win, ierror)
3
4     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
5     INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
6     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
7     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
8     TYPE(MPI_Op), INTENT(IN) :: op
9     TYPE(MPI_Win), INTENT(IN) :: win
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12 MPI_Accumulate(origin_addr, origin_count, origin_datatype, target_rank,
13               target_disp, target_count, target_datatype, op, win, ierror)
14               !(_c)
15     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
16     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
17     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
18     INTEGER, INTENT(IN) :: target_rank
19     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
20     TYPE(MPI_Op), INTENT(IN) :: op
21     TYPE(MPI_Win), INTENT(IN) :: win
22     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
23
24 MPI_Compare_and_swap(origin_addr, compare_addr, result_addr, datatype,
25                      target_rank, target_disp, win, ierror)
26     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr,
27                      compare_addr
28     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr
29     TYPE(MPI_Datatype), INTENT(IN) :: datatype
30     INTEGER, INTENT(IN) :: target_rank
31     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
32     TYPE(MPI_Win), INTENT(IN) :: win
33     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
34
35 MPI_Fetch_and_op(origin_addr, result_addr, datatype, target_rank, target_disp,
36                  op, win, ierror)
37     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
38     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr
39     TYPE(MPI_Datatype), INTENT(IN) :: datatype
40     INTEGER, INTENT(IN) :: target_rank
41     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
42     TYPE(MPI_Op), INTENT(IN) :: op
43     TYPE(MPI_Win), INTENT(IN) :: win
44     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
45
46 MPI_Get(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
47          target_count, target_datatype, win, ierror)
48     TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr
49     INTEGER, INTENT(IN) :: origin_count, target_rank, target_count

```

```

TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype      1
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp              2
TYPE(MPI_Win), INTENT(IN) :: win                                       3
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                               4
MPI_Get(origin_addr, origin_count, origin_datatype, target_rank, target_disp, 5
        target_count, target_datatype, win, ierror) !(_c)              6
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr                    7
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count  8
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype     9
INTEGER, INTENT(IN) :: target_rank                                     10
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp            11
TYPE(MPI_Win), INTENT(IN) :: win                                       12
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                               13
MPI_Get_accumulate(origin_addr, origin_count, origin_datatype, result_addr, 14
        result_count, result_datatype, target_rank, target_disp,
        target_count, target_datatype, op, win, ierror)                15
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr      16
INTEGER, INTENT(IN) :: origin_count, result_count, target_rank,
        target_count                                                    17
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,
        target_datatype                                                 18
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr                  19
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp            20
TYPE(MPI_Op), INTENT(IN) :: op                                         21
TYPE(MPI_Win), INTENT(IN) :: win                                       22
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                               23
MPI_Get_accumulate(origin_addr, origin_count, origin_datatype, result_addr, 24
        result_count, result_datatype, target_rank, target_disp,
        target_count, target_datatype, op, win, ierror) !(_c)        25
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr      26
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, result_count,
        target_count                                                    27
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,
        target_datatype                                                 28
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr                  29
INTEGER, INTENT(IN) :: target_rank                                     30
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp            31
TYPE(MPI_Op), INTENT(IN) :: op                                         32
TYPE(MPI_Win), INTENT(IN) :: win                                       33
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                               34
MPI_Get_accumulate(origin_addr, origin_count, origin_datatype, result_addr, 35
        result_count, result_datatype, target_rank, target_disp,
        target_count, target_datatype, op, win, ierror) !(_c)        36
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr      37
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, result_count,
        target_count                                                    38
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,
        target_datatype                                                 39
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr                  40
INTEGER, INTENT(IN) :: target_rank                                     41
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp            42
TYPE(MPI_Op), INTENT(IN) :: op                                         43
TYPE(MPI_Win), INTENT(IN) :: win                                       44
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                               45
MPI_Put(origin_addr, origin_count, origin_datatype, target_rank, target_disp, 46
        target_count, target_datatype, win, ierror)                    47
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr      48
INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp

```

```

1     TYPE(MPI_Win), INTENT(IN) :: win
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
3
4 MPI_Put(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
5         target_count, target_datatype, win, ierror) !(_c)
6     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
7     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
8     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
9     INTEGER, INTENT(IN) :: target_rank
10    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
11    TYPE(MPI_Win), INTENT(IN) :: win
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13
14 MPI_Raccumulate(origin_addr, origin_count, origin_datatype, target_rank,
15               target_disp, target_count, target_datatype, op, win, request,
16               ierror)
17    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
18    INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
19    TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
20    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
21    TYPE(MPI_Op), INTENT(IN) :: op
22    TYPE(MPI_Win), INTENT(IN) :: win
23    TYPE(MPI_Request), INTENT(OUT) :: request
24    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
25
26 MPI_Raccumulate(origin_addr, origin_count, origin_datatype, target_rank,
27               target_disp, target_count, target_datatype, op, win, request,
28               ierror) !(_c)
29    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
30    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
31    TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
32    INTEGER, INTENT(IN) :: target_rank
33    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
34    TYPE(MPI_Op), INTENT(IN) :: op
35    TYPE(MPI_Win), INTENT(IN) :: win
36    TYPE(MPI_Request), INTENT(OUT) :: request
37    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
38
39 MPI_Rget(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
40         target_count, target_datatype, win, request, ierror)
41    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr
42    INTEGER, INTENT(IN) :: origin_count, target_rank, target_count
43    TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
44    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
45    TYPE(MPI_Win), INTENT(IN) :: win
46    TYPE(MPI_Request), INTENT(OUT) :: request
47    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
48
49 MPI_Rget(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
50         target_count, target_datatype, win, request, ierror) !(_c)
51    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: origin_addr

```

```

INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count      1
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype        2
INTEGER, INTENT(IN) :: target_rank                                       3
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp                4
TYPE(MPI_Win), INTENT(IN) :: win                                         5
TYPE(MPI_Request), INTENT(OUT) :: request                                 6
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  7
                                                                           8
MPI_Rget_accumulate(origin_addr, origin_count, origin_datatype, result_addr,  9
                    result_count, result_datatype, target_rank, target_disp, 10
                    target_count, target_datatype, op, win, request, ierror) 11
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr         12
INTEGER, INTENT(IN) :: origin_count, result_count, target_rank,         13
                    target_count                                         14
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,     15
                    target_datatype                                       16
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr                     17
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp                18
TYPE(MPI_Op), INTENT(IN) :: op                                           19
TYPE(MPI_Win), INTENT(IN) :: win                                         20
TYPE(MPI_Request), INTENT(OUT) :: request                                 21
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  22
                                                                           23
MPI_Rget_accumulate(origin_addr, origin_count, origin_datatype, result_addr,  24
                    result_count, result_datatype, target_rank, target_disp, 25
                    target_count, target_datatype, op, win, request, ierror) !(_c) 26
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr         27
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, result_count,  28
                    target_count                                         29
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, result_datatype,     30
                    target_datatype                                       31
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: result_addr                     32
INTEGER, INTENT(IN) :: target_rank                                       33
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp                34
TYPE(MPI_Op), INTENT(IN) :: op                                           35
TYPE(MPI_Win), INTENT(IN) :: win                                         36
TYPE(MPI_Request), INTENT(OUT) :: request                                 37
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  38
                                                                           39
MPI_Rput(origin_addr, origin_count, origin_datatype, target_rank, target_disp, 40
          target_count, target_datatype, win, request, ierror)           41
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr         42
INTEGER, INTENT(IN) :: origin_count, target_rank, target_count           43
TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype       44
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp                45
TYPE(MPI_Win), INTENT(IN) :: win                                         46
TYPE(MPI_Request), INTENT(OUT) :: request                                 47
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                                  48

```

```

1 MPI_Rput(origin_addr, origin_count, origin_datatype, target_rank, target_disp,
2         target_count, target_datatype, win, request, ierror) !(_c)
3     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: origin_addr
4     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: origin_count, target_count
5     TYPE(MPI_Datatype), INTENT(IN) :: origin_datatype, target_datatype
6     INTEGER, INTENT(IN) :: target_rank
7     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: target_disp
8     TYPE(MPI_Win), INTENT(IN) :: win
9     TYPE(MPI_Request), INTENT(OUT) :: request
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12 MPI_Win_allocate(size, disp_unit, info, comm, baseptr, win, ierror)
13     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
14     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
15     INTEGER, INTENT(IN) :: disp_unit
16     TYPE(MPI_Info), INTENT(IN) :: info
17     TYPE(MPI_Comm), INTENT(IN) :: comm
18     TYPE(C_PTR), INTENT(OUT) :: baseptr
19     TYPE(MPI_Win), INTENT(OUT) :: win
20     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Win_allocate(size, disp_unit, info, comm, baseptr, win, ierror) !(_c)
23     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
24     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size, disp_unit
25     TYPE(MPI_Info), INTENT(IN) :: info
26     TYPE(MPI_Comm), INTENT(IN) :: comm
27     TYPE(C_PTR), INTENT(OUT) :: baseptr
28     TYPE(MPI_Win), INTENT(OUT) :: win
29     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
30
31 MPI_Win_allocate_shared(size, disp_unit, info, comm, baseptr, win, ierror)
32     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
33     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
34     INTEGER, INTENT(IN) :: disp_unit
35     TYPE(MPI_Info), INTENT(IN) :: info
36     TYPE(MPI_Comm), INTENT(IN) :: comm
37     TYPE(C_PTR), INTENT(OUT) :: baseptr
38     TYPE(MPI_Win), INTENT(OUT) :: win
39     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41 MPI_Win_allocate_shared(size, disp_unit, info, comm, baseptr, win, ierror)
42     !(_c)
43     USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
44     INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size, disp_unit
45     TYPE(MPI_Info), INTENT(IN) :: info
46     TYPE(MPI_Comm), INTENT(IN) :: comm
47     TYPE(C_PTR), INTENT(OUT) :: baseptr
48     TYPE(MPI_Win), INTENT(OUT) :: win
49     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
50
51 MPI_Win_attach(win, base, size, ierror)

```



```

TYPE(MPI_Win), INTENT(IN) :: win
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_complete(win, ierror)
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_create(base, size, disp_unit, info, comm, win, ierror)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size
INTEGER, INTENT(IN) :: disp_unit
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Win), INTENT(OUT) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_create(base, size, disp_unit, info, comm, win, ierror) !(_c)
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: size, disp_unit
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Win), INTENT(OUT) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_create_dynamic(info, comm, win, ierror)
TYPE(MPI_Info), INTENT(IN) :: info
TYPE(MPI_Comm), INTENT(IN) :: comm
TYPE(MPI_Win), INTENT(OUT) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_detach(win, base, ierror)
TYPE(MPI_Win), INTENT(IN) :: win
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: base
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_fence(assert, win, ierror)
INTEGER, INTENT(IN) :: assert
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_flush(rank, win, ierror)
INTEGER, INTENT(IN) :: rank
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_flush_all(win, ierror)
TYPE(MPI_Win), INTENT(IN) :: win
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_Win_flush_local(rank, win, ierror)
INTEGER, INTENT(IN) :: rank

```

```
1     TYPE(MPI_Win), INTENT(IN) :: win
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
3
4 MPI_Win_flush_local_all(win, ierror)
5     TYPE(MPI_Win), INTENT(IN) :: win
6     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
7
8 MPI_Win_free(win, ierror)
9     TYPE(MPI_Win), INTENT(INOUT) :: win
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12 MPI_Win_get_group(win, group, ierror)
13    TYPE(MPI_Win), INTENT(IN) :: win
14    TYPE(MPI_Group), INTENT(OUT) :: group
15    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
16
17 MPI_Win_get_info(win, info_used, ierror)
18    TYPE(MPI_Win), INTENT(IN) :: win
19    TYPE(MPI_Info), INTENT(OUT) :: info_used
20    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Win_lock(lock_type, rank, assert, win, ierror)
23    INTEGER, INTENT(IN) :: lock_type, rank, assert
24    TYPE(MPI_Win), INTENT(IN) :: win
25    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27 MPI_Win_lock_all(assert, win, ierror)
28    INTEGER, INTENT(IN) :: assert
29    TYPE(MPI_Win), INTENT(IN) :: win
30    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_Win_post(group, assert, win, ierror)
33    TYPE(MPI_Group), INTENT(IN) :: group
34    INTEGER, INTENT(IN) :: assert
35    TYPE(MPI_Win), INTENT(IN) :: win
36    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Win_set_info(win, info, ierror)
39    TYPE(MPI_Win), INTENT(IN) :: win
40    TYPE(MPI_Info), INTENT(IN) :: info
41    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Win_shared_query(win, rank, size, disp_unit, baseptr, ierror)
44    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
45    TYPE(MPI_Win), INTENT(IN) :: win
46    INTEGER, INTENT(IN) :: rank
47    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: size
48    INTEGER, INTENT(OUT) :: disp_unit
49    TYPE(C_PTR), INTENT(OUT) :: baseptr
50    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
51
52 MPI_Win_shared_query(win, rank, size, disp_unit, baseptr, ierror) !(_c)
53    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
```

TYPE(MPI_Win), INTENT(IN) :: win	1
INTEGER, INTENT(IN) :: rank	2
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: size, disp_unit	3
TYPE(C_PTR), INTENT(OUT) :: baseptr	4
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	5
MPI_Win_start(group, assert, win, ierror)	6
TYPE(MPI_Group), INTENT(IN) :: group	7
INTEGER, INTENT(IN) :: assert	8
TYPE(MPI_Win), INTENT(IN) :: win	10
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	11
MPI_Win_sync(win, ierror)	12
TYPE(MPI_Win), INTENT(IN) :: win	13
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	14
MPI_Win_test(win, flag, ierror)	15
TYPE(MPI_Win), INTENT(IN) :: win	16
LOGICAL, INTENT(OUT) :: flag	17
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	18
MPI_Win_unlock(rank, win, ierror)	20
INTEGER, INTENT(IN) :: rank	21
TYPE(MPI_Win), INTENT(IN) :: win	22
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	23
MPI_Win_unlock_all(win, ierror)	24
TYPE(MPI_Win), INTENT(IN) :: win	25
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	26
MPI_Win_wait(win, ierror)	27
TYPE(MPI_Win), INTENT(IN) :: win	28
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	29
	30
	31
A.4.11 External Interfaces Fortran 2008 Bindings	32
MPI_Grequest_complete(request, ierror)	33
TYPE(MPI_Request), INTENT(IN) :: request	34
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	35
MPI_Grequest_start(query_fn, free_fn, cancel_fn, extra_state, request, ierror)	36
PROCEDURE(MPI_Grequest_query_function) :: query_fn	37
PROCEDURE(MPI_Grequest_free_function) :: free_fn	38
PROCEDURE(MPI_Grequest_cancel_function) :: cancel_fn	39
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state	40
TYPE(MPI_Request), INTENT(OUT) :: request	41
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	42
MPI_Status_set_cancelled(status, flag, ierror)	43
TYPE(MPI_Status), INTENT(INOUT) :: status	44
LOGICAL, INTENT(IN) :: flag	45
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	46
	47
	48

```

1 MPI_Status_set_elements(status, datatype, count, ierror)
2   TYPE(MPI_Status), INTENT(INOUT) :: status
3   TYPE(MPI_Datatype), INTENT(IN) :: datatype
4   INTEGER, INTENT(IN) :: count
5   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
6
7 MPI_Status_set_elements(status, datatype, count, ierror) !(_c)
8   TYPE(MPI_Status), INTENT(INOUT) :: status
9   TYPE(MPI_Datatype), INTENT(IN) :: datatype
10  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
11  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_Status_set_error(status, err, ierror)
14   TYPE(MPI_Status), INTENT(INOUT) :: status
15   INTEGER, INTENT(IN) :: err
16   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
17
18 MPI_Status_set_source(status, source, ierror)
19   TYPE(MPI_Status), INTENT(INOUT) :: status
20   INTEGER, INTENT(IN) :: source
21   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
22
23 MPI_Status_set_tag(status, tag, ierror)
24   TYPE(MPI_Status), INTENT(INOUT) :: status
25   INTEGER, INTENT(IN) :: tag
26   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28 A.4.12 I/O Fortran 2008 Bindings
29
30 MPI_CONVERSION_FN_NULL(userbuf, datatype, count, filebuf, position,
31   extra_state, ierror)
32   USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
33   TYPE(C_PTR), VALUE :: userbuf, filebuf
34   TYPE(MPI_Datatype) :: datatype
35   INTEGER :: count, ierror
36   INTEGER(KIND=MPI_OFFSET_KIND) :: position
37   INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
38
39 MPI_CONVERSION_FN_NULL_C(userbuf, datatype, count, filebuf, position,
40   extra_state, ierror) !(_c)
41   USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR
42   TYPE(C_PTR), VALUE :: userbuf, filebuf
43   TYPE(MPI_Datatype) :: datatype
44   INTEGER(KIND=MPI_COUNT_KIND) :: count
45   INTEGER(KIND=MPI_OFFSET_KIND) :: position
46   INTEGER(KIND=MPI_ADDRESS_KIND) :: extra_state
47   INTEGER :: ierror
48
49 MPI_File_close(fh, ierror)
50   TYPE(MPI_File), INTENT(INOUT) :: fh
51   INTEGER, OPTIONAL, INTENT(OUT) :: ierror

```

MPI_File_delete(filename, info, ierror)	1
CHARACTER(LEN=*), INTENT(IN) :: filename	2
TYPE(MPI_Info), INTENT(IN) :: info	3
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	4
MPI_File_get_amode(fh, amode, ierror)	5
TYPE(MPI_File), INTENT(IN) :: fh	6
INTEGER, INTENT(OUT) :: amode	7
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	8
MPI_File_get_atomicsity(fh, flag, ierror)	9
TYPE(MPI_File), INTENT(IN) :: fh	10
LOGICAL, INTENT(OUT) :: flag	11
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	12
MPI_File_get_byte_offset(fh, offset, disp, ierror)	13
TYPE(MPI_File), INTENT(IN) :: fh	14
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset	15
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: disp	16
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	17
MPI_File_get_group(fh, group, ierror)	18
TYPE(MPI_File), INTENT(IN) :: fh	19
TYPE(MPI_Group), INTENT(OUT) :: group	20
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	21
MPI_File_get_info(fh, info_used, ierror)	22
TYPE(MPI_File), INTENT(IN) :: fh	23
TYPE(MPI_Info), INTENT(OUT) :: info_used	24
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	25
MPI_File_get_position(fh, offset, ierror)	26
TYPE(MPI_File), INTENT(IN) :: fh	27
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: offset	28
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	29
MPI_File_get_position_shared(fh, offset, ierror)	30
TYPE(MPI_File), INTENT(IN) :: fh	31
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: offset	32
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	33
MPI_File_get_size(fh, size, ierror)	34
TYPE(MPI_File), INTENT(IN) :: fh	35
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: size	36
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	37
MPI_File_get_type_extent(fh, datatype, extent, ierror)	38
TYPE(MPI_File), INTENT(IN) :: fh	39
TYPE(MPI_Datatype), INTENT(IN) :: datatype	40
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(OUT) :: extent	41
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	42
MPI_File_get_type_extent(fh, datatype, extent, ierror) !(_c)	43
	44
	45
	46
	47
	48

```

1     TYPE(MPI_File), INTENT(IN) :: fh
2     TYPE(MPI_Datatype), INTENT(IN) :: datatype
3     INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: extent
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6 MPI_File_get_view(fh, disp, etype, filetype, datarep, ierror)
7     TYPE(MPI_File), INTENT(IN) :: fh
8     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(OUT) :: disp
9     TYPE(MPI_Datatype), INTENT(OUT) :: etype, filetype
10    CHARACTER(LEN=*), INTENT(OUT) :: datarep
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_File_iread(fh, buf, count, datatype, request, ierror)
14    TYPE(MPI_File), INTENT(IN) :: fh
15    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
16    INTEGER, INTENT(IN) :: count
17    TYPE(MPI_Datatype), INTENT(IN) :: datatype
18    TYPE(MPI_Request), INTENT(OUT) :: request
19    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21 MPI_File_iread(fh, buf, count, datatype, request, ierror) !(_c)
22    TYPE(MPI_File), INTENT(IN) :: fh
23    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
24    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
25    TYPE(MPI_Datatype), INTENT(IN) :: datatype
26    TYPE(MPI_Request), INTENT(OUT) :: request
27    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
28
29 MPI_File_iread_all(fh, buf, count, datatype, request, ierror)
30    TYPE(MPI_File), INTENT(IN) :: fh
31    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
32    INTEGER, INTENT(IN) :: count
33    TYPE(MPI_Datatype), INTENT(IN) :: datatype
34    TYPE(MPI_Request), INTENT(OUT) :: request
35    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
36
37 MPI_File_iread_all(fh, buf, count, datatype, request, ierror) !(_c)
38    TYPE(MPI_File), INTENT(IN) :: fh
39    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
40    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
41    TYPE(MPI_Datatype), INTENT(IN) :: datatype
42    TYPE(MPI_Request), INTENT(OUT) :: request
43    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45 MPI_File_iread_at(fh, offset, buf, count, datatype, request, ierror)
46    TYPE(MPI_File), INTENT(IN) :: fh
47    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
48    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
49    INTEGER, INTENT(IN) :: count
50    TYPE(MPI_Datatype), INTENT(IN) :: datatype
51    TYPE(MPI_Request), INTENT(OUT) :: request

```

```

INTEGER, OPTIONAL, INTENT(OUT) :: ierror 1
MPI_File_iread_at(fh, offset, buf, count, datatype, request, ierror) !(_c) 2
TYPE(MPI_File), INTENT(IN) :: fh 3
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset 4
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf 5
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count 6
TYPE(MPI_Datatype), INTENT(IN) :: datatype 7
TYPE(MPI_Request), INTENT(OUT) :: request 8
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 9
MPI_File_iread_at_all(fh, offset, buf, count, datatype, request, ierror) 10
TYPE(MPI_File), INTENT(IN) :: fh 11
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset 12
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf 13
INTEGER, INTENT(IN) :: count 14
TYPE(MPI_Datatype), INTENT(IN) :: datatype 15
TYPE(MPI_Request), INTENT(OUT) :: request 16
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 17
MPI_File_iread_at_all(fh, offset, buf, count, datatype, request, ierror) !(_c) 18
TYPE(MPI_File), INTENT(IN) :: fh 19
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset 20
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf 21
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count 22
TYPE(MPI_Datatype), INTENT(IN) :: datatype 23
TYPE(MPI_Request), INTENT(OUT) :: request 24
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 25
MPI_File_iread_shared(fh, buf, count, datatype, request, ierror) 26
TYPE(MPI_File), INTENT(IN) :: fh 27
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf 28
INTEGER, INTENT(IN) :: count 29
TYPE(MPI_Datatype), INTENT(IN) :: datatype 30
TYPE(MPI_Request), INTENT(OUT) :: request 31
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 32
MPI_File_iread_shared(fh, buf, count, datatype, request, ierror) !(_c) 33
TYPE(MPI_File), INTENT(IN) :: fh 34
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf 35
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count 36
TYPE(MPI_Datatype), INTENT(IN) :: datatype 37
TYPE(MPI_Request), INTENT(OUT) :: request 38
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 39
MPI_File_iwrite(fh, buf, count, datatype, request, ierror) 40
TYPE(MPI_File), INTENT(IN) :: fh 41
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf 42
INTEGER, INTENT(IN) :: count 43
TYPE(MPI_Datatype), INTENT(IN) :: datatype 44
TYPE(MPI_Request), INTENT(OUT) :: request 45
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 46

```

```
1     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
2
3 MPI_File_iwrite(fh, buf, count, datatype, request, ierror) !(_c)
4     TYPE(MPI_File), INTENT(IN) :: fh
5     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
6     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
7     TYPE(MPI_Datatype), INTENT(IN) :: datatype
8     TYPE(MPI_Request), INTENT(OUT) :: request
9     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
10
11 MPI_File_iwrite_all(fh, buf, count, datatype, request, ierror)
12     TYPE(MPI_File), INTENT(IN) :: fh
13     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
14     INTEGER, INTENT(IN) :: count
15     TYPE(MPI_Datatype), INTENT(IN) :: datatype
16     TYPE(MPI_Request), INTENT(OUT) :: request
17     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
18
19 MPI_File_iwrite_all(fh, buf, count, datatype, request, ierror) !(_c)
20     TYPE(MPI_File), INTENT(IN) :: fh
21     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
22     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
23     TYPE(MPI_Datatype), INTENT(IN) :: datatype
24     TYPE(MPI_Request), INTENT(OUT) :: request
25     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27 MPI_File_iwrite_at(fh, offset, buf, count, datatype, request, ierror)
28     TYPE(MPI_File), INTENT(IN) :: fh
29     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
30     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
31     INTEGER, INTENT(IN) :: count
32     TYPE(MPI_Datatype), INTENT(IN) :: datatype
33     TYPE(MPI_Request), INTENT(OUT) :: request
34     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
35
36 MPI_File_iwrite_at(fh, offset, buf, count, datatype, request, ierror) !(_c)
37     TYPE(MPI_File), INTENT(IN) :: fh
38     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
39     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
40     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
41     TYPE(MPI_Datatype), INTENT(IN) :: datatype
42     TYPE(MPI_Request), INTENT(OUT) :: request
43     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
44
45 MPI_File_iwrite_at_all(fh, offset, buf, count, datatype, request, ierror)
46     TYPE(MPI_File), INTENT(IN) :: fh
47     INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
48     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
49     INTEGER, INTENT(IN) :: count
50     TYPE(MPI_Datatype), INTENT(IN) :: datatype
51     TYPE(MPI_Request), INTENT(OUT) :: request
```



```

INTEGER, OPTIONAL, INTENT(OUT) :: ierror 1
MPI_File_iwrite_at_all(fh, offset, buf, count, datatype, request, ierror) !(_c) 2
TYPE(MPI_File), INTENT(IN) :: fh 3
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset 4
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf 5
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count 6
TYPE(MPI_Datatype), INTENT(IN) :: datatype 7
TYPE(MPI_Request), INTENT(OUT) :: request 8
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 9
MPI_File_iwrite_shared(fh, buf, count, datatype, request, ierror) 10
TYPE(MPI_File), INTENT(IN) :: fh 11
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf 12
INTEGER, INTENT(IN) :: count 13
TYPE(MPI_Datatype), INTENT(IN) :: datatype 14
TYPE(MPI_Request), INTENT(OUT) :: request 15
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 16
MPI_File_iwrite_shared(fh, buf, count, datatype, request, ierror) !(_c) 17
TYPE(MPI_File), INTENT(IN) :: fh 18
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf 19
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count 20
TYPE(MPI_Datatype), INTENT(IN) :: datatype 21
TYPE(MPI_Request), INTENT(OUT) :: request 22
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 23
MPI_File_open(comm, filename, amode, info, fh, ierror) 24
TYPE(MPI_Comm), INTENT(IN) :: comm 25
CHARACTER(LEN=*), INTENT(IN) :: filename 26
INTEGER, INTENT(IN) :: amode 27
TYPE(MPI_Info), INTENT(IN) :: info 28
TYPE(MPI_File), INTENT(OUT) :: fh 29
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 30
MPI_File_preallocate(fh, size, ierror) 31
TYPE(MPI_File), INTENT(IN) :: fh 32
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: size 33
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 34
MPI_File_read(fh, buf, count, datatype, status, ierror) 35
TYPE(MPI_File), INTENT(IN) :: fh 36
TYPE(*), DIMENSION(..) :: buf 37
INTEGER, INTENT(IN) :: count 38
TYPE(MPI_Datatype), INTENT(IN) :: datatype 39
TYPE(MPI_Status) :: status 40
INTEGER, OPTIONAL, INTENT(OUT) :: ierror 41
MPI_File_read(fh, buf, count, datatype, status, ierror) !(_c) 42
TYPE(MPI_File), INTENT(IN) :: fh 43
TYPE(*), DIMENSION(..) :: buf 44
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count 45

```

```
1     TYPE(MPI_Datatype), INTENT(IN) :: datatype
2     TYPE(MPI_Status) :: status
3     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
4
5 MPI_File_read_all(fh, buf, count, datatype, status, ierror)
6     TYPE(MPI_File), INTENT(IN) :: fh
7     TYPE(*), DIMENSION(..) :: buf
8     INTEGER, INTENT(IN) :: count
9     TYPE(MPI_Datatype), INTENT(IN) :: datatype
10    TYPE(MPI_Status) :: status
11    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
12
13 MPI_File_read_all(fh, buf, count, datatype, status, ierror) !(_c)
14    TYPE(MPI_File), INTENT(IN) :: fh
15    TYPE(*), DIMENSION(..) :: buf
16    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
17    TYPE(MPI_Datatype), INTENT(IN) :: datatype
18    TYPE(MPI_Status) :: status
19    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21 MPI_File_read_all_begin(fh, buf, count, datatype, ierror)
22    TYPE(MPI_File), INTENT(IN) :: fh
23    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
24    INTEGER, INTENT(IN) :: count
25    TYPE(MPI_Datatype), INTENT(IN) :: datatype
26    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28 MPI_File_read_all_begin(fh, buf, count, datatype, ierror) !(_c)
29    TYPE(MPI_File), INTENT(IN) :: fh
30    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
31    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
32    TYPE(MPI_Datatype), INTENT(IN) :: datatype
33    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
34
35 MPI_File_read_all_end(fh, buf, status, ierror)
36    TYPE(MPI_File), INTENT(IN) :: fh
37    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
38    TYPE(MPI_Status) :: status
39    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41 MPI_File_read_at(fh, offset, buf, count, datatype, status, ierror)
42    TYPE(MPI_File), INTENT(IN) :: fh
43    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
44    TYPE(*), DIMENSION(..) :: buf
45    INTEGER, INTENT(IN) :: count
46    TYPE(MPI_Datatype), INTENT(IN) :: datatype
47    TYPE(MPI_Status) :: status
48    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
49
50 MPI_File_read_at(fh, offset, buf, count, datatype, status, ierror) !(_c)
51    TYPE(MPI_File), INTENT(IN) :: fh
52    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
```

```

TYPE(*), DIMENSION(..) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_read_at_all(fh, offset, buf, count, datatype, status, ierror)
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..) :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_read_at_all(fh, offset, buf, count, datatype, status, ierror) !(_c)
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_read_at_all_begin(fh, offset, buf, count, datatype, ierror)
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_read_at_all_begin(fh, offset, buf, count, datatype, ierror) !(_c)
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_read_at_all_end(fh, buf, status, ierror)
TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_read_ordered(fh, buf, count, datatype, status, ierror)
TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..) :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype

```

```
1     TYPE(MPI_Status) :: status
2     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
3
4 MPI_File_read_ordered(fh, buf, count, datatype, status, ierror) !(_c)
5     TYPE(MPI_File), INTENT(IN) :: fh
6     TYPE(*), DIMENSION(..) :: buf
7     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
8     TYPE(MPI_Datatype), INTENT(IN) :: datatype
9     TYPE(MPI_Status) :: status
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12 MPI_File_read_ordered_begin(fh, buf, count, datatype, ierror)
13    TYPE(MPI_File), INTENT(IN) :: fh
14    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
15    INTEGER, INTENT(IN) :: count
16    TYPE(MPI_Datatype), INTENT(IN) :: datatype
17    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
18
19 MPI_File_read_ordered_begin(fh, buf, count, datatype, ierror) !(_c)
20    TYPE(MPI_File), INTENT(IN) :: fh
21    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
22    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
23    TYPE(MPI_Datatype), INTENT(IN) :: datatype
24    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
25
26 MPI_File_read_ordered_end(fh, buf, status, ierror)
27    TYPE(MPI_File), INTENT(IN) :: fh
28    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf
29    TYPE(MPI_Status) :: status
30    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
31
32 MPI_File_read_shared(fh, buf, count, datatype, status, ierror)
33    TYPE(MPI_File), INTENT(IN) :: fh
34    TYPE(*), DIMENSION(..) :: buf
35    INTEGER, INTENT(IN) :: count
36    TYPE(MPI_Datatype), INTENT(IN) :: datatype
37    TYPE(MPI_Status) :: status
38    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
39
40 MPI_File_read_shared(fh, buf, count, datatype, status, ierror) !(_c)
41    TYPE(MPI_File), INTENT(IN) :: fh
42    TYPE(*), DIMENSION(..) :: buf
43    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
44    TYPE(MPI_Datatype), INTENT(IN) :: datatype
45    TYPE(MPI_Status) :: status
46    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
47
48 MPI_File_seek(fh, offset, whence, ierror)
49    TYPE(MPI_File), INTENT(IN) :: fh
50    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
51    INTEGER, INTENT(IN) :: whence
52    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
```

MPI_File_seek_shared(fh, offset, whence, ierror)	1
TYPE(MPI_File), INTENT(IN) :: fh	2
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset	3
INTEGER, INTENT(IN) :: whence	4
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	5
	6
MPI_File_set_atomicity(fh, flag, ierror)	7
TYPE(MPI_File), INTENT(IN) :: fh	8
LOGICAL, INTENT(IN) :: flag	9
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	10
	11
MPI_File_set_info(fh, info, ierror)	11
TYPE(MPI_File), INTENT(IN) :: fh	12
TYPE(MPI_Info), INTENT(IN) :: info	13
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	14
	15
MPI_File_set_size(fh, size, ierror)	16
TYPE(MPI_File), INTENT(IN) :: fh	17
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: size	18
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	19
	20
MPI_File_set_view(fh, disp, etype, filetype, datarep, info, ierror)	20
TYPE(MPI_File), INTENT(IN) :: fh	21
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: disp	22
TYPE(MPI_Datatype), INTENT(IN) :: etype, filetype	23
CHARACTER(LEN=*), INTENT(IN) :: datarep	24
TYPE(MPI_Info), INTENT(IN) :: info	25
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	26
	27
MPI_File_sync(fh, ierror)	28
TYPE(MPI_File), INTENT(IN) :: fh	29
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	30
	31
MPI_File_write(fh, buf, count, datatype, status, ierror)	31
TYPE(MPI_File), INTENT(IN) :: fh	32
TYPE(*), DIMENSION(..), INTENT(IN) :: buf	33
INTEGER, INTENT(IN) :: count	34
TYPE(MPI_Datatype), INTENT(IN) :: datatype	35
TYPE(MPI_Status) :: status	36
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	37
	38
MPI_File_write(fh, buf, count, datatype, status, ierror) !(_c)	38
TYPE(MPI_File), INTENT(IN) :: fh	39
TYPE(*), DIMENSION(..), INTENT(IN) :: buf	40
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count	41
TYPE(MPI_Datatype), INTENT(IN) :: datatype	42
TYPE(MPI_Status) :: status	43
INTEGER, OPTIONAL, INTENT(OUT) :: ierror	44
	45
MPI_File_write_all(fh, buf, count, datatype, status, ierror)	46
TYPE(MPI_File), INTENT(IN) :: fh	47
TYPE(*), DIMENSION(..), INTENT(IN) :: buf	48

```
1     INTEGER, INTENT(IN) :: count
2     TYPE(MPI_Datatype), INTENT(IN) :: datatype
3     TYPE(MPI_Status) :: status
4     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
5
6 MPI_File_write_all(fh, buf, count, datatype, status, ierror) !(_c)
7     TYPE(MPI_File), INTENT(IN) :: fh
8     TYPE(*), DIMENSION(..), INTENT(IN) :: buf
9     INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
10    TYPE(MPI_Datatype), INTENT(IN) :: datatype
11    TYPE(MPI_Status) :: status
12    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
13
14 MPI_File_write_all_begin(fh, buf, count, datatype, ierror)
15    TYPE(MPI_File), INTENT(IN) :: fh
16    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
17    INTEGER, INTENT(IN) :: count
18    TYPE(MPI_Datatype), INTENT(IN) :: datatype
19    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
20
21 MPI_File_write_all_begin(fh, buf, count, datatype, ierror) !(_c)
22    TYPE(MPI_File), INTENT(IN) :: fh
23    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
24    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
25    TYPE(MPI_Datatype), INTENT(IN) :: datatype
26    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
27
28 MPI_File_write_all_end(fh, buf, status, ierror)
29    TYPE(MPI_File), INTENT(IN) :: fh
30    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
31    TYPE(MPI_Status) :: status
32    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
33
34 MPI_File_write_at(fh, offset, buf, count, datatype, status, ierror)
35    TYPE(MPI_File), INTENT(IN) :: fh
36    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
37    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
38    INTEGER, INTENT(IN) :: count
39    TYPE(MPI_Datatype), INTENT(IN) :: datatype
40    TYPE(MPI_Status) :: status
41    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_File_write_at(fh, offset, buf, count, datatype, status, ierror) !(_c)
44    TYPE(MPI_File), INTENT(IN) :: fh
45    INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
46    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
47    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
48    TYPE(MPI_Datatype), INTENT(IN) :: datatype
49    TYPE(MPI_Status) :: status
50    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
51
52 MPI_File_write_at_all(fh, offset, buf, count, datatype, status, ierror)
```

```

TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_at_all(fh, offset, buf, count, datatype, status, ierror) !(_c)
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_at_all_begin(fh, offset, buf, count, datatype, ierror)
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_at_all_begin(fh, offset, buf, count, datatype, ierror) !(_c)
TYPE(MPI_File), INTENT(IN) :: fh
INTEGER(KIND=MPI_OFFSET_KIND), INTENT(IN) :: offset
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_at_all_end(fh, buf, status, ierror)
TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_ordered(fh, buf, count, datatype, status, ierror)
TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER, INTENT(IN) :: count
TYPE(MPI_Datatype), INTENT(IN) :: datatype
TYPE(MPI_Status) :: status
INTEGER, OPTIONAL, INTENT(OUT) :: ierror
MPI_File_write_ordered(fh, buf, count, datatype, status, ierror) !(_c)
TYPE(MPI_File), INTENT(IN) :: fh
TYPE(*), DIMENSION(..), INTENT(IN) :: buf
INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count

```

```

1     TYPE(MPI_Datatype), INTENT(IN) :: datatype
2     TYPE(MPI_Status) :: status
3     INTEGER, OPTIONAL, INTENT(OUT) :: ierror
4
5     MPI_File_write_ordered_begin(fh, buf, count, datatype, ierror)
6     TYPE(MPI_File), INTENT(IN) :: fh
7     TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
8     INTEGER, INTENT(IN) :: count
9     TYPE(MPI_Datatype), INTENT(IN) :: datatype
10    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
11
12    MPI_File_write_ordered_begin(fh, buf, count, datatype, ierror) !(_c)
13    TYPE(MPI_File), INTENT(IN) :: fh
14    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
15    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
16    TYPE(MPI_Datatype), INTENT(IN) :: datatype
17    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
18
19    MPI_File_write_ordered_end(fh, buf, status, ierror)
20    TYPE(MPI_File), INTENT(IN) :: fh
21    TYPE(*), DIMENSION(..), INTENT(IN), ASYNCHRONOUS :: buf
22    TYPE(MPI_Status) :: status
23    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
24
25    MPI_File_write_shared(fh, buf, count, datatype, status, ierror)
26    TYPE(MPI_File), INTENT(IN) :: fh
27    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
28    INTEGER, INTENT(IN) :: count
29    TYPE(MPI_Datatype), INTENT(IN) :: datatype
30    TYPE(MPI_Status) :: status
31    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33    MPI_File_write_shared(fh, buf, count, datatype, status, ierror) !(_c)
34    TYPE(MPI_File), INTENT(IN) :: fh
35    TYPE(*), DIMENSION(..), INTENT(IN) :: buf
36    INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
37    TYPE(MPI_Datatype), INTENT(IN) :: datatype
38    TYPE(MPI_Status) :: status
39    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
40
41    MPI_Register_datarep(datarep, read_conversion_fn, write_conversion_fn,
42                        dtype_file_extent_fn, extra_state, ierror)
43    CHARACTER(LEN=*), INTENT(IN) :: datarep
44    PROCEDURE(MPI_Datarep_conversion_function) :: read_conversion_fn,
45    write_conversion_fn
46    PROCEDURE(MPI_Datarep_extent_function) :: dtype_file_extent_fn
47    INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state
48    INTEGER, OPTIONAL, INTENT(OUT) :: ierror
49
50    MPI_Register_datarep_c(datarep, read_conversion_fn, write_conversion_fn,
51                          dtype_file_extent_fn, extra_state, ierror) !(_c)
52    CHARACTER(LEN=*), INTENT(IN) :: datarep

```



```

PROCEDURE(MPI_Datarep_conversion_function_c) :: read_conversion_fn,      1
    write_conversion_fn                                              2
PROCEDURE(MPI_Datarep_extent_function) :: dtype_file_extent_fn      3
INTEGER(KIND=MPI_ADDRESS_KIND), INTENT(IN) :: extra_state          4
INTEGER, OPTIONAL, INTENT(OUT) :: ierror                            5
                                                                    6
                                                                    7

```

A.4.13 Language Bindings Fortran 2008 Bindings

```

MPI_F_sync_reg(buf)                                                  9
    TYPE(*), DIMENSION(..), ASYNCHRONOUS :: buf                    10
                                                                    11
MPI_Status_f082f(f08_status, f_status, ierror)                    12
    TYPE(MPI_Status), INTENT(IN) :: f08_status                    13
    INTEGER, INTENT(OUT) :: f_status(MPI_STATUS_SIZE)             14
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       15
                                                                    16
MPI_Status_f2f08(f_status, f08_status, ierror)                    17
    INTEGER, INTENT(IN) :: f_status(MPI_STATUS_SIZE)              18
    TYPE(MPI_Status), INTENT(OUT) :: f08_status                    19
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       20
                                                                    21
MPI_Type_create_f90_complex(p, r, newtype, ierror)                 22
    INTEGER, INTENT(IN) :: p, r                                    23
    TYPE(MPI_Datatype), INTENT(OUT) :: newtype                     24
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       25
                                                                    26
MPI_Type_create_f90_integer(r, newtype, ierror)                    27
    INTEGER, INTENT(IN) :: r                                        28
    TYPE(MPI_Datatype), INTENT(OUT) :: newtype                     29
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       30
                                                                    31
MPI_Type_create_f90_real(p, r, newtype, ierror)                    32
    INTEGER, INTENT(IN) :: p, r                                    33
    TYPE(MPI_Datatype), INTENT(OUT) :: newtype                     34
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       35
                                                                    36
MPI_Type_match_size(typeclass, size, datatype, ierror)            37
    INTEGER, INTENT(IN) :: typeclass, size                         38
    TYPE(MPI_Datatype), INTENT(OUT) :: datatype                   39
    INTEGER, OPTIONAL, INTENT(OUT) :: ierror                       40
                                                                    41

```

A.4.14 Tools / Profiling Interface Fortran 2008 Bindings

```

MPI_Pcontrol(level)                                                41
    INTEGER, INTENT(IN) :: level                                    42
                                                                    43
                                                                    44
                                                                    45
                                                                    46
                                                                    47
                                                                    48

```

A.4.15 Deprecated Fortran 2008 Bindings

```
1 MPI_Get_elements_x(status, datatype, count, ierror)
2   TYPE(MPI_Status), INTENT(IN) :: status
3   TYPE(MPI_Datatype), INTENT(IN) :: datatype
4   INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: count
5   INTEGER, OPTIONAL, INTENT(OUT) :: ierror
6
7 MPI_Info_get(info, key, valuelen, value, flag, ierror)
8   TYPE(MPI_Info), INTENT(IN) :: info
9   CHARACTER(LEN=*), INTENT(IN) :: key
10  INTEGER, INTENT(IN) :: valuelen
11  CHARACTER(LEN=valuelen), INTENT(OUT) :: value
12  LOGICAL, INTENT(OUT) :: flag
13  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
14
15 MPI_Info_get_valuelen(info, key, valuelen, flag, ierror)
16  TYPE(MPI_Info), INTENT(IN) :: info
17  CHARACTER(LEN=*), INTENT(IN) :: key
18  INTEGER, INTENT(OUT) :: valuelen
19  LOGICAL, INTENT(OUT) :: flag
20  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
21
22 MPI_Sizeof(x, size, ierror)
23  TYPE(*), DIMENSION(..) :: x
24  INTEGER, INTENT(OUT) :: size
25  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
26
27 MPI_Status_set_elements_x(status, datatype, count, ierror)
28  TYPE(MPI_Status), INTENT(INOUT) :: status
29  TYPE(MPI_Datatype), INTENT(IN) :: datatype
30  INTEGER(KIND=MPI_COUNT_KIND), INTENT(IN) :: count
31  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
32
33 MPI_Type_get_extent_x(datatype, lb, extent, ierror)
34  TYPE(MPI_Datatype), INTENT(IN) :: datatype
35  INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: lb, extent
36  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
37
38 MPI_Type_get_true_extent_x(datatype, true_lb, true_extent, ierror)
39  TYPE(MPI_Datatype), INTENT(IN) :: datatype
40  INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: true_lb, true_extent
41  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
42
43 MPI_Type_size_x(datatype, size, ierror)
44  TYPE(MPI_Datatype), INTENT(IN) :: datatype
45  INTEGER(KIND=MPI_COUNT_KIND), INTENT(OUT) :: size
46  INTEGER, OPTIONAL, INTENT(OUT) :: ierror
47
48
```

A.5 Fortran Bindings with mpif.h or the mpi Module

A.5.1 Point-to-Point Communication Fortran Bindings

MPI_BSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)	1
<type> BUF(*)	2
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR	3
MPI_BSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	4
<type> BUF(*)	5
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR	6
MPI_BUFFER_ATTACH(BUFFER, SIZE, IERROR)	7
<type> BUFFER(*)	8
INTEGER SIZE, IERROR	9
MPI_BUFFER_DETACH(BUFFER_ADDR, SIZE, IERROR)	10
<type> BUFFER_ADDR(*)	11
INTEGER SIZE, IERROR	12
MPI_BUFFER_FLUSH(IERROR)	13
INTEGER IERROR	14
MPI_BUFFER_IFLUSH(REQUEST, IERROR)	15
INTEGER REQUEST, IERROR	16
MPI_CANCEL(REQUEST, IERROR)	17
INTEGER REQUEST, IERROR	18
MPI_COMM_ATTACH_BUFFER(COMM, BUFFER, SIZE, IERROR)	19
INTEGER COMM, SIZE, IERROR	20
<type> BUFFER(*)	21
MPI_COMM_DETACH_BUFFER(COMM, BUFFER_ADDR, SIZE, IERROR)	22
INTEGER COMM, SIZE, IERROR	23
<type> BUFFER_ADDR(*)	24
MPI_COMM_FLUSH_BUFFER(COMM, IERROR)	25
INTEGER COMM, IERROR	26
MPI_COMM_IFLUSH_BUFFER(COMM, REQUEST, IERROR)	27
INTEGER COMM, REQUEST, IERROR	28
MPI_GET_COUNT(STATUS, DATATYPE, COUNT, IERROR)	29
INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR	30
MPI_IBSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	31
<type> BUF(*)	32
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR	33
MPI_IMPROBE(SOURCE, TAG, COMM, FLAG, MESSAGE, STATUS, IERROR)	34
INTEGER SOURCE, TAG, COMM, MESSAGE, STATUS(MPI_STATUS_SIZE), IERROR	35
LOGICAL FLAG	36
MPI_IMRECV(BUF, COUNT, DATATYPE, MESSAGE, REQUEST, IERROR)	37
<type> BUF(*)	38

```
1     INTEGER COUNT, DATATYPE, MESSAGE, REQUEST, IERROR
2
3 MPI_IProbe(SOURCE, TAG, COMM, FLAG, STATUS, IERROR)
4     INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
5     LOGICAL FLAG
6
7 MPI_Irecv(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR)
8     <type> BUF(*)
9     INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR
10
11 MPI_Irsend(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
12     <type> BUF(*)
13     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
14
15 MPI_Isend(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
16     <type> BUF(*)
17     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
18
19 MPI_Isendrecv(Sendbuf, Sendcount, Sendtype, Dest, Sendtag, Recvbuf, Recvcount,
20               Recvtype, Source, Recvtag, COMM, REQUEST, IERROR)
21     <type> Sendbuf(*), Recvbuf(*)
22     INTEGER Sendcount, Sendtype, Dest, Sendtag, Recvcount, Recvtype, Source,
23               Recvtag, COMM, REQUEST, IERROR
24
25 MPI_Isendrecv_replace(BUF, COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG,
26                       COMM, REQUEST, IERROR)
27     <type> BUF(*)
28     INTEGER COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG, COMM, REQUEST,
29               IERROR
30
31 MPI_Issend(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
32     <type> BUF(*)
33     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
34
35 MPI_Mprobe(SOURCE, TAG, COMM, MESSAGE, STATUS, IERROR)
36     INTEGER SOURCE, TAG, COMM, MESSAGE, STATUS(MPI_STATUS_SIZE), IERROR
37
38 MPI_Mrecv(BUF, COUNT, DATATYPE, MESSAGE, STATUS, IERROR)
39     <type> BUF(*)
40     INTEGER COUNT, DATATYPE, MESSAGE, STATUS(MPI_STATUS_SIZE), IERROR
41
42 MPI_Probe(SOURCE, TAG, COMM, STATUS, IERROR)
43     INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
44
45 MPI_Recv(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, STATUS, IERROR)
46     <type> BUF(*)
47     INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
48
49 MPI_Recv_init(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR)
50     <type> BUF(*)
51     INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR
52
53 MPI_Request_free(REQUEST, IERROR)
54     INTEGER REQUEST, IERROR
55
56 MPI_Request_get_status(REQUEST, FLAG, STATUS, IERROR)
```

INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR	1
LOGICAL FLAG	2
MPI_REQUEST_GET_STATUS_ALL(COUNT, ARRAY_OF_REQUESTS, FLAG, ARRAY_OF_STATUSES,	3
IERROR)	4
INTEGER COUNT, ARRAY_OF_REQUESTS(*), ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *),	5
IERROR	6
LOGICAL FLAG	7
MPI_REQUEST_GET_STATUS_ANY(COUNT, ARRAY_OF_REQUESTS, INDEX, FLAG, STATUS,	8
IERROR)	9
INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE), IERROR	10
LOGICAL FLAG	11
MPI_REQUEST_GET_STATUS_SOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT,	12
ARRAY_OF_INDICES, ARRAY_OF_STATUSES, IERROR)	13
INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),	14
ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *), IERROR	15
MPI_RSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)	16
<type> BUF(*)	17
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR	18
MPI_RSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	19
<type> BUF(*)	20
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR	21
MPI_SEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)	22
<type> BUF(*)	23
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR	24
MPI_SEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	25
<type> BUF(*)	26
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR	27
MPI_SENDRECV(SENDBUF, SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVBUF, RECVCOUNT,	28
RECVTYPE, SOURCE, RECVTAG, COMM, STATUS, IERROR)	29
<type> SENDBUF(*), RECVBUF(*)	30
INTEGER SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVCOUNT, RECVTYPE, SOURCE,	31
RECVTAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR	32
MPI_SENDRECV_REPLACE(BUF, COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG,	33
COMM, STATUS, IERROR)	34
<type> BUF(*)	35
INTEGER COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG, COMM,	36
STATUS(MPI_STATUS_SIZE), IERROR	37
MPI_SESSION_ATTACH_BUFFER(SESSION, BUFFER, SIZE, IERROR)	38
INTEGER SESSION, SIZE, IERROR	39
<type> BUFFER(*)	40
MPI_SESSION_DETACH_BUFFER(SESSION, BUFFER_ADDR, SIZE, IERROR)	41
INTEGER SESSION, SIZE, IERROR	42
<type> BUFFER_ADDR(*)	43
	44
	45
	46
	47
	48

```
1 MPI_SESSION_FLUSH_BUFFER(SESSION, IERROR)
2     INTEGER SESSION, IERROR
3
4 MPI_SESSION_IFLUSH_BUFFER(SESSION, REQUEST, IERROR)
5     INTEGER SESSION, REQUEST, IERROR
6
7 MPI_SSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
8     <type> BUF(*)
9     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR
10
11 MPI_SSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
12     <type> BUF(*)
13     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
14
15 MPI_START(REQUEST, IERROR)
16     INTEGER REQUEST, IERROR
17
18 MPI_STARTALL(COUNT, ARRAY_OF_REQUESTS, IERROR)
19     INTEGER COUNT, ARRAY_OF_REQUESTS(*), IERROR
20
21 MPI_STATUS_GET_ERROR(STATUS, ERR, IERROR)
22     INTEGER STATUS(MPI_STATUS_SIZE), ERR, IERROR
23
24 MPI_STATUS_GET_SOURCE(STATUS, SOURCE, IERROR)
25     INTEGER STATUS(MPI_STATUS_SIZE), SOURCE, IERROR
26
27 MPI_STATUS_GET_TAG(STATUS, TAG, IERROR)
28     INTEGER STATUS(MPI_STATUS_SIZE), TAG, IERROR
29
30 MPI_TEST(REQUEST, FLAG, STATUS, IERROR)
31     INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR
32     LOGICAL FLAG
33
34 MPI_TEST_CANCELLED(STATUS, FLAG, IERROR)
35     INTEGER STATUS(MPI_STATUS_SIZE), IERROR
36     LOGICAL FLAG
37
38 MPI_TESTALL(COUNT, ARRAY_OF_REQUESTS, FLAG, ARRAY_OF_STATUSES, IERROR)
39     INTEGER COUNT, ARRAY_OF_REQUESTS(*), ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *),
40     IERROR
41     LOGICAL FLAG
42
43 MPI_TESTANY(COUNT, ARRAY_OF_REQUESTS, INDEX, FLAG, STATUS, IERROR)
44     INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE), IERROR
45     LOGICAL FLAG
46
47 MPI_TESTSOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT, ARRAY_OF_INDICES,
48     ARRAY_OF_STATUSES, IERROR)
49     INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),
50     ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *), IERROR
51
52 MPI_WAIT(REQUEST, STATUS, IERROR)
53     INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR
54
55 MPI_WAITALL(COUNT, ARRAY_OF_REQUESTS, ARRAY_OF_STATUSES, IERROR)
56     INTEGER COUNT, ARRAY_OF_REQUESTS(*), ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *),
57     IERROR
```

```

MPI_WAITANY(COUNT, ARRAY_OF_REQUESTS, INDEX, STATUS, IERROR)           1
    INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE), IERROR  2
MPI_WAITSOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT, ARRAY_OF_INDICES,  3
    ARRAY_OF_STATUSES, IERROR)                                         4
    INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),  5
    ARRAY_OF_STATUSES(MPI_STATUS_SIZE, *), IERROR                       6

```

A.5.2 Partitioned Communication Fortran Bindings

```

MPI_PARRIVED(REQUEST, PARTITION, FLAG, IERROR)                         9
    INTEGER REQUEST, PARTITION, IERROR                                  10
    LOGICAL FLAG                                                       11
MPI_PREADY(PARTITION, REQUEST, IERROR)                                 12
    INTEGER PARTITION, REQUEST, IERROR                                 13
MPI_PREADY_LIST(LENGTH, ARRAY_OF_PARTITIONS, REQUEST, IERROR)         14
    INTEGER LENGTH, ARRAY_OF_PARTITIONS(*), REQUEST, IERROR          15
MPI_PREADY_RANGE(PARTITION_LOW, PARTITION_HIGH, REQUEST, IERROR)      16
    INTEGER PARTITION_LOW, PARTITION_HIGH, REQUEST, IERROR           17
MPI_PRECV_INIT(BUF, PARTITIONS, COUNT, DATATYPE, SOURCE, TAG, COMM, INFO,  18
    REQUEST, IERROR)                                                   19
    <type> BUF(*)                                                       20
    INTEGER PARTITIONS, DATATYPE, SOURCE, TAG, COMM, INFO, REQUEST, IERROR  21
    INTEGER(KIND=MPI_COUNT_KIND) COUNT                                22
MPI_PSEND_INIT(BUF, PARTITIONS, COUNT, DATATYPE, DEST, TAG, COMM, INFO,  23
    REQUEST, IERROR)                                                   24
    <type> BUF(*)                                                       25
    INTEGER PARTITIONS, DATATYPE, DEST, TAG, COMM, INFO, REQUEST, IERROR  26
    INTEGER(KIND=MPI_COUNT_KIND) COUNT                                27

```

A.5.3 Datatypes Fortran Bindings

```

INTEGER(KIND=MPI_ADDRESS_KIND) MPI_AINT_ADD(BASE, DISP)                28
    INTEGER(KIND=MPI_ADDRESS_KIND) BASE, DISP                          29
INTEGER(KIND=MPI_ADDRESS_KIND) MPI_AINT_DIFF(ADDR1, ADDR2)            30
    INTEGER(KIND=MPI_ADDRESS_KIND) ADDR1, ADDR2                         31
MPI_GET_ADDRESS(LOCATION, ADDRESS, IERROR)                              32
    <type> LOCATION(*)                                                  33
    INTEGER(KIND=MPI_ADDRESS_KIND) ADDRESS                              34
    INTEGER IERROR                                                       35
MPI_GET_ELEMENTS(STATUS, DATATYPE, COUNT, IERROR)                     36
    INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR           37
MPI_PACK(INBUF, INCOUNT, DATATYPE, OUTBUF, OUTSIZE, POSITION, COMM, IERROR)  38
    <type> INBUF(*), OUTBUF(*)                                          39

```

```

1     INTEGER INCOUNT, DATATYPE, OUTSIZE, POSITION, COMM, IERROR
2
3     MPI_PACK_EXTERNAL(DATAREP, INBUF, INCOUNT, DATATYPE, OUTBUF, OUTSIZE, POSITION,
4         IERROR)
5     CHARACTER*(*) DATAREP
6     <type> INBUF(*), OUTBUF(*)
7     INTEGER INCOUNT, DATATYPE, IERROR
8     INTEGER(KIND=MPI_ADDRESS_KIND) OUTSIZE, POSITION
9
10    MPI_PACK_EXTERNAL_SIZE(DATAREP, INCOUNT, DATATYPE, SIZE, IERROR)
11    CHARACTER*(*) DATAREP
12    INTEGER INCOUNT, DATATYPE, IERROR
13    INTEGER(KIND=MPI_ADDRESS_KIND) SIZE
14
15    MPI_PACK_SIZE(INCOUNT, DATATYPE, COMM, SIZE, IERROR)
16    INTEGER INCOUNT, DATATYPE, COMM, SIZE, IERROR
17
18    MPI_TYPE_COMMIT(DATATYPE, IERROR)
19    INTEGER DATATYPE, IERROR
20
21    MPI_TYPE_CONTIGUOUS(COUNT, OLDTYPE, NEWTYPE, IERROR)
22    INTEGER COUNT, OLDTYPE, NEWTYPE, IERROR
23
24    MPI_TYPE_CREATE_DARRAY(SIZE, RANK, NDIMS, ARRAY_OF_GSIZES, ARRAY_OF_DISTRIBS,
25        ARRAY_OF_DARGS, ARRAY_OF_PSIZEs, ORDER, OLDTYPE, NEWTYPE, IERROR)
26    INTEGER SIZE, RANK, NDIMS, ARRAY_OF_GSIZES(*), ARRAY_OF_DISTRIBS(*),
27        ARRAY_OF_DARGS(*), ARRAY_OF_PSIZEs(*), ORDER, OLDTYPE, NEWTYPE,
28        IERROR
29
30    MPI_TYPE_CREATE_HINDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
31        OLDTYPE, NEWTYPE, IERROR)
32    INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), OLDTYPE, NEWTYPE, IERROR
33    INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_DISPLACEMENTS(*)
34
35    MPI_TYPE_CREATE_HINDEXED_BLOCK(COUNT, BLOCKLENGTH, ARRAY_OF_DISPLACEMENTS,
36        OLDTYPE, NEWTYPE, IERROR)
37    INTEGER COUNT, BLOCKLENGTH, OLDTYPE, NEWTYPE, IERROR
38    INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_DISPLACEMENTS(*)
39
40    MPI_TYPE_CREATE_HVECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR)
41    INTEGER COUNT, BLOCKLENGTH, OLDTYPE, NEWTYPE, IERROR
42    INTEGER(KIND=MPI_ADDRESS_KIND) STRIDE
43
44    MPI_TYPE_CREATE_INDEXED_BLOCK(COUNT, BLOCKLENGTH, ARRAY_OF_DISPLACEMENTS,
45        OLDTYPE, NEWTYPE, IERROR)
46    INTEGER COUNT, BLOCKLENGTH, ARRAY_OF_DISPLACEMENTS(*), OLDTYPE, NEWTYPE,
47        IERROR
48
49    MPI_TYPE_CREATE_RESIZED(OLDTYPE, LB, EXTENT, NEWTYPE, IERROR)
50    INTEGER OLDTYPE, NEWTYPE, IERROR
51    INTEGER(KIND=MPI_ADDRESS_KIND) LB, EXTENT
52
53    MPI_TYPE_CREATE_STRUCT(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
54        ARRAY_OF_TYPES, NEWTYPE, IERROR)
55    INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_TYPES(*), NEWTYPE, IERROR

```



```

INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_DISPLACEMENTS(*)           1
MPI_TYPE_CREATE_SUBARRAY(NDIMS, ARRAY_OF_SIZES, ARRAY_OF_SUBSIZES,  2
    ARRAY_OF_STARTS, ORDER, OLDTYPE, NEWTYPE, IERROR)             3
INTEGER NDIMS, ARRAY_OF_SIZES(*), ARRAY_OF_SUBSIZES(*), ARRAY_OF_STARTS(*),  4
    ORDER, OLDTYPE, NEWTYPE, IERROR                               5
MPI_TYPE_DUP(OLDTYPE, NEWTYPE, IERROR)                             6
INTEGER OLDTYPE, NEWTYPE, IERROR                                  7
MPI_TYPE_FREE(DATATYPE, IERROR)                                    8
INTEGER DATATYPE, IERROR                                         9
MPI_TYPE_GET_CONTENTS(DATATYPE, MAX_INTEGERS, MAX_ADDRESSES, MAX_DATATYPES, 10
    ARRAY_OF_INTEGERS, ARRAY_OF_ADDRESSES, ARRAY_OF_DATATYPES, 11
    IERROR)                                                       12
INTEGER DATATYPE, MAX_INTEGERS, MAX_ADDRESSES, MAX_DATATYPES, 13
    ARRAY_OF_INTEGERS(*), ARRAY_OF_DATATYPES(*), IERROR         14
INTEGER(KIND=MPI_ADDRESS_KIND) ARRAY_OF_ADDRESSES(*)             15
MPI_TYPE_GET_ENVELOPE(DATATYPE, NUM_INTEGERS, NUM_ADDRESSES, NUM_DATATYPES, 16
    COMBINER, IERROR)                                             17
INTEGER DATATYPE, NUM_INTEGERS, NUM_ADDRESSES, NUM_DATATYPES, COMBINER, 18
    IERROR                                                         19
MPI_TYPE_GET_EXTENT(DATATYPE, LB, EXTENT, IERROR)                 20
INTEGER DATATYPE, IERROR                                         21
INTEGER(KIND=MPI_ADDRESS_KIND) LB, EXTENT                         22
MPI_TYPE_GET_TRUE_EXTENT(DATATYPE, TRUE_LB, TRUE_EXTENT, IERROR) 23
INTEGER DATATYPE, IERROR                                         24
INTEGER(KIND=MPI_ADDRESS_KIND) TRUE_LB, TRUE_EXTENT              25
MPI_TYPE_INDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS, OLDTYPE, 26
    NEWTYPE, IERROR)                                             27
INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*), 28
    OLDTYPE, NEWTYPE, IERROR                                     29
MPI_TYPE_SIZE(DATATYPE, SIZE, IERROR)                             30
INTEGER DATATYPE, SIZE, IERROR                                  31
MPI_TYPE_VECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR) 32
INTEGER COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR 33
MPI_UNPACK(INBUF, INSIZE, POSITION, OUTBUF, OUTCOUNT, DATATYPE, COMM, IERROR) 34
<type> INBUF(*), OUTBUF(*)                                       35
INTEGER INSIZE, POSITION, OUTCOUNT, DATATYPE, COMM, IERROR     36
MPI_UNPACK_EXTERNAL(DATAREP, INBUF, INSIZE, POSITION, OUTBUF, OUTCOUNT, 37
    DATATYPE, IERROR)                                             38
CHARACTER*(*) DATAREP                                           39
<type> INBUF(*), OUTBUF(*)                                       40
INTEGER(KIND=MPI_ADDRESS_KIND) INSIZE, POSITION                   41
INTEGER OUTCOUNT, DATATYPE, IERROR                             42

```

A.5.4 Collective Communication Fortran Bindings

1 MPI_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE, COMM,
2 IERROR)

3 <type> SENDBUF(*), RECVBUF(*)

4 INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, IERROR

5 MPI_ALLGATHER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE,
6 COMM, INFO, REQUEST, IERROR)

7 <type> SENDBUF(*), RECVBUF(*)

8 INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, INFO, REQUEST,
9 IERROR

10 MPI_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNTS, DISPLS,
11 RECVTYPE, COMM, IERROR)

12 <type> SENDBUF(*), RECVBUF(*)

13 INTEGER SENDCOUNT, SENDTYPE, REVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
14 IERROR

15 MPI_ALLGATHERV_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNTS, DISPLS,
16 RECVTYPE, COMM, INFO, REQUEST, IERROR)

17 <type> SENDBUF(*), RECVBUF(*)

18 INTEGER SENDCOUNT, SENDTYPE, REVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
19 INFO, REQUEST, IERROR

20 MPI_ALLREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)

21 <type> SENDBUF(*), RECVBUF(*)

22 INTEGER COUNT, DATATYPE, OP, COMM, IERROR

23 MPI_ALLREDUCE_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, INFO, REQUEST,
24 IERROR)

25 <type> SENDBUF(*), RECVBUF(*)

26 INTEGER COUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR

27 MPI_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE, COMM,
28 IERROR)

29 <type> SENDBUF(*), RECVBUF(*)

30 INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, IERROR

31 MPI_ALLTOALL_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE,
32 COMM, INFO, REQUEST, IERROR)

33 <type> SENDBUF(*), RECVBUF(*)

34 INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, INFO, REQUEST,
35 IERROR

36 MPI_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, REVCOUNTS,
37 RDISPLS, RECVTYPE, COMM, IERROR)

38 <type> SENDBUF(*), RECVBUF(*)

39 INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, REVCOUNTS(*), RDISPLS(*),
40 RECVTYPE, COMM, IERROR

41 MPI_ALLTOALLV_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, REVCOUNTS,
42 RDISPLS, RECVTYPE, COMM, INFO, REQUEST, IERROR)

```

<type> SENDBUF(*), RECVBUF(*) 1
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*), 2
RECVTYPE, COMM, INFO, REQUEST, IERROR 3
MPI_ALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF, RECVCOUNTS, 4
RDISPLS, RECVTYPES, COMM, IERROR) 5
<type> SENDBUF(*), RECVBUF(*) 6
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPES(*), RECVCOUNTS(*), RDISPLS(*), 7
RECVTYPES(*), COMM, IERROR 8
MPI_ALLTOALLW_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF, 9
RECVCOUNTS, RDISPLS, RECVTYPES, COMM, INFO, REQUEST, IERROR) 10
<type> SENDBUF(*), RECVBUF(*) 11
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPES(*), RECVCOUNTS(*), RDISPLS(*), 12
RECVTYPES(*), COMM, INFO, REQUEST, IERROR 13
MPI_BARRIER(COMM, IERROR) 14
INTEGER COMM, IERROR 15
MPI_BARRIER_INIT(COMM, INFO, REQUEST, IERROR) 16
INTEGER COMM, INFO, REQUEST, IERROR 17
MPI_BCAST(BUFFER, COUNT, DATATYPE, ROOT, COMM, IERROR) 18
<type> BUFFER(*) 19
INTEGER COUNT, DATATYPE, ROOT, COMM, IERROR 20
MPI_BCAST_INIT(BUFFER, COUNT, DATATYPE, ROOT, COMM, INFO, REQUEST, IERROR) 21
<type> BUFFER(*) 22
INTEGER COUNT, DATATYPE, ROOT, COMM, INFO, REQUEST, IERROR 23
MPI_EXSCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR) 24
<type> SENDBUF(*), RECVBUF(*) 25
INTEGER COUNT, DATATYPE, OP, COMM, IERROR 26
MPI_EXSCAN_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, INFO, REQUEST, 27
IERROR) 28
<type> SENDBUF(*), RECVBUF(*) 29
INTEGER COUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR 30
MPI_GATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, ROOT, 31
COMM, IERROR) 32
<type> SENDBUF(*), RECVBUF(*) 33
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, IERROR 34
MPI_GATHER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, 35
ROOT, COMM, INFO, REQUEST, IERROR) 36
<type> SENDBUF(*), RECVBUF(*) 37
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, INFO, 38
REQUEST, IERROR 39
MPI_GATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS, 40
RECVTYPE, ROOT, COMM, IERROR) 41
<type> SENDBUF(*), RECVBUF(*) 42

```

```
1     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT,
2         COMM, IERROR
3
4     MPI_GATHERV_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
5         RECVTYPE, ROOT, COMM, INFO, REQUEST, IERROR)
6     <type> SENDBUF(*), RECVBUF(*)
7     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT,
8         COMM, INFO, REQUEST, IERROR
9
10    MPI_IALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
11        COMM, REQUEST, IERROR)
12    <type> SENDBUF(*), RECVBUF(*)
13    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, REQUEST, IERROR
14
15    MPI_IALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
16        RECVTYPE, COMM, REQUEST, IERROR)
17    <type> SENDBUF(*), RECVBUF(*)
18    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
19        REQUEST, IERROR
20
21    MPI_IALLREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, REQUEST, IERROR)
22    <type> SENDBUF(*), RECVBUF(*)
23    INTEGER COUNT, DATATYPE, OP, COMM, REQUEST, IERROR
24
25    MPI_IALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, COMM,
26        REQUEST, IERROR)
27    <type> SENDBUF(*), RECVBUF(*)
28    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, REQUEST, IERROR
29
30    MPI_IALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, RECVCOUNTS,
31        RDISPLS, RECVTYPE, COMM, REQUEST, IERROR)
32    <type> SENDBUF(*), RECVBUF(*)
33    INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*),
34        RECVTYPE, COMM, REQUEST, IERROR
35
36    MPI_IALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF, RECVCOUNTS,
37        RDISPLS, RECVTYPES, COMM, REQUEST, IERROR)
38    <type> SENDBUF(*), RECVBUF(*)
39    INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPES(*), RECVCOUNTS(*), RDISPLS(*),
40        RECVTYPES(*), COMM, REQUEST, IERROR
41
42    MPI_IBARRIER(COMM, REQUEST, IERROR)
43    INTEGER COMM, REQUEST, IERROR
44
45    MPI_IBCAST(BUFFER, COUNT, DATATYPE, ROOT, COMM, REQUEST, IERROR)
46    <type> BUFFER(*)
47    INTEGER COUNT, DATATYPE, ROOT, COMM, REQUEST, IERROR
48
49    MPI_IXSCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, REQUEST, IERROR)
50    <type> SENDBUF(*), RECVBUF(*)
51    INTEGER COUNT, DATATYPE, OP, COMM, REQUEST, IERROR
52
53    MPI_IGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, ROOT,
54        COMM, REQUEST, IERROR)
```

<type> SENDBUF(*), RECVBUF(*)	1
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, REQUEST,	2
IERROR	3
MPI_IGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,	4
RECVTYPE, ROOT, COMM, REQUEST, IERROR)	5
<type> SENDBUF(*), RECVBUF(*)	6
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT,	7
COMM, REQUEST, IERROR	8
	9
MPI_IREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, REQUEST, IERROR)	10
<type> SENDBUF(*), RECVBUF(*)	11
INTEGER COUNT, DATATYPE, OP, ROOT, COMM, REQUEST, IERROR	12
MPI_IREDUCE_SCATTER(SENDBUF, RECVBUF, RECVCOUNTS, DATATYPE, OP, COMM, REQUEST,	13
IERROR)	14
<type> SENDBUF(*), RECVBUF(*)	15
INTEGER RECVCOUNTS(*), DATATYPE, OP, COMM, REQUEST, IERROR	16
	17
MPI_IREDUCE_SCATTER_BLOCK(SENDBUF, RECVBUF, RECVCOUNT, DATATYPE, OP, COMM,	18
REQUEST, IERROR)	19
<type> SENDBUF(*), RECVBUF(*)	20
INTEGER RECVCOUNT, DATATYPE, OP, COMM, REQUEST, IERROR	21
	22
MPI_ISCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, REQUEST, IERROR)	23
<type> SENDBUF(*), RECVBUF(*)	24
INTEGER COUNT, DATATYPE, OP, COMM, REQUEST, IERROR	25
MPI_ISCATTER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, ROOT,	26
COMM, REQUEST, IERROR)	27
<type> SENDBUF(*), RECVBUF(*)	28
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, REQUEST,	29
IERROR	30
	31
MPI_ISCATTERV(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,	32
RECVTYPE, ROOT, COMM, REQUEST, IERROR)	33
<type> SENDBUF(*), RECVBUF(*)	34
INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,	35
COMM, REQUEST, IERROR	36
	37
MPI_OP_COMMUTATIVE(OP, COMMUTE, IERROR)	38
INTEGER OP, IERROR	39
LOGICAL COMMUTE	40
	41
MPI_OP_CREATE(USER_FN, COMMUTE, OP, IERROR)	42
EXTERNAL USER_FN	43
LOGICAL COMMUTE	44
INTEGER OP, IERROR	45
	46
MPI_OP_FREE(OP, IERROR)	47
INTEGER OP, IERROR	48
MPI_REDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, IERROR)	
<type> SENDBUF(*), RECVBUF(*)	

```
1     INTEGER COUNT, DATATYPE, OP, ROOT, COMM, IERROR
2
3     MPI_REDUCE_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, INFO,
4         REQUEST, IERROR)
5     <type> SENDBUF(*), RECVBUF(*)
6     INTEGER COUNT, DATATYPE, OP, ROOT, COMM, INFO, REQUEST, IERROR
7
8     MPI_REDUCE_LOCAL(INBUF, INOUTBUF, COUNT, DATATYPE, OP, IERROR)
9     <type> INBUF(*), INOUTBUF(*)
10    INTEGER COUNT, DATATYPE, OP, IERROR
11
12    MPI_REDUCE_SCATTER(SENDBUF, RECVBUF, RECVCOUNTS, DATATYPE, OP, COMM, IERROR)
13    <type> SENDBUF(*), RECVBUF(*)
14    INTEGER RECVCOUNTS(*), DATATYPE, OP, COMM, IERROR
15
16    MPI_REDUCE_SCATTER_BLOCK(SENDBUF, RECVBUF, RECVCOUNT, DATATYPE, OP, COMM,
17        IERROR)
18    <type> SENDBUF(*), RECVBUF(*)
19    INTEGER RECVCOUNT, DATATYPE, OP, COMM, IERROR
20
21    MPI_REDUCE_SCATTER_BLOCK_INIT(SENDBUF, RECVBUF, RECVCOUNT, DATATYPE, OP, COMM,
22        INFO, REQUEST, IERROR)
23    <type> SENDBUF(*), RECVBUF(*)
24    INTEGER RECVCOUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR
25
26    MPI_REDUCE_SCATTER_INIT(SENDBUF, RECVBUF, RECVCOUNTS, DATATYPE, OP, COMM, INFO,
27        REQUEST, IERROR)
28    <type> SENDBUF(*), RECVBUF(*)
29    INTEGER RECVCOUNTS(*), DATATYPE, OP, COMM, INFO, REQUEST, IERROR
30
31    MPI_SCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
32    <type> SENDBUF(*), RECVBUF(*)
33    INTEGER COUNT, DATATYPE, OP, COMM, IERROR
34
35    MPI_SCAN_INIT(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, INFO, REQUEST,
36        IERROR)
37    <type> SENDBUF(*), RECVBUF(*)
38    INTEGER COUNT, DATATYPE, OP, COMM, INFO, REQUEST, IERROR
39
40    MPI_SCATTER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, ROOT,
41        COMM, IERROR)
42    <type> SENDBUF(*), RECVBUF(*)
43    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, IERROR
44
45    MPI_SCATTER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
46        ROOT, COMM, INFO, REQUEST, IERROR)
47    <type> SENDBUF(*), RECVBUF(*)
48    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, INFO,
49        REQUEST, IERROR
50
51    MPI_SCATTERV(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,
52        RECVTYPE, ROOT, COMM, IERROR)
53    <type> SENDBUF(*), RECVBUF(*)
```

```

    INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,
        COMM, IERROR
1
2
3
MPI_SCATTERV_INIT(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

```

INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,
 COMM, IERROR
 MPI_SCATTERV_INIT(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,
 RECVTYPE, ROOT, COMM, INFO, REQUEST, IERROR)
 <type> SENDBUF(*), RECVBUF(*)
 INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,
 COMM, INFO, REQUEST, IERROR
 MPI_TYPE_GET_VALUE_INDEX(VALUE_TYPE, INDEX_TYPE, PAIR_TYPE, IERROR)
 INTEGER VALUE_TYPE, INDEX_TYPE, PAIR_TYPE, IERROR

A.5.5 Groups, Contexts, Communicators, and Caching Fortran Bindings

```

MPI_COMM_COMPARE(COMM1, COMM2, RESULT, IERROR)
    INTEGER COMM1, COMM2, RESULT, IERROR
MPI_COMM_CREATE(COMM, GROUP, NEWCOMM, IERROR)
    INTEGER COMM, GROUP, NEWCOMM, IERROR
MPI_COMM_CREATE_FROM_GROUP(GROUP, STRINGTAG, INFO, ERRHANDLER, NEWCOMM, IERROR)
    INTEGER GROUP, INFO, ERRHANDLER, NEWCOMM, IERROR
    CHARACTER*(*) STRINGTAG
MPI_COMM_CREATE_GROUP(COMM, GROUP, TAG, NEWCOMM, IERROR)
    INTEGER COMM, GROUP, TAG, NEWCOMM, IERROR
MPI_COMM_CREATE_KEYVAL(COMM_COPY_ATTR_FN, COMM_DELETE_ATTR_FN, COMM_KEYVAL,
    EXTRA_STATE, IERROR)
    EXTERNAL COMM_COPY_ATTR_FN, COMM_DELETE_ATTR_FN
    INTEGER COMM_KEYVAL, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
MPI_COMM_DELETE_ATTR(COMM, COMM_KEYVAL, IERROR)
    INTEGER COMM, COMM_KEYVAL, IERROR
MPI_COMM_DUP(COMM, NEWCOMM, IERROR)
    INTEGER COMM, NEWCOMM, IERROR
MPI_COMM_DUP_FN(OLDCOMM, COMM_KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
    ATTRIBUTE_VAL_OUT, FLAG, IERROR)
    INTEGER OLDCOMM, COMM_KEYVAL, IERROR
    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
    ATTRIBUTE_VAL_OUT
    LOGICAL FLAG
MPI_COMM_DUP_WITH_INFO(COMM, INFO, NEWCOMM, IERROR)
    INTEGER COMM, INFO, NEWCOMM, IERROR
MPI_COMM_FREE(COMM, IERROR)
    INTEGER COMM, IERROR
MPI_COMM_FREE_KEYVAL(COMM_KEYVAL, IERROR)
    INTEGER COMM_KEYVAL, IERROR

```

```
1 MPI_COMM_GET_ATTR(COMM, COMM_KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)
2   INTEGER COMM, COMM_KEYVAL, IERROR
3   INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL
4   LOGICAL FLAG
5
6 MPI_COMM_GET_INFO(COMM, INFO_USED, IERROR)
7   INTEGER COMM, INFO_USED, IERROR
8
9 MPI_COMM_GET_NAME(COMM, COMM_NAME, RESULTLEN, IERROR)
10  INTEGER COMM, RESULTLEN, IERROR
11  CHARACTER*(*) COMM_NAME
12
13 MPI_COMM_GROUP(COMM, GROUP, IERROR)
14  INTEGER COMM, GROUP, IERROR
15
16 MPI_COMM_IDUP(COMM, NEWCOMM, REQUEST, IERROR)
17  INTEGER COMM, NEWCOMM, REQUEST, IERROR
18
19 MPI_COMM_IDUP_WITH_INFO(COMM, INFO, NEWCOMM, REQUEST, IERROR)
20  INTEGER COMM, INFO, NEWCOMM, REQUEST, IERROR
21
22 MPI_COMM_NULL_COPY_FN(OLDCOMM, COMM_KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
23   ATTRIBUTE_VAL_OUT, FLAG, IERROR)
24  INTEGER OLDCOMM, COMM_KEYVAL, IERROR
25  INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
26   ATTRIBUTE_VAL_OUT
27  LOGICAL FLAG
28
29 MPI_COMM_NULL_DELETE_FN(COMM, COMM_KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERROR)
30  INTEGER COMM, COMM_KEYVAL, IERROR
31  INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE
32
33 MPI_COMM_RANK(COMM, RANK, IERROR)
34  INTEGER COMM, RANK, IERROR
35
36 MPI_COMM_REMOTE_GROUP(COMM, GROUP, IERROR)
37  INTEGER COMM, GROUP, IERROR
38
39 MPI_COMM_REMOTE_SIZE(COMM, SIZE, IERROR)
40  INTEGER COMM, SIZE, IERROR
41
42 MPI_COMM_SET_ATTR(COMM, COMM_KEYVAL, ATTRIBUTE_VAL, IERROR)
43  INTEGER COMM, COMM_KEYVAL, IERROR
44  INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL
45
46 MPI_COMM_SET_INFO(COMM, INFO, IERROR)
47  INTEGER COMM, INFO, IERROR
48
49 MPI_COMM_SET_NAME(COMM, COMM_NAME, IERROR)
50  INTEGER COMM, IERROR
51  CHARACTER*(*) COMM_NAME
52
53 MPI_COMM_SIZE(COMM, SIZE, IERROR)
54  INTEGER COMM, SIZE, IERROR
55
56 MPI_COMM_SPLIT(COMM, COLOR, KEY, NEWCOMM, IERROR)
```


INTEGER COMM, COLOR, KEY, NEWCOMM, IERROR	1
MPI_COMM_SPLIT_TYPE(COMM, SPLIT_TYPE, KEY, INFO, NEWCOMM, IERROR)	2
INTEGER COMM, SPLIT_TYPE, KEY, INFO, NEWCOMM, IERROR	3
MPI_COMM_TEST_INTER(COMM, FLAG, IERROR)	4
INTEGER COMM, IERROR	5
LOGICAL FLAG	6
MPI_GROUP_COMPARE(GROUP1, GROUP2, RESULT, IERROR)	7
INTEGER GROUP1, GROUP2, RESULT, IERROR	8
MPI_GROUP_DIFFERENCE(GROUP1, GROUP2, NEWGROUP, IERROR)	9
INTEGER GROUP1, GROUP2, NEWGROUP, IERROR	10
MPI_GROUP_EXCL(GROUP, N, RANKS, NEWGROUP, IERROR)	11
INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR	12
MPI_GROUP_FREE(GROUP, IERROR)	13
INTEGER GROUP, IERROR	14
MPI_GROUP_FROM_SESSION_PSET(SESSION, PSET_NAME, NEWGROUP, IERROR)	15
INTEGER SESSION, NEWGROUP, IERROR	16
CHARACTER*(*) PSET_NAME	17
MPI_GROUP_INCL(GROUP, N, RANKS, NEWGROUP, IERROR)	18
INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR	19
MPI_GROUP_INTERSECTION(GROUP1, GROUP2, NEWGROUP, IERROR)	20
INTEGER GROUP1, GROUP2, NEWGROUP, IERROR	21
MPI_GROUP_RANGE_EXCL(GROUP, N, RANGES, NEWGROUP, IERROR)	22
INTEGER GROUP, N, RANGES(3, *), NEWGROUP, IERROR	23
MPI_GROUP_RANGE_INCL(GROUP, N, RANGES, NEWGROUP, IERROR)	24
INTEGER GROUP, N, RANGES(3, *), NEWGROUP, IERROR	25
MPI_GROUP_RANK(GROUP, RANK, IERROR)	26
INTEGER GROUP, RANK, IERROR	27
MPI_GROUP_SIZE(GROUP, SIZE, IERROR)	28
INTEGER GROUP, SIZE, IERROR	29
MPI_GROUP_TRANSLATE_RANKS(GROUP1, N, RANKS1, GROUP2, RANKS2, IERROR)	30
INTEGER GROUP1, N, RANKS1(*), GROUP2, RANKS2(*), IERROR	31
MPI_GROUP_UNION(GROUP1, GROUP2, NEWGROUP, IERROR)	32
INTEGER GROUP1, GROUP2, NEWGROUP, IERROR	33
MPI_INTERCOMM_CREATE(LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG, NEWINTERCOMM, IERROR)	34
INTEGER LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG, NEWINTERCOMM, IERROR	35
MPI_INTERCOMM_CREATE_FROM_GROUPS(LOCAL_GROUP, LOCAL_LEADER, REMOTE_GROUP, REMOTE_LEADER, STRINGTAG, INFO, ERRHANDLER, NEWINTERCOMM, IERROR)	36
INTEGER LOCAL_GROUP, LOCAL_LEADER, REMOTE_GROUP, REMOTE_LEADER, INFO, ERRHANDLER, NEWINTERCOMM, IERROR	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48

```
1     CHARACTER*(*) STRINGTAG
2
3     MPI_INTERCOMM_MERGE(INTERCOMM, HIGH, NEWINTRACOMM, IERROR)
4     INTEGER INTERCOMM, NEWINTRACOMM, IERROR
5     LOGICAL HIGH
6
7     MPI_TYPE_CREATE_KEYVAL(TYPE_COPY_ATTR_FN, TYPE_DELETE_ATTR_FN, TYPE_KEYVAL,
8         EXTRA_STATE, IERROR)
9     EXTERNAL TYPE_COPY_ATTR_FN, TYPE_DELETE_ATTR_FN
10    INTEGER TYPE_KEYVAL, IERROR
11    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
12
13    MPI_TYPE_DELETE_ATTR(DATATYPE, TYPE_KEYVAL, IERROR)
14    INTEGER DATATYPE, TYPE_KEYVAL, IERROR
15
16    MPI_TYPE_DUP_FN(OLDTYPE, TYPE_KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
17        ATTRIBUTE_VAL_OUT, FLAG, IERROR)
18    INTEGER OLDTYPE, TYPE_KEYVAL, IERROR
19    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
20        ATTRIBUTE_VAL_OUT
21    LOGICAL FLAG
22
23    MPI_TYPE_FREE_KEYVAL(TYPE_KEYVAL, IERROR)
24    INTEGER TYPE_KEYVAL, IERROR
25
26    MPI_TYPE_GET_ATTR(DATATYPE, TYPE_KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)
27    INTEGER DATATYPE, TYPE_KEYVAL, IERROR
28    INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL
29    LOGICAL FLAG
30
31    MPI_TYPE_GET_NAME(DATATYPE, TYPE_NAME, RESULTLEN, IERROR)
32    INTEGER DATATYPE, RESULTLEN, IERROR
33    CHARACTER*(*) TYPE_NAME
34
35    MPI_TYPE_NULL_COPY_FN(OLDTYPE, TYPE_KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
36        ATTRIBUTE_VAL_OUT, FLAG, IERROR)
37    INTEGER OLDTYPE, TYPE_KEYVAL, IERROR
38    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN,
39        ATTRIBUTE_VAL_OUT
40    LOGICAL FLAG
41
42    MPI_TYPE_NULL_DELETE_FN(DATATYPE, TYPE_KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE,
43        IERROR)
44    INTEGER DATATYPE, TYPE_KEYVAL, IERROR
45    INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE
46
47    MPI_TYPE_SET_ATTR(DATATYPE, TYPE_KEYVAL, ATTRIBUTE_VAL, IERROR)
48    INTEGER DATATYPE, TYPE_KEYVAL, IERROR
49    INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL
50
51    MPI_TYPE_SET_NAME(DATATYPE, TYPE_NAME, IERROR)
52    INTEGER DATATYPE, IERROR
53    CHARACTER*(*) TYPE_NAME
```

MPI_WIN_CREATE_KEYVAL(WIN_COPY_ATTR_FN, WIN_DELETE_ATTR_FN, WIN_KEYVAL, EXTRA_STATE, IERROR)	1
EXTERNAL WIN_COPY_ATTR_FN, WIN_DELETE_ATTR_FN	2
INTEGER WIN_KEYVAL, IERROR	3
INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE	4
MPI_WIN_DELETE_ATTR(WIN, WIN_KEYVAL, IERROR)	5
INTEGER WIN, WIN_KEYVAL, IERROR	6
MPI_WIN_DUP_FN(OLDWIN, WIN_KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)	7
INTEGER OLDWIN, WIN_KEYVAL, IERROR	8
INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT	9
LOGICAL FLAG	10
MPI_WIN_FREE_KEYVAL(WIN_KEYVAL, IERROR)	11
INTEGER WIN_KEYVAL, IERROR	12
MPI_WIN_GET_ATTR(WIN, WIN_KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)	13
INTEGER WIN, WIN_KEYVAL, IERROR	14
INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL	15
LOGICAL FLAG	16
MPI_WIN_GET_NAME(WIN, WIN_NAME, RESULTLEN, IERROR)	17
INTEGER WIN, RESULTLEN, IERROR	18
CHARACTER*(*) WIN_NAME	19
MPI_WIN_NULL_COPY_FN(OLDWIN, WIN_KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT, FLAG, IERROR)	20
INTEGER OLDWIN, WIN_KEYVAL, IERROR	21
INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT	22
LOGICAL FLAG	23
MPI_WIN_NULL_DELETE_FN(WIN, WIN_KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERROR)	24
INTEGER WIN, WIN_KEYVAL, IERROR	25
INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL, EXTRA_STATE	26
MPI_WIN_SET_ATTR(WIN, WIN_KEYVAL, ATTRIBUTE_VAL, IERROR)	27
INTEGER WIN, WIN_KEYVAL, IERROR	28
INTEGER(KIND=MPI_ADDRESS_KIND) ATTRIBUTE_VAL	29
MPI_WIN_SET_NAME(WIN, WIN_NAME, IERROR)	30
INTEGER WIN, IERROR	31
CHARACTER*(*) WIN_NAME	32
A.5.6 Virtual Topologies for MPI Processes Fortran Bindings	33
MPI_CART_COORDS(COMM, RANK, MAXDIMS, COORDS, IERROR)	34
INTEGER COMM, RANK, MAXDIMS, COORDS(*), IERROR	35
MPI_CART_CREATE(COMM_OLD, NDIMS, DIMS, PERIODS, REORDER, COMM_CART, IERROR)	36
	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48

```
1     INTEGER COMM_OLD, NDIMS, DIMS(*), COMM_CART, IERROR
2     LOGICAL PERIODS(*), REORDER
3
4     MPI_CART_GET(COMM, MAXDIMS, DIMS, PERIODS, COORDS, IERROR)
5     INTEGER COMM, MAXDIMS, DIMS(*), COORDS(*), IERROR
6     LOGICAL PERIODS(*)
7
8     MPI_CART_MAP(COMM, NDIMS, DIMS, PERIODS, NEWRANK, IERROR)
9     INTEGER COMM, NDIMS, DIMS(*), NEWRANK, IERROR
10    LOGICAL PERIODS(*)
11
12    MPI_CART_RANK(COMM, COORDS, RANK, IERROR)
13    INTEGER COMM, COORDS(*), RANK, IERROR
14
15    MPI_CART_SHIFT(COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR)
16    INTEGER COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR
17
18    MPI_CART_SUB(COMM, REMAIN_DIMS, NEWCOMM, IERROR)
19    INTEGER COMM, NEWCOMM, IERROR
20    LOGICAL REMAIN_DIMS(*)
21
22    MPI_CARTDIM_GET(COMM, NDIMS, IERROR)
23    INTEGER COMM, NDIMS, IERROR
24
25    MPI_DIMS_CREATE(NNODES, NDIMS, DIMS, IERROR)
26    INTEGER NNODES, NDIMS, DIMS(*), IERROR
27
28    MPI_DIST_GRAPH_CREATE(COMM_OLD, N, SOURCES, DEGREES, DESTINATIONS, WEIGHTS,
29    INFO, REORDER, COMM_DIST_GRAPH, IERROR)
30    INTEGER COMM_OLD, N, SOURCES(*), DEGREES(*), DESTINATIONS(*), WEIGHTS(*),
31    INFO, COMM_DIST_GRAPH, IERROR
32    LOGICAL REORDER
33
34    MPI_DIST_GRAPH_CREATE_ADJACENT(COMM_OLD, INDEGREE, SOURCES, SOURCEWEIGHTS,
35    OUTDEGREE, DESTINATIONS, DESTWEIGHTS, INFO, REORDER,
36    COMM_DIST_GRAPH, IERROR)
37    INTEGER COMM_OLD, INDEGREE, SOURCES(*), SOURCEWEIGHTS(*), OUTDEGREE,
38    DESTINATIONS(*), DESTWEIGHTS(*), INFO, COMM_DIST_GRAPH, IERROR
39    LOGICAL REORDER
40
41    MPI_DIST_GRAPH_NEIGHBORS(COMM, MAXINDEGREE, SOURCES, SOURCEWEIGHTS,
42    MAXOUTDEGREE, DESTINATIONS, DESTWEIGHTS, IERROR)
43    INTEGER COMM, MAXINDEGREE, SOURCES(*), SOURCEWEIGHTS(*), MAXOUTDEGREE,
44    DESTINATIONS(*), DESTWEIGHTS(*), IERROR
45
46    MPI_DIST_GRAPH_NEIGHBORS_COUNT(COMM, INDEGREE, OUTDEGREE, WEIGHTED, IERROR)
47    INTEGER COMM, INDEGREE, OUTDEGREE, IERROR
48    LOGICAL WEIGHTED
49
50    MPI_GRAPH_CREATE(COMM_OLD, NNODES, INDEX, EDGES, REORDER, COMM_GRAPH, IERROR)
51    INTEGER COMM_OLD, NNODES, INDEX(*), EDGES(*), COMM_GRAPH, IERROR
52    LOGICAL REORDER
53
54    MPI_GRAPH_GET(COMM, MAXINDEX, MAXEDGES, INDEX, EDGES, IERROR)
55    INTEGER COMM, MAXINDEX, MAXEDGES, INDEX(*), EDGES(*), IERROR
```

MPI_GRAPH_MAP(COMM, NNODES, INDEX, EDGES, NEWRANK, IERROR)	1
INTEGER COMM, NNODES, INDEX(*), EDGES(*), NEWRANK, IERROR	2
	3
MPI_GRAPH_NEIGHBORS(COMM, RANK, MAXNEIGHBORS, NEIGHBORS, IERROR)	4
INTEGER COMM, RANK, MAXNEIGHBORS, NEIGHBORS(*), IERROR	5
	6
MPI_GRAPH_NEIGHBORS_COUNT(COMM, RANK, NNEIGHBORS, IERROR)	7
INTEGER COMM, RANK, NNEIGHBORS, IERROR	8
	9
MPI_GRAPHDIMS_GET(COMM, NNODES, NEDGES, IERROR)	10
INTEGER COMM, NNODES, NEDGES, IERROR	11
	12
MPI_INEIGHBOR_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUFF, RECVCOUNT,	13
RECVMODE, COMM, REQUEST, IERROR)	14
<type> SENDBUF(*), RECVBUFF(*)	15
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVMODE, COMM, REQUEST, IERROR	16
	17
MPI_INEIGHBOR_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUFF, RECVCOUNTS,	18
DISPLS, RECVMODE, COMM, REQUEST, IERROR)	19
<type> SENDBUF(*), RECVBUFF(*)	20
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVMODE, COMM,	21
REQUEST, IERROR	22
	23
MPI_INEIGHBOR_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUFF, RECVCOUNT,	24
RECVMODE, COMM, REQUEST, IERROR)	25
<type> SENDBUF(*), RECVBUFF(*)	26
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVMODE, COMM, REQUEST, IERROR	27
	28
MPI_INEIGHBOR_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUFF,	29
RECVCOUNTS, RDISPLS, RECVMODE, COMM, REQUEST, IERROR)	30
<type> SENDBUF(*), RECVBUFF(*)	31
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*),	32
RECVMODE, COMM, REQUEST, IERROR	33
	34
MPI_INEIGHBOR_ALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUFF,	35
RECVCOUNTS, RDISPLS, RECVMODES, COMM, REQUEST, IERROR)	36
<type> SENDBUF(*), RECVBUFF(*)	37
INTEGER SENDCOUNTS(*), SENDTYPES(*), RECVCOUNTS(*), RECVMODES(*), COMM,	38
REQUEST, IERROR	39
INTEGER(KIND=MPI_ADDRESS_KIND) SDISPLS(*), RDISPLS(*)	40
	41
MPI_NEIGHBOR_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUFF, RECVCOUNT,	42
RECVMODE, COMM, IERROR)	43
<type> SENDBUF(*), RECVBUFF(*)	44
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVMODE, COMM, IERROR	45
	46
MPI_NEIGHBOR_ALLGATHER_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUFF, RECVCOUNT,	47
RECVMODE, COMM, INFO, REQUEST, IERROR)	48
<type> SENDBUF(*), RECVBUFF(*)	
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVMODE, COMM, INFO, REQUEST,	
IERROR	
MPI_NEIGHBOR_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUFF, RECVCOUNTS,	
DISPLS, RECVMODE, COMM, IERROR)	

```

1      <type> SENDBUF(*), RECVBUF(*)
2      INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
3          IERROR
4
5      MPI_NEIGHBOR_ALLGATHERV_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS,
6          DISPLS, RECVTYPE, COMM, INFO, REQUEST, IERROR)
7      <type> SENDBUF(*), RECVBUF(*)
8      INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
9          INFO, REQUEST, IERROR
10
11     MPI_NEIGHBOR_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT,
12         RECVTYPE, COMM, IERROR)
13     <type> SENDBUF(*), RECVBUF(*)
14     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, IERROR
15
16     MPI_NEIGHBOR_ALLTOALL_INIT(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT,
17         RECVTYPE, COMM, INFO, REQUEST, IERROR)
18     <type> SENDBUF(*), RECVBUF(*)
19     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, INFO, REQUEST,
20         IERROR
21
22     MPI_NEIGHBOR_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF,
23         RECVCOUNTS, RDISPLS, RECVTYPE, COMM, IERROR)
24     <type> SENDBUF(*), RECVBUF(*)
25     INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*),
26         RECVTYPE, COMM, IERROR
27
28     MPI_NEIGHBOR_ALLTOALLV_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF,
29         RECVCOUNTS, RDISPLS, RECVTYPE, COMM, INFO, REQUEST, IERROR)
30     <type> SENDBUF(*), RECVBUF(*)
31     INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*),
32         RECVTYPE, COMM, INFO, REQUEST, IERROR
33
34     MPI_NEIGHBOR_ALLTOALLW(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF,
35         RECVCOUNTS, RDISPLS, RECVTYPES, COMM, IERROR)
36     <type> SENDBUF(*), RECVBUF(*)
37     INTEGER SENDCOUNTS(*), SENDTYPES(*), RECVCOUNTS(*), RECVTYPES(*), COMM,
38         IERROR
39     INTEGER(KIND=MPI_ADDRESS_KIND) SDISPLS(*), RDISPLS(*)
40
41     MPI_NEIGHBOR_ALLTOALLW_INIT(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPES, RECVBUF,
42         RECVCOUNTS, RDISPLS, RECVTYPES, COMM, INFO, REQUEST, IERROR)
43     <type> SENDBUF(*), RECVBUF(*)
44     INTEGER SENDCOUNTS(*), SENDTYPES(*), RECVCOUNTS(*), RECVTYPES(*), COMM,
45         INFO, REQUEST, IERROR
46     INTEGER(KIND=MPI_ADDRESS_KIND) SDISPLS(*), RDISPLS(*)
47
48     MPI_TOPO_TEST(COMM, STATUS, IERROR)
49     INTEGER COMM, STATUS, IERROR

```

A.5.7 MPI Environmental Management Fortran Bindings

DOUBLE PRECISION MPI_WTICK()	1
DOUBLE PRECISION MPI_WTIME()	2
MPI_ADD_ERROR_CLASS(ERRORCLASS, IERROR)	3
INTEGER ERRORCLASS, IERROR	4
MPI_ADD_ERROR_CODE(ERRORCLASS, ERRORCODE, IERROR)	5
INTEGER ERRORCLASS, ERRORCODE, IERROR	6
MPI_ADD_ERROR_STRING(ERRORCODE, STRING, IERROR)	7
INTEGER ERRORCODE, IERROR	8
CHARACTER*(*) STRING	9
MPI_ALLOC_MEM(SIZE, INFO, BASEPTR, IERROR)	10
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR	11
INTEGER INFO, IERROR	12
If the Fortran compiler provides TYPE(C_PTR), then overloaded by:	13
INTERFACE MPI_ALLOC_MEM	14
SUBROUTINE MPI_ALLOC_MEM(SIZE, INFO, BASEPTR, IERROR)	15
IMPORT :: MPI_ADDRESS_KIND	16
INTEGER :: INFO, IERROR	17
INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR	18
END SUBROUTINE	19
SUBROUTINE MPI_ALLOC_MEM_CPTR(SIZE, INFO, BASEPTR, IERROR)	20
USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR	21
IMPORT :: MPI_ADDRESS_KIND	22
INTEGER :: INFO, IERROR	23
INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE	24
TYPE(C_PTR) :: BASEPTR	25
END SUBROUTINE	26
END INTERFACE	27
MPI_COMM_CALL_ERRHANDLER(COMM, ERRORCODE, IERROR)	28
INTEGER COMM, ERRORCODE, IERROR	29
MPI_COMM_CREATE_ERRHANDLER(COMM_ERRHANDLER_FN, ERRHANDLER, IERROR)	30
EXTERNAL COMM_ERRHANDLER_FN	31
INTEGER ERRHANDLER, IERROR	32
MPI_COMM_GET_ERRHANDLER(COMM, ERRHANDLER, IERROR)	33
INTEGER COMM, ERRHANDLER, IERROR	34
MPI_COMM_SET_ERRHANDLER(COMM, ERRHANDLER, IERROR)	35
INTEGER COMM, ERRHANDLER, IERROR	36
MPI_ERRHANDLER_FREE(ERRHANDLER, IERROR)	37
INTEGER ERRHANDLER, IERROR	38
MPI_ERROR_CLASS(ERRORCODE, ERRORCLASS, IERROR)	39
INTEGER ERRORCODE, ERRORCLASS, IERROR	40
MPI_ERROR_STRING(ERRORCODE, STRING, RESULTLEN, IERROR)	41
INTEGER ERRORCODE, IERROR	42
CHARACTER*(*) STRING	43
INTEGER RESULTLEN	44

```
1     INTEGER ERRORCODE, RESULTLEN, IERROR
2     CHARACTER*(*) STRING
3
4     MPI_FILE_CALL_ERRHANDLER(FH, ERRORCODE, IERROR)
5     INTEGER FH, ERRORCODE, IERROR
6
7     MPI_FILE_CREATE_ERRHANDLER(FILE_ERRHANDLER_FN, ERRHANDLER, IERROR)
8     EXTERNAL FILE_ERRHANDLER_FN
9     INTEGER ERRHANDLER, IERROR
10
11    MPI_FILE_GET_ERRHANDLER(FILE, ERRHANDLER, IERROR)
12    INTEGER FILE, ERRHANDLER, IERROR
13
14    MPI_FILE_SET_ERRHANDLER(FILE, ERRHANDLER, IERROR)
15    INTEGER FILE, ERRHANDLER, IERROR
16
17    MPI_FREE_MEM(BASE, IERROR)
18    <type> BASE(*)
19    INTEGER IERROR
20
21    MPI_GET_HW_RESOURCE_INFO(HW_INFO, IERROR)
22    INTEGER HW_INFO, IERROR
23
24    MPI_GET_LIBRARY_VERSION(VERSION, RESULTLEN, IERROR)
25    CHARACTER*(*) VERSION
26    INTEGER RESULTLEN, IERROR
27
28    MPI_GET_PROCESSOR_NAME(NAME, RESULTLEN, IERROR)
29    CHARACTER*(*) NAME
30    INTEGER RESULTLEN, IERROR
31
32    MPI_GET_VERSION(VERSION, SUBVERSION, IERROR)
33    INTEGER VERSION, SUBVERSION, IERROR
34
35    MPI_REMOVE_ERROR_CLASS(ERRORCLASS, IERROR)
36    INTEGER ERRORCLASS, IERROR
37
38    MPI_REMOVE_ERROR_CODE(ERRORCODE, IERROR)
39    INTEGER ERRORCODE, IERROR
40
41    MPI_REMOVE_ERROR_STRING(ERRORCODE, IERROR)
42    INTEGER ERRORCODE, IERROR
43
44    MPI_SESSION_CALL_ERRHANDLER(SESSION, ERRORCODE, IERROR)
45    INTEGER SESSION, ERRORCODE, IERROR
46
47    MPI_SESSION_CREATE_ERRHANDLER(SESSION_ERRHANDLER_FN, ERRHANDLER, IERROR)
48    EXTERNAL SESSION_ERRHANDLER_FN
49    INTEGER ERRHANDLER, IERROR
50
51    MPI_SESSION_GET_ERRHANDLER(SESSION, ERRHANDLER, IERROR)
52    INTEGER SESSION, ERRHANDLER, IERROR
53
54    MPI_SESSION_SET_ERRHANDLER(SESSION, ERRHANDLER, IERROR)
55    INTEGER SESSION, ERRHANDLER, IERROR
56
57    MPI_WIN_CALL_ERRHANDLER(WIN, ERRORCODE, IERROR)
```


INTEGER WIN, ERRORCODE, IERROR	1
MPI_WIN_CREATE_ERRHANDLER(WIN_ERRHANDLER_FN, ERRHANDLER, IERROR)	2
EXTERNAL WIN_ERRHANDLER_FN	3
INTEGER ERRHANDLER, IERROR	4
	5
MPI_WIN_GET_ERRHANDLER(WIN, ERRHANDLER, IERROR)	6
INTEGER WIN, ERRHANDLER, IERROR	7
	8
MPI_WIN_SET_ERRHANDLER(WIN, ERRHANDLER, IERROR)	9
INTEGER WIN, ERRHANDLER, IERROR	10
	11
A.5.8 The Info Object Fortran Bindings	12
MPI_INFO_CREATE(INFO, IERROR)	13
INTEGER INFO, IERROR	14
	15
MPI_INFO_CREATE_ENV(INFO, IERROR)	16
INTEGER INFO, IERROR	17
	18
MPI_INFO_DELETE(INFO, KEY, IERROR)	19
INTEGER INFO, IERROR	20
CHARACTER*(*) KEY	21
	22
MPI_INFO_DUP(INFO, NEWINFO, IERROR)	23
INTEGER INFO, NEWINFO, IERROR	24
	25
MPI_INFO_FREE(INFO, IERROR)	26
INTEGER INFO, IERROR	27
	28
MPI_INFO_GET_NKEYS(INFO, NKEYS, IERROR)	29
INTEGER INFO, NKEYS, IERROR	30
	31
MPI_INFO_GET_NTHKEY(INFO, N, KEY, IERROR)	32
INTEGER INFO, N, IERROR	33
CHARACTER*(*) KEY	34
	35
MPI_INFO_GET_STRING(INFO, KEY, BUFLen, VALUE, FLAG, IERROR)	36
INTEGER INFO, BUFLen, IERROR	37
CHARACTER*(*) KEY, VALUE	38
LOGICAL FLAG	39
	40
MPI_INFO_SET(INFO, KEY, VALUE, IERROR)	41
INTEGER INFO, IERROR	42
CHARACTER*(*) KEY, VALUE	43
	44
A.5.9 Process Creation and Management Fortran Bindings	45
MPI_ABORT(COMM, ERRORCODE, IERROR)	46
INTEGER COMM, ERRORCODE, IERROR	47
	48
MPI_CLOSE_PORT(PORT_NAME, IERROR)	49
CHARACTER*(*) PORT_NAME	50
INTEGER IERROR	51

```
1 MPI_COMM_ACCEPT(PORT_NAME, INFO, ROOT, COMM, NEWCOMM, IERROR)
2   CHARACTER*(*) PORT_NAME
3   INTEGER INFO, ROOT, COMM, NEWCOMM, IERROR
4
5 MPI_COMM_CONNECT(PORT_NAME, INFO, ROOT, COMM, NEWCOMM, IERROR)
6   CHARACTER*(*) PORT_NAME
7   INTEGER INFO, ROOT, COMM, NEWCOMM, IERROR
8
9 MPI_COMM_DISCONNECT(COMM, IERROR)
10  INTEGER COMM, IERROR
11
12 MPI_COMM_GET_PARENT(PARENT, IERROR)
13  INTEGER PARENT, IERROR
14
15 MPI_COMM_JOIN(FD, INTERCOMM, IERROR)
16  INTEGER FD, INTERCOMM, IERROR
17
18 MPI_COMM_SPAWN(COMMAND, ARGV, MAXPROCS, INFO, ROOT, COMM, INTERCOMM,
19               ARRAY_OF_ERRCODES, IERROR)
20  CHARACTER*(*) COMMAND, ARGV(*)
21  INTEGER MAXPROCS, INFO, ROOT, COMM, INTERCOMM, ARRAY_OF_ERRCODES(*), IERROR
22
23 MPI_COMM_SPAWN_MULTIPLE(COUNT, ARRAY_OF_COMMANDS, ARRAY_OF_ARGV,
24                          ARRAY_OF_MAXPROCS, ARRAY_OF_INFO, ROOT, COMM, INTERCOMM,
25                          ARRAY_OF_ERRCODES, IERROR)
26  INTEGER COUNT, ARRAY_OF_MAXPROCS(*), ARRAY_OF_INFO(*), ROOT, COMM,
27  INTERCOMM, ARRAY_OF_ERRCODES(*), IERROR
28  CHARACTER*(*) ARRAY_OF_COMMANDS(*), ARRAY_OF_ARGV(COUNT, *)
29
30 MPI_FINALIZE(IERROR)
31  INTEGER IERROR
32
33 MPI_FINALIZED(FLAG, IERROR)
34  LOGICAL FLAG
35  INTEGER IERROR
36
37 MPI_INIT(IERROR)
38  INTEGER IERROR
39
40 MPI_INIT_THREAD(REQUIRED, PROVIDED, IERROR)
41  INTEGER REQUIRED, PROVIDED, IERROR
42
43 MPI_INITIALIZED(FLAG, IERROR)
44  LOGICAL FLAG
45  INTEGER IERROR
46
47 MPI_IS_THREAD_MAIN(FLAG, IERROR)
48  LOGICAL FLAG
49  INTEGER IERROR
50
51 MPI_LOOKUP_NAME(SERVICE_NAME, INFO, PORT_NAME, IERROR)
52  CHARACTER*(*) SERVICE_NAME, PORT_NAME
53  INTEGER INFO, IERROR
54
55 MPI_OPEN_PORT(INFO, PORT_NAME, IERROR)
56  INTEGER INFO, IERROR
```

CHARACTER*(*) PORT_NAME	1
MPI_PUBLISH_NAME(SERVICE_NAME, INFO, PORT_NAME, IERROR)	2
CHARACTER*(*) SERVICE_NAME, PORT_NAME	3
INTEGER INFO, IERROR	4
MPI_QUERY_THREAD(PROVIDED, IERROR)	5
INTEGER PROVIDED, IERROR	6
MPI_SESSION_FINALIZE(SESSION, IERROR)	7
INTEGER SESSION, IERROR	8
MPI_SESSION_GET_INFO(SESSION, INFO_USED, IERROR)	9
INTEGER SESSION, INFO_USED, IERROR	10
MPI_SESSION_GET_NTH_PSET(SESSION, INFO, N, PSET_LEN, PSET_NAME, IERROR)	11
INTEGER SESSION, INFO, N, PSET_LEN, IERROR	12
CHARACTER*(*) PSET_NAME	13
MPI_SESSION_GET_NUM_PSETS(SESSION, INFO, NPSET_NAMES, IERROR)	14
INTEGER SESSION, INFO, NPSET_NAMES, IERROR	15
MPI_SESSION_GET_PSET_INFO(SESSION, PSET_NAME, INFO, IERROR)	16
INTEGER SESSION, INFO, IERROR	17
CHARACTER*(*) PSET_NAME	18
MPI_SESSION_INIT(INFO, ERRHANDLER, SESSION, IERROR)	19
INTEGER INFO, ERRHANDLER, SESSION, IERROR	20
MPI_UNPUBLISH_NAME(SERVICE_NAME, INFO, PORT_NAME, IERROR)	21
CHARACTER*(*) SERVICE_NAME, PORT_NAME	22
INTEGER INFO, IERROR	23
A.5.10 One-Sided Communications Fortran Bindings	24
MPI_ACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK,	25
TARGET_DISP, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, IERROR)	26
<type> ORIGIN_ADDR(*)	27
INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,	28
TARGET_DATATYPE, OP, WIN, IERROR	29
INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP	30
MPI_COMPARE_AND_SWAP(ORIGIN_ADDR, COMPARE_ADDR, RESULT_ADDR, DATATYPE,	31
TARGET_RANK, TARGET_DISP, WIN, IERROR)	32
<type> ORIGIN_ADDR(*), COMPARE_ADDR(*), RESULT_ADDR(*)	33
INTEGER DATATYPE, TARGET_RANK, WIN, IERROR	34
INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP	35
MPI_FETCH_AND_OP(ORIGIN_ADDR, RESULT_ADDR, DATATYPE, TARGET_RANK, TARGET_DISP,	36
OP, WIN, IERROR)	37
<type> ORIGIN_ADDR(*), RESULT_ADDR(*)	38
INTEGER DATATYPE, TARGET_RANK, OP, WIN, IERROR	39
INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP	40
MPI_FETCH_AND_OP(ORIGIN_ADDR, RESULT_ADDR, DATATYPE, TARGET_RANK, TARGET_DISP,	41
OP, WIN, IERROR)	42
<type> ORIGIN_ADDR(*), RESULT_ADDR(*)	43
INTEGER DATATYPE, TARGET_RANK, OP, WIN, IERROR	44
INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP	45
	46
	47
	48

```
1 MPI_GET(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
2         TARGET_COUNT, TARGET_DATATYPE, WIN, IERROR)
3     <type> ORIGIN_ADDR(*)
4     INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
5         TARGET_DATATYPE, WIN, IERROR
6     INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
7
8 MPI_GET_ACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_ADDR,
9         RESULT_COUNT, RESULT_DATATYPE, TARGET_RANK, TARGET_DISP,
10        TARGET_COUNT, TARGET_DATATYPE, OP, WIN, IERROR)
11    <type> ORIGIN_ADDR(*), RESULT_ADDR(*)
12    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_COUNT, RESULT_DATATYPE,
13        TARGET_RANK, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, IERROR
14    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
15
16 MPI_PUT(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
17        TARGET_COUNT, TARGET_DATATYPE, WIN, IERROR)
18    <type> ORIGIN_ADDR(*)
19    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
20        TARGET_DATATYPE, WIN, IERROR
21    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
22
23 MPI_RACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK,
24        TARGET_DISP, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, REQUEST,
25        IERROR)
26    <type> ORIGIN_ADDR(*)
27    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
28        TARGET_DATATYPE, OP, WIN, REQUEST, IERROR
29    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
30
31 MPI_RGET(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
32        TARGET_COUNT, TARGET_DATATYPE, WIN, REQUEST, IERROR)
33    <type> ORIGIN_ADDR(*)
34    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
35        TARGET_DATATYPE, WIN, REQUEST, IERROR
36    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
37
38 MPI_RGET_ACCUMULATE(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_ADDR,
39        RESULT_COUNT, RESULT_DATATYPE, TARGET_RANK, TARGET_DISP,
40        TARGET_COUNT, TARGET_DATATYPE, OP, WIN, REQUEST, IERROR)
41    <type> ORIGIN_ADDR(*), RESULT_ADDR(*)
42    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, RESULT_COUNT, RESULT_DATATYPE,
43        TARGET_RANK, TARGET_COUNT, TARGET_DATATYPE, OP, WIN, REQUEST,
44        IERROR
45    INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP
46
47 MPI_RPUT(ORIGIN_ADDR, ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_DISP,
48        TARGET_COUNT, TARGET_DATATYPE, WIN, REQUEST, IERROR)
49    <type> ORIGIN_ADDR(*)
50    INTEGER ORIGIN_COUNT, ORIGIN_DATATYPE, TARGET_RANK, TARGET_COUNT,
51        TARGET_DATATYPE, WIN, REQUEST, IERROR
```

```

INTEGER(KIND=MPI_ADDRESS_KIND) TARGET_DISP 1
MPI_WIN_ALLOCATE(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, WIN, IERROR) 2
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR 3
INTEGER DISP_UNIT, INFO, COMM, WIN, IERROR 4
If the Fortran compiler provides TYPE(C_PTR), then overloaded by: 5
INTERFACE MPI_WIN_ALLOCATE 6
  SUBROUTINE MPI_WIN_ALLOCATE(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, & 7
    WIN, IERROR) 8
    IMPORT :: MPI_ADDRESS_KIND 9
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR 10
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR 11
  END SUBROUTINE 12
  SUBROUTINE MPI_WIN_ALLOCATE_CPTR(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, & 13
    WIN, IERROR) 14
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR 15
    IMPORT :: MPI_ADDRESS_KIND 16
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR 17
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE 18
    TYPE(C_PTR) :: BASEPTR 19
  END SUBROUTINE 20
END INTERFACE 21
MPI_WIN_ALLOCATE_SHARED(SIZE, DISP_UNIT, INFO, COMM, BASEPTR, WIN, IERROR) 22
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR 23
INTEGER DISP_UNIT, INFO, COMM, WIN, IERROR 24
If the Fortran compiler provides TYPE(C_PTR), then overloaded by: 25
INTERFACE MPI_WIN_ALLOCATE_SHARED 26
  SUBROUTINE MPI_WIN_ALLOCATE_SHARED(SIZE, DISP_UNIT, INFO, COMM, & 27
    BASEPTR, WIN, IERROR) 28
    IMPORT :: MPI_ADDRESS_KIND 29
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR 30
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR 31
  END SUBROUTINE 32
  SUBROUTINE MPI_WIN_ALLOCATE_SHARED_CPTR(SIZE, DISP_UNIT, INFO, COMM, & 33
    BASEPTR, WIN, IERROR) 34
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR 35
    IMPORT :: MPI_ADDRESS_KIND 36
    INTEGER :: DISP_UNIT, INFO, COMM, WIN, IERROR 37
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE 38
    TYPE(C_PTR) :: BASEPTR 39
  END SUBROUTINE 40
END INTERFACE 41
MPI_WIN_ATTACH(WIN, BASE, SIZE, IERROR) 42
INTEGER WIN, IERROR 43
<type> BASE(*) 44
INTEGER(KIND=MPI_ADDRESS_KIND) SIZE 45

```

```
1 MPI_WIN_COMPLETE(WIN, IERROR)
2   INTEGER WIN, IERROR
3
4 MPI_WIN_CREATE(BASE, SIZE, DISP_UNIT, INFO, COMM, WIN, IERROR)
5   <type> BASE(*)
6   INTEGER(KIND=MPI_ADDRESS_KIND) SIZE
7   INTEGER DISP_UNIT, INFO, COMM, WIN, IERROR
8
9 MPI_WIN_CREATE_DYNAMIC(INFO, COMM, WIN, IERROR)
10  INTEGER INFO, COMM, WIN, IERROR
11
12 MPI_WIN_DETACH(WIN, BASE, IERROR)
13  INTEGER WIN, IERROR
14  <type> BASE(*)
15
16 MPI_WIN_FENCE(ASSERT, WIN, IERROR)
17  INTEGER ASSERT, WIN, IERROR
18
19 MPI_WIN_FLUSH(RANK, WIN, IERROR)
20  INTEGER RANK, WIN, IERROR
21
22 MPI_WIN_FLUSH_ALL(WIN, IERROR)
23  INTEGER WIN, IERROR
24
25 MPI_WIN_FLUSH_LOCAL(RANK, WIN, IERROR)
26  INTEGER RANK, WIN, IERROR
27
28 MPI_WIN_FLUSH_LOCAL_ALL(WIN, IERROR)
29  INTEGER WIN, IERROR
30
31 MPI_WIN_FREE(WIN, IERROR)
32  INTEGER WIN, IERROR
33
34 MPI_WIN_GET_GROUP(WIN, GROUP, IERROR)
35  INTEGER WIN, GROUP, IERROR
36
37 MPI_WIN_GET_INFO(WIN, INFO_USED, IERROR)
38  INTEGER WIN, INFO_USED, IERROR
39
40 MPI_WIN_LOCK(LOCK_TYPE, RANK, ASSERT, WIN, IERROR)
41  INTEGER LOCK_TYPE, RANK, ASSERT, WIN, IERROR
42
43 MPI_WIN_LOCK_ALL(ASSERT, WIN, IERROR)
44  INTEGER ASSERT, WIN, IERROR
45
46 MPI_WIN_POST(GROUP, ASSERT, WIN, IERROR)
47  INTEGER GROUP, ASSERT, WIN, IERROR
48
49 MPI_WIN_SET_INFO(WIN, INFO, IERROR)
50  INTEGER WIN, INFO, IERROR
51
52 MPI_WIN_SHARED_QUERY(WIN, RANK, SIZE, DISP_UNIT, BASEPTR, IERROR)
53  INTEGER WIN, RANK, DISP_UNIT, IERROR
54  INTEGER(KIND=MPI_ADDRESS_KIND) SIZE, BASEPTR
55
56 If the Fortran compiler provides TYPE(C_PTR), then overloaded by:
```

```

INTERFACE MPI_WIN_SHARED_QUERY                                1
  SUBROUTINE MPI_WIN_SHARED_QUERY(WIN, RANK, SIZE, DISP_UNIT, & 2
    BASEPTR, IERROR)                                        3
    IMPORT :: MPI_ADDRESS_KIND                              4
    INTEGER :: WIN, RANK, DISP_UNIT, IERROR                5
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE, BASEPTR        6
  END SUBROUTINE                                          7
  SUBROUTINE MPI_WIN_SHARED_QUERY_CPTR(WIN, RANK, SIZE, DISP_UNIT, & 8
    BASEPTR, IERROR)                                        9
    USE, INTRINSIC :: ISO_C_BINDING, ONLY : C_PTR          10
    IMPORT :: MPI_ADDRESS_KIND                              11
    INTEGER :: WIN, RANK, DISP_UNIT, IERROR                12
    INTEGER(KIND=MPI_ADDRESS_KIND) :: SIZE                 13
    TYPE(C_PTR) :: BASEPTR                                 14
  END SUBROUTINE                                          15
END INTERFACE                                            16

MPI_WIN_START(GROUP, ASSERT, WIN, IERROR)                 17
  INTEGER GROUP, ASSERT, WIN, IERROR                       18

MPI_WIN_SYNC(WIN, IERROR)                                  19
  INTEGER WIN, IERROR                                      20

MPI_WIN_TEST(WIN, FLAG, IERROR)                            21
  INTEGER WIN, IERROR                                      22
  LOGICAL FLAG                                            23

MPI_WIN_UNLOCK(RANK, WIN, IERROR)                          24
  INTEGER RANK, WIN, IERROR                                25

MPI_WIN_UNLOCK_ALL(WIN, IERROR)                            26
  INTEGER WIN, IERROR                                      27

MPI_WIN_WAIT(WIN, IERROR)                                   28
  INTEGER WIN, IERROR                                      29

A.5.11 External Interfaces Fortran Bindings                30

MPI_GREQUEST_COMPLETE(REQUEST, IERROR)                     31
  INTEGER REQUEST, IERROR                                  32

MPI_GREQUEST_START(QUERY_FN, FREE_FN, CANCEL_FN, EXTRA_STATE, REQUEST, IERROR) 33
  EXTERNAL QUERY_FN, FREE_FN, CANCEL_FN                   34
  INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE              35
  INTEGER REQUEST, IERROR                                  36

MPI_STATUS_SET_CANCELLED(STATUS, FLAG, IERROR)             37
  INTEGER STATUS(MPI_STATUS_SIZE), IERROR                 38
  LOGICAL FLAG                                             39

MPI_STATUS_SET_ELEMENTS(STATUS, DATATYPE, COUNT, IERROR)  40
  INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR 41

```

```
1 MPI_STATUS_SET_ERROR(STATUS, ERR, IERROR)
2     INTEGER STATUS(MPI_STATUS_SIZE), ERR, IERROR
3
4 MPI_STATUS_SET_SOURCE(STATUS, SOURCE, IERROR)
5     INTEGER STATUS(MPI_STATUS_SIZE), SOURCE, IERROR
6
7 MPI_STATUS_SET_TAG(STATUS, TAG, IERROR)
8     INTEGER STATUS(MPI_STATUS_SIZE), TAG, IERROR
9
```

10 A.5.12 I/O Fortran Bindings

```
11 MPI_CONVERSION_FN_NULL(USERBUF, DATATYPE, COUNT, FILEBUF, POSITION,
12     EXTRA_STATE, IERROR)
13     <TYPE> USERBUF(*), FILEBUF(*)
14     INTEGER DATATYPE, COUNT, IERROR
15     INTEGER(KIND=MPI_OFFSET_KIND) POSITION
16     INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
17
18 MPI_FILE_CLOSE(FH, IERROR)
19     INTEGER FH, IERROR
20
21 MPI_FILE_DELETE(FILENAME, INFO, IERROR)
22     CHARACTER*(*) FILENAME
23     INTEGER INFO, IERROR
24
25 MPI_FILE_GET_AMODE(FH, AMODE, IERROR)
26     INTEGER FH, AMODE, IERROR
27
28 MPI_FILE_GET_ATOMICITY(FH, FLAG, IERROR)
29     INTEGER FH, IERROR
30     LOGICAL FLAG
31
32 MPI_FILE_GET_BYTE_OFFSET(FH, OFFSET, DISP, IERROR)
33     INTEGER FH, IERROR
34     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET, DISP
35
36 MPI_FILE_GET_GROUP(FH, GROUP, IERROR)
37     INTEGER FH, GROUP, IERROR
38
39 MPI_FILE_GET_INFO(FH, INFO_USED, IERROR)
40     INTEGER FH, INFO_USED, IERROR
41
42 MPI_FILE_GET_POSITION(FH, OFFSET, IERROR)
43     INTEGER FH, IERROR
44     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
45
46 MPI_FILE_GET_POSITION_SHARED(FH, OFFSET, IERROR)
47     INTEGER FH, IERROR
48     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
49
50 MPI_FILE_GET_SIZE(FH, SIZE, IERROR)
51     INTEGER FH, IERROR
52     INTEGER(KIND=MPI_OFFSET_KIND) SIZE
53
54 MPI_FILE_GET_TYPE_EXTENT(FH, DATATYPE, EXTENT, IERROR)
```


INTEGER FH, DATATYPE, IERROR	1
INTEGER(KIND=MPI_ADDRESS_KIND) EXTENT	2
MPI_FILE_GET_VIEW(FH, DISP, ETYPE, FILETYPE, DATAREP, IERROR)	3
INTEGER FH, ETYPE, FILETYPE, IERROR	4
INTEGER(KIND=MPI_OFFSET_KIND) DISP	5
CHARACTER*(*) DATAREP	6
MPI_FILE_IREAD(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)	7
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	8
<type> BUF(*)	9
MPI_FILE_IREAD_ALL(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)	10
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	11
<type> BUF(*)	12
MPI_FILE_IREAD_AT(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)	13
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	14
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	15
<type> BUF(*)	16
MPI_FILE_IREAD_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)	17
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	18
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	19
<type> BUF(*)	20
MPI_FILE_IREAD_SHARED(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)	21
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	22
<type> BUF(*)	23
MPI_FILE_IWRITE(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)	24
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	25
<type> BUF(*)	26
MPI_FILE_IWRITE_ALL(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)	27
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	28
<type> BUF(*)	29
MPI_FILE_IWRITE_AT(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)	30
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	31
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	32
<type> BUF(*)	33
MPI_FILE_IWRITE_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, REQUEST, IERROR)	34
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	35
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	36
<type> BUF(*)	37
MPI_FILE_IWRITE_SHARED(FH, BUF, COUNT, DATATYPE, REQUEST, IERROR)	38
INTEGER FH, COUNT, DATATYPE, REQUEST, IERROR	39
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	40
<type> BUF(*)	41
MPI_FILE_OPEN(COMM, FILENAME, AMODE, INFO, FH, IERROR)	42
INTEGER COMM, AMODE, INFO, FH, IERROR	43
CHARACTER*(*) FILENAME	44
	45
	46
	47
	48

```
1 MPI_FILE_PREALLOCATE(FH, SIZE, IERROR)
2   INTEGER FH, IERROR
3   INTEGER(KIND=MPI_OFFSET_KIND) SIZE
4
5 MPI_FILE_READ(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
6   INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
7   <type> BUF(*)
8
9 MPI_FILE_READ_ALL(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
10  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
11  <type> BUF(*)
12
13 MPI_FILE_READ_ALL_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)
14  INTEGER FH, COUNT, DATATYPE, IERROR
15  <type> BUF(*)
16
17 MPI_FILE_READ_ALL_END(FH, BUF, STATUS, IERROR)
18  INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
19  <type> BUF(*)
20
21 MPI_FILE_READ_AT(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)
22  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
23  INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
24  <type> BUF(*)
25
26 MPI_FILE_READ_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)
27  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
28  INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
29  <type> BUF(*)
30
31 MPI_FILE_READ_AT_ALL_BEGIN(FH, OFFSET, BUF, COUNT, DATATYPE, IERROR)
32  INTEGER FH, COUNT, DATATYPE, IERROR
33  INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
34  <type> BUF(*)
35
36 MPI_FILE_READ_AT_ALL_END(FH, BUF, STATUS, IERROR)
37  INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
38  <type> BUF(*)
39
40 MPI_FILE_READ_ORDERED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
41  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
42  <type> BUF(*)
43
44 MPI_FILE_READ_ORDERED_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)
45  INTEGER FH, COUNT, DATATYPE, IERROR
46  <type> BUF(*)
47
48 MPI_FILE_READ_ORDERED_END(FH, BUF, STATUS, IERROR)
49  INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
50  <type> BUF(*)
51
52 MPI_FILE_READ_SHARED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
53  INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
54  <type> BUF(*)
```

MPI_FILE_SEEK(FH, OFFSET, WHENCE, IERROR)	1
INTEGER FH, WHENCE, IERROR	2
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	3
MPI_FILE_SEEK_SHARED(FH, OFFSET, WHENCE, IERROR)	4
INTEGER FH, WHENCE, IERROR	5
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	6
MPI_FILE_SET_ATOMICITY(FH, FLAG, IERROR)	7
INTEGER FH, IERROR	8
LOGICAL FLAG	9
MPI_FILE_SET_INFO(FH, INFO, IERROR)	10
INTEGER FH, INFO, IERROR	11
MPI_FILE_SET_SIZE(FH, SIZE, IERROR)	12
INTEGER FH, IERROR	13
INTEGER(KIND=MPI_OFFSET_KIND) SIZE	14
MPI_FILE_SET_VIEW(FH, DISP, ETYPE, FILETYPE, DATAREP, INFO, IERROR)	15
INTEGER FH, ETYPE, FILETYPE, INFO, IERROR	16
INTEGER(KIND=MPI_OFFSET_KIND) DISP	17
CHARACTER*(*) DATAREP	18
MPI_FILE_SYNC(FH, IERROR)	19
INTEGER FH, IERROR	20
MPI_FILE_WRITE(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)	21
INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR	22
<type> BUF(*)	23
MPI_FILE_WRITE_ALL(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)	24
INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR	25
<type> BUF(*)	26
MPI_FILE_WRITE_ALL_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)	27
INTEGER FH, COUNT, DATATYPE, IERROR	28
<type> BUF(*)	29
MPI_FILE_WRITE_ALL_END(FH, BUF, STATUS, IERROR)	30
INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR	31
<type> BUF(*)	32
MPI_FILE_WRITE_AT(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)	33
INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR	34
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	35
<type> BUF(*)	36
MPI_FILE_WRITE_AT_ALL(FH, OFFSET, BUF, COUNT, DATATYPE, STATUS, IERROR)	37
INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR	38
INTEGER(KIND=MPI_OFFSET_KIND) OFFSET	39
<type> BUF(*)	40
MPI_FILE_WRITE_AT_ALL_BEGIN(FH, OFFSET, BUF, COUNT, DATATYPE, IERROR)	41
INTEGER FH, COUNT, DATATYPE, IERROR	42
<type> BUF(*)	43
MPI_FILE_WRITE_AT_ALL_END(FH, OFFSET, BUF, STATUS, IERROR)	44
INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR	45
<type> BUF(*)	46
MPI_FILE_WRITE_AT_ALL_BEGIN(FH, OFFSET, BUF, COUNT, DATATYPE, IERROR)	47
INTEGER FH, COUNT, DATATYPE, IERROR	48

```

1     INTEGER(KIND=MPI_OFFSET_KIND) OFFSET
2     <type> BUF(*)
3
4     MPI_FILE_WRITE_AT_ALL_END(FH, BUF, STATUS, IERROR)
5     INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
6     <type> BUF(*)
7
8     MPI_FILE_WRITE_ORDERED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
9     INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
10    <type> BUF(*)
11
12    MPI_FILE_WRITE_ORDERED_BEGIN(FH, BUF, COUNT, DATATYPE, IERROR)
13    INTEGER FH, COUNT, DATATYPE, IERROR
14    <type> BUF(*)
15
16    MPI_FILE_WRITE_ORDERED_END(FH, BUF, STATUS, IERROR)
17    INTEGER FH, STATUS(MPI_STATUS_SIZE), IERROR
18    <type> BUF(*)
19
20    MPI_FILE_WRITE_SHARED(FH, BUF, COUNT, DATATYPE, STATUS, IERROR)
21    INTEGER FH, COUNT, DATATYPE, STATUS(MPI_STATUS_SIZE), IERROR
22    <type> BUF(*)
23
24    MPI_REGISTER_DATAREP(DATAREP, READ_CONVERSION_FN, WRITE_CONVERSION_FN,
25    DTYPE_FILE_EXTENT_FN, EXTRA_STATE, IERROR)
26    CHARACTER*(*) DATAREP
27    EXTERNAL READ_CONVERSION_FN, WRITE_CONVERSION_FN, DTYPE_FILE_EXTENT_FN
28    INTEGER(KIND=MPI_ADDRESS_KIND) EXTRA_STATE
29    INTEGER IERROR

```

A.5.13 Language Bindings Fortran Bindings

```

28    MPI_F_SYNC_REG(BUF)
29    <type> BUF(*)

```

The following procedure is not available with mpif.h:

```

30    MPI_STATUS_F082F(F08_STATUS, F_STATUS, IERROR)
31    TYPE(MPI_Status) :: F08_STATUS
32    INTEGER :: F_STATUS(MPI_STATUS_SIZE), IERROR

```

The following procedure is not available with mpif.h:

```

33    MPI_STATUS_F2F08(F_STATUS, F08_STATUS, IERROR)
34    INTEGER :: F_STATUS(MPI_STATUS_SIZE), IERROR
35    TYPE(MPI_Status) :: F08_STATUS

```

```

36    MPI_TYPE_CREATE_F90_COMPLEX(P, R, NEWTYPE, IERROR)
37    INTEGER P, R, NEWTYPE, IERROR

```

```

38    MPI_TYPE_CREATE_F90_INTEGER(R, NEWTYPE, IERROR)
39    INTEGER R, NEWTYPE, IERROR

```

```

40    MPI_TYPE_CREATE_F90_REAL(P, R, NEWTYPE, IERROR)
41    INTEGER P, R, NEWTYPE, IERROR

```

```

42    MPI_TYPE_MATCH_SIZE(TYPECLASS, SIZE, DATATYPE, IERROR)
43    INTEGER TYPECLASS, SIZE, DATATYPE, IERROR

```

A.5.14 Tools / Profiling Interface Fortran Bindings

MPI_PCONTROL(LEVEL)
 INTEGER LEVEL

A.5.15 Deprecated Fortran Bindings

MPI_ATTR_DELETE(COMM, KEYVAL, IERROR)
 INTEGER COMM, KEYVAL, IERROR

MPI_ATTR_GET(COMM, KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)
 INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR
 LOGICAL FLAG

MPI_ATTR_PUT(COMM, KEYVAL, ATTRIBUTE_VAL, IERROR)
 INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR

MPI_DUP_FN(OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT,
 FLAG, IERR)
 INTEGER OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT,
 IERR
 LOGICAL FLAG

MPI_GET_ELEMENTS_X(STATUS, DATATYPE, COUNT, IERROR)
 INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, IERROR
 INTEGER(KIND=MPI_COUNT_KIND) COUNT

MPI_INFO_GET(INFO, KEY, VALUELEN, VALUE, FLAG, IERROR)
 INTEGER INFO, VALUELEN, IERROR
 CHARACTER*(*) KEY, VALUE
 LOGICAL FLAG

MPI_INFO_GET_VALUELEN(INFO, KEY, VALUELEN, FLAG, IERROR)
 INTEGER INFO, VALUELEN, IERROR
 CHARACTER*(*) KEY
 LOGICAL FLAG

MPI_KEYVAL_CREATE(COPY_FN, DELETE_FN, KEYVAL, EXTRA_STATE, IERROR)
 EXTERNAL COPY_FN, DELETE_FN
 INTEGER KEYVAL, EXTRA_STATE, IERROR

MPI_KEYVAL_FREE(KEYVAL, IERROR)
 INTEGER KEYVAL, IERROR

MPI_NULL_COPY_FN(OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
 ATTRIBUTE_VAL_OUT, FLAG, IERR)
 INTEGER OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN, ATTRIBUTE_VAL_OUT,
 IERR
 LOGICAL FLAG

MPI_NULL_DELETE_FN(COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERROR)
 INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERROR

MPI_SIZEOF(X, SIZE, IERROR)

```
1      <type> X
2      INTEGER SIZE, IERROR
3
4      MPI_STATUS_SET_ELEMENTS_X(STATUS, DATATYPE, COUNT, IERROR)
5      INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, IERROR
6      INTEGER(KIND=MPI_COUNT_KIND) COUNT
7
8      MPI_TYPE_GET_EXTENT_X(DATATYPE, LB, EXTENT, IERROR)
9      INTEGER DATATYPE, IERROR
10     INTEGER(KIND=MPI_COUNT_KIND) LB, EXTENT
11
12     MPI_TYPE_GET_TRUE_EXTENT_X(DATATYPE, TRUE_LB, TRUE_EXTENT, IERROR)
13     INTEGER DATATYPE, IERROR
14     INTEGER(KIND=MPI_COUNT_KIND) TRUE_LB, TRUE_EXTENT
15
16     MPI_TYPE_SIZE_X(DATATYPE, SIZE, IERROR)
17     INTEGER DATATYPE, IERROR
18     INTEGER(KIND=MPI_COUNT_KIND) SIZE
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
```

Appendix B

Change-Log

Annex B.1 summarizes changes from the previous version of the MPI standard to the version presented by this document. Only significant changes (i.e., clarifications and new features) that might either require implementation effort in the MPI libraries or change the understanding of MPI from a user's perspective are presented. Editorial modifications, formatting, typo corrections and minor clarifications are not shown. If not otherwise noted, the section and page references refer to the locations of the change or new functionality in this version of the standard. Changes in Annexes B.2–B.6 were already introduced in the corresponding sections in previous versions of this standard.

B.1 Changes from Version 4.0 to Version 4.1

B.1.1 Fixes to Errata in Previous Versions of MPI

1. Sections 2.4.1, 3.5, 3.7, 3.8, 6.12, 7.1.2, 7.2.2, 7.4.2, 11.6.2, 14.2.5, 14.6.5, 19.1.17, 19.1.20, on pages 11, 55, 69, 94, 250, 305, 308, 321, 516, 644, 705, 822, 832 and MPI-4.0 Sections 2.4.1, 3.5, 3.7, 3.8, 6.12, 7.1.2, 7.2.2, 7.4.2, 11.6.2, 14.2.5, 14.6.5, 19.1.17, 19.1.20 on pages 13, 54, 60, 84, 250, 312, 314, 327, 518, 650, 713, 826, and 834.
The term *pending* communication or I/O operation is defined as *in the active operation state*. If the phrase *pending communication operation* in MPI-1.1 to MPI-4.0 additionally includes *decoupled MPI activities*, then this has been added explicitly. If this phrase had a different meaning, it was replaced accordingly, see item 4 in this list and item 15 in Section B.1.2.
2. Sections 2.5.4 and A.1.1 on pages 19 and 849, and MPI-4.0 Sections 2.5.4 and A.1.1 on pages 20 and 857.
The implementation of named MPI *constants* in C and Fortran and implied usage restrictions were clarified.
3. Section 2.5.4 on page 19, and MPI-4.0 Section 2.5.4 on page 20.
Add MPI_MAX_PSET_NAME_LEN and MPI_MAX_STRINGTAG_LEN to list of named constants.
4. Sections 3.9, 11.2.2, 11.3.1 and 11.10.4 on pages 104, 488, 494 and 542 and MPI-4.0 Sections 3.9, 11.2.2, 11.3.1 and 11.10.4 on pages 94, 494, 501 and 546.
The requirements for calling MPI_FINALIZE, MPI_SESSION_FINALIZE, and MPI_COMM_DISCONNECT and the outcome of MPI_COMM_DISCONNECT, especially for related inactive persistent request handles, were clarified.
5. Section 4.3 on page 121, and MPI-4.0 Section 4.3.3. on page 115. Example 4.4 on page 125, and MPI-4.0 Example 4.4 on page 115.

1 Fixed and simplified erroneous MPI-4.0 Example 4.4. The example could deadlock
2 due to incorrect use of the flag variable in multiple MPI test procedure calls or thread
3 concurrent access. The example was also simplified by removing unnecessary code
4 and updated according to current best practice in OpenMP.

- 5
- 6 6. Section 5.1.9 on page 156 and MPI-4.0 Section 5.1.9 on page 150.
7 The relationship between MPI_TYPE_COMMIT and initialization/finalization of a
8 session (or the World Model) with MPI_INIT, MPI_INIT_THREAD, MPI_FINALIZE,
9 MPI_SESSION_INIT, and MPI_SESSION_FINALIZE was clarified.
- 10
- 11 7. Sections 7.4.2 on page 321 and 7.6.2 on page 353, and MPI-4.0 Section 7.4.2 on page 343
12 and Section 7.6.2 on page 360.
13 Use of the errhandler argument to MPI_COMM_CREATE_FROM_GROUP and
14 MPI_INTERCOMM_CREATE_FROM_GROUPS is clarified. The error handler invoked
15 when an error is encountered during invocation of these functions is also clarified.
- 16
- 17 8. Section 8.5.5 on page 397 and MPI-4.0 Section 8.5.5 on page 403.
18 The unintended change in the specification of argument coords in
19 MPI_CART_COORDS in MPI-4.0 is reverted to the original meaning in MPI-1.1 to
20 MPI-3.1. It is clarified that the outcome of MPI_CART_GET and MPI_CART_COORDS
21 is unspecified for the case that maxdims is less than ndims.
- 22
- 23 9. Section 9.3 on page 449 and MPI-4.0 Section 9.3 on page 459.
24 The fallback error-handler for the Sessions Model was clarified.
- 25
- 26 10. Section 9.5 on page 464 and MPI-4.0 Section 9.5 on page 473.
27 It was clarified that MPI_LASTUSEDPCODE is only available in the World Model.
- 28
- 29 11. Sections 11.1 and 11.7 on pages 481 and 518, and MPI-4.0 Sections 11.1 and 11.7 on
30 pages 487 and 521.
31 It was clarified that the usage of the Dynamic Process Model requires the World Model
32 to be initialized.
- 33
- 34 12. Section 12.5.2 on page 593 and MPI-4 Section 12.5.2 on page 598.
35 The definition of MPI_WIN_TEST was clarified.
- 36
- 37 13. Section 12.5.4 on page 602, MPI-3 Section 11.5.4 on page 449, MPI-3.1 Section 11.5.4
38 on page 448, and MPI-4 Section 12.5.4 on page 605.
39 The description of MPI_WIN_SYNC was clarified to include its use for ordering load-
40 /store accesses to shared memory. A statement was added to highlight that a call
41 to MPI_WIN_SYNC does not complete pending RMA operations and that a call to
42 MPI_WIN_SYNC does not guarantee any progress of MPI operations.
- 43
- 44 14. Section 7.4.2 on page 322, MPI-3 Section 6.4.2 on page 237, MPI-3.1 Section 6.4.2 on
45 page 237, and MPI-4 Section 7.4.2 on page 327.
46 The description of MPI_COMM_DUP now clarifies that error handlers are also copied
47 in the output communicator produced when this procedure is called.
- 48
- 49 15. Section 13.3 on page 633, and MPI-4 Section 13.3 on page 640.
50 Large count interface of MPI_STATUS_SET_ELEMENTS had been missing and was
51 added.

16. Sections 15.3.6 and 15.3.7 on pages 729 and 736.
The intent of handle arguments of the language independent definition of MPI_T_CVAR_WRITE, MPI_T_PVAR_HANDLE_ALLOC, MPI_T_PVAR_HANDLE_FREE, MPI_T_PVAR_START, MPI_T_PVAR_STOP, MPI_T_PVAR_WRITE, and MPI_T_PVAR_RESET were marked as INOUT in accordance with the special rule for handles described in Section 2.3.
17. Section 4.3.1 on page 122, MPI-4 Section 4.3.1 on page 112.
Example 4.2 was corrected.

B.1.2 Changes in MPI-4.1

1. Section 1.13 on page 6.
Introduced the concept of side documents.
2. Sections 2.4.1 and 2.4.2 on pages 11 and 14.
Added the definition of the *enabled operation state*.
3. Sections 2.5.5, 2.6.2, and 19.1 on pages 20, 23, and 787.
The MPI standard now reflects that TS 29113 was superceded by Fortran 2018.
4. Sections 2.6.1, 5.1.5, 5.1.7, 5.1.8, 5.1.11, 13.3, and 16.4 on pages 22, 148, 153, 155, 158, 633, and 777.
MPI_TYPE_SIZE_X, MPI_TYPE_GET_EXTENT_X, MPI_TYPE_GET_TRUE_EXTENT_X, MPI_GET_ELEMENTS_X, and MPI_STATUS_SET_ELEMENTS_X were deprecated and may be removed in a future version of the MPI specification.
5. Sections 2.6.1, 7.2.4 and 9.1.2 on pages 22, 309 and 442.
MPI_HOST has been deprecated, and a mention to host process has been removed.
6. Sections 2.6.1, 19.1.1, 19.1.4, 16.4 on pages 22, 787, 793, 777.
Deprecated the use of `mpif.h`.
7. Section 2.6.4 on page 25.
Removed the functions MPI_WTIME, PMPI_WTIME, MPI_WTICK, and PMPI_WTICK from the list of functions that may be implemented as a macro.
8. Section 2.6.4 on page 25 and Section 19.3.4 on page 836.
Removed the ability to implement MPI handle conversion functions as a macro.
9. Sections 2.9, 3.7.6, 12.5.2, and Example 12.13 on pages 27, 90, 593, and 617.
The *progress rules* were clarified in general and for MPI_REQUEST_GET_STATUS and MPI_WIN_TEST. The terms *strong* and *weak progress* were introduced. An example showing restrictions on the use of MPI shared memory for synchronizing purposes was introduced.
10. Sections 3.2.5 and 13.3 on pages 39 and 633.
Added procedures MPI_STATUS_GET_SOURCE, MPI_STATUS_GET_TAG, and MPI_STATUS_GET_ERROR to query MPI status fields and procedures MPI_STATUS_SET_SOURCE, MPI_STATUS_SET_TAG, and

1 MPI_STATUS_SET_ERROR to set MPI status fields. Direct access to these fields
2 remains available.

3
4 11. Section 3.6 on page 58.

5 Automatic (unlimited) buffering is added, which can be enabled by using
6 MPI_BUFFER_AUTOMATIC in any of the buffer attach procedures. New procedures
7 MPI_COMM_ATTACH_BUFFER, MPI_SESSION_ATTACH_BUFFER,
8 MPI_COMM_DETACH_BUFFER and MPI_SESSION_DETACH_BUFFER have been
9 added. The buffers attached with the existing functions MPI_BUFFER_ATTACH and
10 MPI_BUFFER_DETACH now only apply to communicators that have no buffer at-
11 tached at the communicator or session level. New procedures
12 MPI_COMM_FLUSH_BUFFER, MPI_SESSION_FLUSH_BUFFER, and
13 MPI_BUFFER_FLUSH were added as a combination of detach and attach, as well as
14 the corresponding nonblocking variants MPI_COMM_IFLUSH_BUFFER,
15 MPI_SESSION_IFLUSH_BUFFER, and MPI_BUFFER_IFLUSH.

16
17 12. Subsection 3.7.6 on page 90

18 Added new procedures MPI_REQUEST_GET_STATUS_ANY,
19 MPI_REQUEST_GET_STATUS_SOME, and MPI_REQUEST_GET_STATUS_ALL to
20 query the statuses of multiple requests without freeing them.

21
22 13. Section 6.9.4 and 5.1.13 on pages 229 and 162.

23 Added procedure MPI_TYPE_GET_VALUE_INDEX to query predefined datatype han-
24 dles for pairs of value and index types to be usable in conjunction with MPI_MINLOC
25 and MPI_MAXLOC. Added combiner MPI_COMBINER_VALUE_INDEX for unnamed type
26 handles returned by MPI_TYPE_GET_VALUE_INDEX.

27
28 14. Section 7.4.2 on page 321.

29 MPI_COMM_TYPE_RESOURCE_GUIDED was added as a new possible value for the
30 split_type parameter of the MPI_COMM_SPLIT_TYPE procedure, as well as a new
31 info key "mpi_pset_name".

32
33 15. Section 7.4.3 on page 339.

34 The definition of MPI_COMM_FREE was clarified.

35
36 16. Section 7.4.4 on page 341.

37 A new info key was added, namely "mpi_assert_strict_persistent_collective_ordering".

38
39 17. Sections 7.4.4, 11.2.1, 11.3, 11.4.3, 11.8.4, 12.2.1, and 14.2.8 on pages 340, 482, 493,
40 508, 527, 548, and 646.

41 Added the ability to request support for, query support of, and assert usage of memory
42 allocation kinds via two new info keys, "mpi_memory_alloc_kinds" and
43 "mpi_assert_memory_alloc_kinds".

44
45 18. Section 7.8 on page 376 was amended to allow MPI_COMM_NULL,
46 MPI_DATATYPE_NULL, and MPI_WIN_NULL to be passed to MPI_COMM_GET_NAME,
47 MPI_TYPE_GET_NAME and MPI_WIN_GET_NAME, respectively.

48
49 19. Section 9.1.2 on page 445.

Added new procedure MPI_GET_HW_RESOURCE_INFO.

20. Section 9.4 on page 460. 1
Added new error class MPI_ERR_ERRHANDLER. 2
21. Section 9.5 on page 464. 3
Add procedures MPI_REMOVE_ERROR_CLASS, MPI_REMOVE_ERROR_CODE, 4
MPI_REMOVE_ERROR_STRING to complement the procedures adding error classes/- 5
codes/strings. 6
22. Section 12.2 on page 548. 7
Relaxed the constraints on the windows for which shared memory can be queried using 8
MPI_WIN_SHARED_QUERY to allow windows with flavor MPI_WIN_FLAVOR_CREATE 9
and MPI_WIN_FLAVOR_ALLOCATE. 10
23. Section 12.2.1 on page 548. 11
Added new info key "mpi_accumulate_granularity" to specify a desired synchronization 12
granularity of accumulate operations. 13
24. Section 12.2.5 on page 562. 14
Implementations may avoid synchronization of processes in MPI_WIN_FREE if the 15
"no_locks" info key is set to "true". 16
25. Section 15.3.2 on page 724. 17
The text specifying when entities of the MPI Tool Information Interface can be bound 18
to objects during the object's lifetime was clarified. 19
26. Section 15.3.8 on page 760 and Section 15.3.9 on page 767. 20
Behavior specified when the count of dropped events or category changes overflow, 21
respectively. 22
27. Annex A.2 on page 875. 23
The annex has been completed with the operation-related MPI procedures for one- 24
sided communication and some other rarely used scenarios. It is now integrated into 25
the MPI standard. 26
28. Section 12.5.2 on page 593. 27
In Example 12.4, MPI_PUT has been removed from the list of procedures that may 28
delay their return waiting for the call to MPI_WIN_POST to occur at the target. 29
MPI RMA communication procedures are generally not intended to delay their return 30
waiting for synchronization procedure calls to occur at the target. 31
29. Section 12.7 on page 609 and Section 12.8 on page 622. 32
Clarified the use of MPI_WIN_SYNC for memory synchronization on shared memory. 33

B.2 Changes from Version 3.1 to Version 4.0 34

B.2.1 Fixes to Errata in Previous Versions of MPI 35

1. Sections 8.6.1, 8.6.2 and 8.9 on pages 411, 415 and 436, and MPI-3.1 Sections 7.6.1, 36
7.6.2 and 7.8 on pages 315, 318 and 329. 37
MPI_NEIGHBOR_ALLTOALL_{|V|W} and MPI_NEIGHBOR_ALLGATHER_{|V} for Car- 38
tesian virtual grids were clarified. An advice to implementors was added to illustrate 39

1 a correct implementation for the case of `periods[d]=1` or `.TRUE.` and `dims[d]=1` or 2
2 in a direction `d`.

- 3
4 2. Section 19.3.5 on page 838, and MPI-3.1 Section 17.2.5 on page 657 line 11.
5 Clarified that the `MPI_STATUS_F2F08` and `MPI_STATUS_F082F` routines and the
6 declaration for `TYPE(MPI_Status)` are not supposed to appear with `mpif.h`.
- 7
8 3. Sections 2.5.4, 19.3.5, and A.1.1 on pages 20, 838, and 852, and MPI-3.1 Sections
9 2.5.4, 17.2.5, and A.1.1 on pages 15, 656, and 669.
10 Define the C constants `MPI_F_STATUS_SIZE`, `MPI_F_SOURCE`, `MPI_F_TAG`, and
11 `MPI_F_ERROR`.
- 12
13 4. Section 19.3.5 on page 839, and MPI-3.1 Section 17.2.5 on page 658.
14 Added missing `const` to `IN` parameters for `MPI_STATUS_F2F08` and
15 `MPI_STATUS_F082F`.

16 B.2.2 Changes in MPI-4.0

- 17
18 1. Sections 2.2, 18.2.2, and 19.1.5 on pages 9, 786, and 794.
19 The limit for the maximum length of MPI identifiers was removed.
- 20
21 2. Section 2.4, 3.4, 3.7.2, 3.7.3, 3.8.1, 3.8.2, 6.13, 14.4.5, and Annex A.2 on pages 11, 50,
22 71, 78, 94, 97, 273, 681, and 875.
23 The semantic terms were updated.
- 24
25 3. Sections 2.5.8 and 19.2 on pages 21 and 833, and throughout the entire document.
26 New large count functions `MPI_{...}_c` in C and through function overloading in
27 the Fortran `mpi_f08` module, (with the exception of the explicit Fortran procedures
28 `MPI_Op_create_c` and `MPI_Register_datarep_c`) and the new large count callbacks
29 `MPI_User_function_c` and `MPI_Datarep_conversion_function_c` together with the prede-
30 fined function `MPI_CONVERSION_FN_NULL_C` were introduced to accomodate large
31 buffers and/or datatypes.
32 Clarifications were added to the behavior of `INOUT/OUT` parameters that cannot
33 represent the value to be returned for the `MPI_BUFFER_DETACH` and
34 `MPI_FILE_GET_TYPE_EXTENT` functions.
35 A new error class `MPI_ERR_VALUE_TOO_LARGE` was introduced.
- 36
37 4. Sections 2.8, 9.3, 9.5, and 11.2.1 on pages 26, 449, 464, and 482.
38 MPI calls that are not related to any objects are considered to be attached to the
39 communicator `MPI_COMM_SELF` instead of `MPI_COMM_WORLD`. The definition of
40 `MPI_ERRORS_ARE_FATAL` was clarified to cover all connected processes, and a new
41 error handler, `MPI_ERRORS_ABORT`, was created to limit the scope of aborting.
- 42
43 5. Section 3.7 on page 69.
44 The introduction of MPI nonblocking communication was changed to describe cor-
45 rectness and performance reasons for the use of nonblocking communication.
- 46
47 6. Section 3.7.2 on page 71.
48 Addition of `MPI_ISENDRECV` and `MPI_ISENDRECV_REPLACE`.

7. Sections [3.7.3](#), [3.9](#), [6.13](#), [8.8](#), and [8.9](#) on pages [78](#), [104](#), [273](#), [429](#), and [436](#).
Persistent collective communication `MPI_{ALLGATHER|...}_INIT` including persistent collective neighborhood communication `MPI_NEIGHBOR_{ALLGATHER|...}_INIT` was added to the standard.
8. Sections [3.8.4](#) and [16.3](#) on pages [102](#) and [774](#).
Cancelling a send request by calling `MPI_CANCEL` has been deprecated and may be removed in a future version of the MPI specification.
9. Chapter [4](#) on page [113](#).
A new chapter on partitioned communication with the new MPI procedures `MPI_{PARRIVED|PREADY{...}}` and `MPI_{PRECV|PSEND}_INIT` was added.
10. Section [7.4.2](#) on page [321](#).
`MPI_COMM_TYPE_HW_UNGUIDED` was added as a new possible value for the `split_type` parameter of the `MPI_COMM_SPLIT_TYPE` function.
11. Section [7.4.2](#) on page [321](#).
`MPI_COMM_TYPE_HW_GUIDED` was added as a new possible value for the `split_type` parameter of the `MPI_COMM_SPLIT_TYPE` function, as well as a new info key `"mpi_hw_resource_type"`. A specific value associated with this new info key is also defined: `"mpi_shared_memory"`.
12. Section [7.4.2](#) on page [321](#).
The functions `MPI_COMM_DUP` and `MPI_COMM_IDUP` were updated to no longer propagate info hints.
This change may affect backward compatibility.
13. Section [7.4.2](#) on page [321](#).
The `MPI_COMM_IDUP_WITH_INFO` function was added.
14. Sections [7.4.4](#), [12.2.7](#), and [14.2.8](#) on pages [340](#), [564](#), and [646](#).
The definition of info hints was updated to allow applications to provide assertions regarding their usage of MPI objects and operations.
15. Section [7.4.4](#) on page [340](#).
The new info hints `"mpi_assert_no_any_tag"`, `"mpi_assert_no_any_source"`, `"mpi_assert_exact_length"`, and `"mpi_assert_allow_overtaking"` were added for use with communicators.
16. Sections [7.4.4](#), [12.2.7](#), and [14.2.8](#) on pages [340](#), [564](#), and [646](#).
The semantics of the `MPI_COMM_SET_INFO`, `MPI_COMM_GET_INFO`, `MPI_WIN_SET_INFO`, `MPI_WIN_GET_INFO`, `MPI_FILE_SET_INFO`, and `MPI_FILE_GET_INFO` were clarified.
17. Section [8.5](#) on page [386](#).
`MPI_DIMS_CREATE` is now guaranteed to return `MPI_SUCCESS` if the number of dimensions passed to the routine is set to 0 and the number of nodes is set to 1.
18. Sections [9.2](#), [12.2.2](#), and [12.2.3](#) on pages [446](#), [552](#), and [554](#).
Introduced alignment requirements for memory allocated through `MPI_ALLOC_MEM`, `MPI_WIN_ALLOCATE`, and `MPI_WIN_ALLOCATE_SHARED` and added a new info key `"mpi_minimum_memory_alignment"` to specify a desired alternative minimum alignment.

- 1 19. Sections [9.3](#) and [9.4](#) on pages [449](#) and [460](#).
2 Clarified definition of errors to say that MPI should continue whenever possible and
3 allow the user to recover from errors.
4
- 5 20. Section [9.4](#) on page [460](#).
6 Added text to clarify what is implied about the status of MPI and user visible buffers
7 when MPI functions return MPI_SUCCESS or other error codes.
8
- 9 21. Section [9.4](#) on page [462](#).
10 The error class MPI_ERR_PROC_ABORTED has been added.
11
- 12 22. Section [10](#) on page [473](#).
13 Added a new function MPI_INFO_GET_STRING that takes a buffer length argument
14 for returning info value strings. This function returns the required buffer length for
15 the requested string and guarantees null termination for C strings where buffer size
16 is greater than 0.
17
- 18 23. Section [10](#) on page [473](#) and Section [16.3](#) on page [774](#).
19 MPI_INFO_GET and MPI_INFO_GET_VALUELEN were deprecated.
20
- 21 24. Chapter [11](#), [3.2.3](#), [7.2.4](#), [7.3.2](#), [7.4.2](#), [7.6.2](#), [9.1.1](#), [9.1.2](#), [9.3](#), [9.3.4](#), [9.5](#), [11.6](#), [14.2.1](#),
22 [14.2.7](#), [14.7](#), [15.3.4](#), [19.3.4](#), [19.3.6](#), and Annex [A](#) on pages [481](#), [35](#), [309](#), [312](#), [321](#), [353](#),
23 [441](#), [442](#), [449](#), [457](#), [464](#), [514](#), [639](#), [645](#), [711](#), [726](#), [836](#), [840](#), and [849](#).
24 The Sessions Model was added to the standard. New MPI procedures are
25 MPI_SESSION_{INIT|FINALIZE}, MPI_SESSION_GET_{...},
26 MPI_SESSION_{...}_ERRHANDLER, MPI_GROUP_FROM_SESSION_PSET,
27 MPI_COMM_CREATE_FROM_GROUP,
28 MPI_INTERCOMM_CREATE_FROM_GROUPS, and new conversion functions are
29 MPI_SESSION_{C2F|F2C}. New declarations are MPI_Session in C and
30 TYPE(MPI_Session) together with the related overloaded operators .EQ., .NE., == and
31 /= in the Fortran mpi_f08 and mpi modules, and the callback function prototype
32 MPI_Session_errhandler_function. New constants are MPI_SESSION_NULL,
33 MPI_ERR_SESSION, MPI_MAX_PSET_NAME_LEN, MPI_MAX_STRINGTAG_LEN,
34 MPI_T_BIND_MPI_SESSION and the predefined info key "mpi_size".
35
- 36 25. Section [11.2.1](#) on page [482](#).
37 A new function MPI_INFO_CREATE_ENV was added.
38
- 39 26. Sections [11.2.1](#) and [11.10.4](#) on pages [482](#) and [542](#).
40 Clarified the semantic of failure and error reporting before (and during) MPI_INIT
41 and after MPI_FINALIZE.
42
- 43 27. Section [11.8.4](#) on page [527](#).
44 Added the "mpi_initial_errhandler" reserved info key with the reserved values
45 "mpi_errors_abort", "mpi_errors_are_fatal", and "mpi_errors_return" to the launch keys in
46 MPI_COMM_SPAWN, MPI_COMM_SPAWN_MULTIPLE, and mpiexec.
47
- 48 28. Section [12.5.3](#) on page [597](#).
49 RMA passive target synchronization using locks can now be used portably in memory
50 allocated via MPI_WIN_ALLOCATE_SHARED.

29. Section 13.3 on page 633. 1
 The `mpi_f08` binding incorrectly had the dummy parameter `flag` in the MPI F08 binding for `MPI_STATUS_SET_CANCELLED` marked as `INTENT(OUT)`. It has been fixed to be `INTENT(IN)`. 2
30. Sections 15.3 and 15.3.8 on pages 722 and 748. 3
 A callback-driven event interface with the `MPI_T_{SOURCE|EVENT}_{...}` and `MPI_T_CATEGORY_{GET|GET_NUM}_EVENTS` routines, the declaration types `MPI_T_cb_safety`, `MPI_T_event_{instance|registration}`, `MPI_T_source_order`, and the callback function prototypes `MPI_T_event_{cb|dropped_cb|free_cb}_function`, were added to the MPI tool information interface. 4
31. Section 15.3.9 on page 767. 5
 The argument `stamp` (previously described as a virtual time stamp) from `MPI_T_CATEGORY_CHANGED` was renamed to `update_number` and its intended implementation and use was clarified. 6
32. Section 15.3.10, Table 15.7, and Section 16.3 on pages 768, 769, and 774. 7
`MPI_T_ERR_INVALID_ITEM` is deprecated. MPI routines should return `MPI_T_ERR_INVALID_INDEX` instead of `MPI_T_ERR_INVALID_ITEM`. 8
33. Section 16.3 on page 777. 9
`MPI_SIZEOF` was deprecated. 10
34. Section 19.1.5 on page 794. 11
 An exception was added for the specific Fortran names in the case of TS 29113 interface specifications in `mpif.h` for `MPI_NEIGHBOR_ALLTOALLW_INIT`, `MPI_NEIGHBOR_ALLTOALLV_INIT`, and `MPI_NEIGHBOR_ALLGATHERV_INIT`. 12

B.3 Changes from Version 3.0 to Version 3.1 13

B.3.1 Fixes to Errata in Previous Versions of MPI 14

1. Chapters 3–19, Annex A.4 on page 921, and Example 6.22 on page 240, and MPI-3.0 Chapters 3–17, Annex A.3 on page 707, and Example 5.21 on page 187. 15
 Within the `mpi_f08` Fortran support method, `BIND(C)` was removed from all `SUBROUTINE`, `FUNCTION`, and `ABSTRACT INTERFACE` definitions. 16
2. Section 3.2.5 on page 39, and MPI-3.0 Section 3.2.5 on page 30. 17
 The three public fields `MPI_SOURCE`, `MPI_TAG`, and `MPI_ERROR` of the Fortran derived type `TYPE(MPI_Status)` must be of type `INTEGER`. 18
3. Section 3.8.2 on page 97, and MPI-3.0 Section 3.8.2 on page 67. 19
 The flag arguments of the Fortran interfaces of `MPI_IMPROBE` were originally incorrectly defined as `INTEGER` (instead as `LOGICAL`). 20
4. Section 7.4.2 on page 321, and MPI-3.0 Section 6.4.2 on page 237. 21
 In the `mpi_f08` binding of `MPI_COMM_IDUP`, the output argument `newcomm` is declared as `ASYNCHRONOUS`. 22

- 1 5. Section 7.4.4 on page 340, and MPI-3.0 Section 6.4.4 on page 248.
2 In the `mpi_f08` binding of `MPI_COMM_SET_INFO`, the `INTENT` of `comm` is `IN`, and the
3 optional output argument `ierror` was missing.
4
- 5 6. Section 8.6 on page 410, and MPI-3.0 Sections 7.6, on pages 314.
6 In the case of virtual general graph topologies (created with `MPI_CART_CREATE`), the
7 use of neighborhood collective communication is restricted to adjacency matrices with
8 the number of edges between any two processes is defined to be the same for both
9 processes (i.e., with a symmetric adjacency matrix).
10
- 11 7. Section 9.1.1 on page 441, and MPI-3.0 Section 8.1.1 on page 335.
12 In the `mpi_f08` binding of `MPI_GET_LIBRARY_VERSION`, a typo in the
13 `resultlen` argument was corrected.
- 14 8. Sections 9.2 (`MPI_ALLOC_MEM` and `MPI_ALLOC_MEM_CPTR`),
15 12.2.2 (`MPI_WIN_ALLOCATE` and `MPI_WIN_ALLOCATE_CPTR`),
16 12.2.3 (`MPI_WIN_ALLOCATE_SHARED` and `MPI_WIN_ALLOCATE_SHARED_CPTR`),
17 12.2.3 (`MPI_WIN_SHARED_QUERY` and `MPI_WIN_SHARED_QUERY_CPTR`),
18 15.2.1 and 15.2.6 (Profiling interface), and corresponding sections in MPI-3.0.
19 The linker name concept was substituted by defining specific procedure names.
20
- 21 9. Section 12.2.1 on page 548, and MPI-3.0 Section 11.2.2 on page 407.
22 The "same_size" info key can be used with all window flavors, and requires that all
23 processes in the process group of the communicator have provided this info key with
24 the same value.
25
- 26 10. Section 12.3.4 on page 573, and MPI-3.0 Section 11.3.4 on page 424.
27 Origin buffer arguments to `MPI_GET_ACCUMULATE` are ignored when the
28 `MPI_NO_OP` operation is used.
- 29 11. Section 12.3.4 on page 573, and MPI-3.0 Section 11.3.4 on page 424.
30 Clarify the roles of origin, result, and target communication parameters in
31 `MPI_GET_ACCUMULATE`.
32
- 33 12. Section 15.3 on page 722, and MPI-3.0 Section 14.3 on page 561
34 New paragraph and advice to users clarifying intent of variable names in the tools
35 information interface.
36
- 37 13. Section 15.3.3 on page 725, and MPI-3.0 Section 14.3.3 on page 563.
38 New paragraph clarifying variable name equivalence in the tools information interface.
- 39 14. Sections 15.3.6, 15.3.7, and 15.3.9 on pages 729, 736, and 764, and
40 MPI-3.0 Sections 14.3.6, 14.3.7, and 14.3.8 on pages 567, 573, and 584.
41 In functions `MPI_T_CVAR_GET_INFO`, `MPI_T_PVAR_GET_INFO`, and
42 `MPI_T_CATEGORY_GET_INFO`, clarification of parameters that must be identical for
43 equivalent control variable / performance variable / category names across connected
44 processes.
45
- 46 15. Section 15.3.7 on page 736, and MPI-3.0 Section 14.3.7 on page 573.
47 Clarify return code of `MPI_T_PVAR_{START,STOP,RESET}` routines.
48

16. Section 15.3.7 on page 736, and MPI-3.0 Section 14.3.7 on page 579, line 7. Clarify the return code when bad handle is passed to an MPI_T_PVAR_* routine.
17. Section 19.1.4 on page 793, and MPI-3.0 Section 17.1.4 on page 603. The advice to implementors at the end of the section was rewritten and moved into the following section.
18. Section 19.1.5 on page 794, and MPI-3.0 Section 17.1.5 on page 605. The section was fully rewritten. The linker name concept was substituted by defining specific procedure names.
19. Section 19.1.6 on page 799, and MPI-3.0 Section 17.1.6 on page 611. The requirements on BIND(C) procedure interfaces were removed.
20. Annexes A.3, A.4, and A.5 on pages 884, 921, and 1009, and MPI-3.0 Annexes A.2, A.3, and A.4 on pages 685, 707, and 756. The predefined callback MPI_CONVERSION_FN_NULL was added to all three annexes.
21. Annex A.4.5 on page 963, and MPI-3.0 Annex A.3.4 on page 724. In the mpi_f08 binding of MPI_{COMM|TYPE|WIN}_{DUP|NULL_COPY|NULL_DELETE}_FN, all INTENT(...) information was removed.

B.3.2 Changes in MPI-3.1

1. Sections 2.6.4 and 5.1.5 on pages 25 and 148. The use of the intrinsic operators “+” and “-” for absolute addresses is substituted by MPI_AINT_ADD and MPI_AINT_DIFF. In C, they can be implemented as macros.
2. Sections 9.1.1, 11.2.1, and 11.6 on pages 441, 482, and 514. The routines MPI_INITIALIZED, MPI_FINALIZED, MPI_QUERY_THREAD, MPI_IS_THREAD_MAIN, MPI_GET_VERSION, and MPI_GET_LIBRARY_VERSION are callable from threads without restriction (in the sense of MPI_THREAD_MULTIPLE), irrespective of the actual level of thread support provided, in the case where the implementation supports threads.
3. Section 12.2.1 on page 548. The “same_disp_unit” info key was added for use in RMA window creation routines.
4. Sections 14.4.2 and 14.4.3 on pages 656 and 663. Added MPI_FILE_IREAD_AT_ALL, MPI_FILE_IWRITE_AT_ALL, MPI_FILE_IREAD_ALL, and MPI_FILE_IWRITE_ALL
5. Sections 15.3.6, 15.3.7, and 15.3.9 on pages 729, 736, and 764. Clarified that NULL parameters can be provided in MPI_T_{CVAR|PVAR|CATEGORY}_GET_INFO routines.
6. Sections 15.3.6, 15.3.7, 15.3.9, and 15.3.10 on pages 729, 736, 764, and 768. New routines MPI_T_CVAR_GET_INDEX, MPI_T_PVAR_GET_INDEX, MPI_T_CATEGORY_GET_INDEX, were added to support retrieving indices of variables and categories. The error codes MPI_T_ERR_INVALID and MPI_T_ERR_INVALID_NAME were added to indicate invalid uses of the interface.

B.4 Changes from Version 2.2 to Version 3.0

B.4.1 Fixes to Errata in Previous Versions of MPI

1. Sections [2.6.2](#) and [2.6.3](#) on pages [23](#) and [24](#), and MPI-2.2 Section 2.6.2 on page 17, lines 41–42, Section 2.6.3 on page 18, lines 15–16, and Section 2.6.4 on page 18, lines 40–41.

This is an MPI-2 erratum: The scope for the reserved prefix `MPI_` and the C++ namespace `MPI` is now any name as originally intended in MPI-1.

2. Sections [3.2.2](#), [6.9.2](#), [14.5.2](#) Table [14.2](#), and Annex [A.1.1](#) on pages [33](#), [227](#), [695](#), and [849](#), and MPI-2.2 Sections 3.2.2, 5.9.2, 13.5.2 Table 13.2, 16.1.16 Table 16.1, and Annex A.1.1 on pages 27, 164, 433, 472 and 513

This is an MPI-2.2 erratum: New named predefined datatypes `MPI_CXX_BOOL`, `MPI_CXX_FLOAT_COMPLEX`, `MPI_CXX_DOUBLE_COMPLEX`, and `MPI_CXX_LONG_DOUBLE_COMPLEX` were added in C and Fortran corresponding to the C++ types `bool`, `std::complex<float>`, `std::complex<double>`, and `std::complex<long double>`. These datatypes also correspond to the deprecated C++ predefined datatypes `MPI::BOOL`, `MPI::COMPLEX`, `MPI::DOUBLE_COMPLEX`, and `MPI::LONG_DOUBLE_COMPLEX`, which were removed in MPI-3.0. The nonstandard C++ types `Complex<...>` were substituted by the standard types `std::complex<...>`.

3. Sections [6.9.2](#) on pages [227](#) and MPI-2.2 Section 5.9.2, page 165, line 47.

This is an MPI-2.2 erratum: `MPI_C_COMPLEX` was added to the “Complex” reduction group.

4. Section [8.5.5](#) on page [397](#), and MPI-2.2, Section 7.5.5 on page 257, C++ interface on page 264, line 3.

This is an MPI-2.2 erratum: The argument `rank` was removed and `in/outdegree` are now defined as `int& indegree` and `int& outdegree` in the C++ interface of `MPI_DIST_GRAPH_NEIGHBORS_COUNT`.

5. Section [14.5.2](#), Table [14.2](#) on page [695](#), and MPI-2.2, Section 13.5.3, Table 13.2 on page 433.

This was an MPI-2.2 erratum: The `MPI_C_BOOL` “external32” representation is corrected to a 1-byte size.

6. MPI-2.2 Section 16.1.16 on page 471, line 45.

This is an MPI-2.2 erratum: The constant `MPI::_LONG_LONG` should be `MPI::LONG_LONG`.

7. Annex [A.1.1](#) on page [849](#), Table “Optional datatypes (Fortran),” and MPI-2.2, Annex A.1.1, Table on page 517, lines 34, and 37–41.

This is an MPI-2.2 erratum: The C++ datatype handles `MPI::INTEGER16`, `MPI::REAL16`, `MPI::F_COMPLEX4`, `MPI::F_COMPLEX8`, `MPI::F_COMPLEX16`, `MPI::F_COMPLEX32` were added to the table.

B.4.2 Changes in MPI-3.0

1. Section [2.6.1](#) on page [22](#), Section [17.2](#) on page [784](#) and all other chapters.

The C++ bindings were removed from the standard. See errata in Section [B.4.1](#) on

- page 1056 for the latest changes to the MPI C++ binding defined in MPI-2.2. This change may affect backward compatibility.
2. Section 2.6.1 on page 22, Section 16.1 on page 771 and Section 17.1 on page 783. The deprecated functions `MPI_TYPE_HVECTOR`, `MPI_TYPE_HINDEXED`, `MPI_TYPE_STRUCT`, `MPI_ADDRESS`, `MPI_TYPE_EXTENT`, `MPI_TYPE_LB`, `MPI_TYPE_UB`, `MPI_ERRHANDLER_CREATE` (and its callback function prototype `MPI_Handler_function`), `MPI_ERRHANDLER_SET`, `MPI_ERRHANDLER_GET`, the deprecated special datatype handles `MPI_LB`, `MPI_UB`, and the constants `MPI_COMBINER_HINDEXED_INTEGER`, `MPI_COMBINER_HVECTOR_INTEGER`, `MPI_COMBINER_STRUCT_INTEGER` were removed from the standard. This change may affect backward compatibility.
 3. Section 2.3 on page 10. Clarified parameter usage for IN parameters. C bindings are now const-correct where backward compatibility is preserved.
 4. Section 2.5.4 on page 19 and Section 8.5.4 on page 390. The recommended C implementation value for `MPI_UNWEIGHTED` changed from `NULL` to non-`NULL`. An additional weight array constant (`MPI_WEIGHTS_EMPTY`) was introduced.
 5. Section 2.5.4 on page 19 and Section 9.1.1 on page 441. Added the new routine `MPI_GET_LIBRARY_VERSION` to query library specific versions, and the new constant `MPI_MAX_LIBRARY_VERSION_STRING`.
 6. Sections 2.5.8, 3.2.2, 3.3, 6.9.2, on pages 21, 33, 35, 227, Sections 5.1, 5.1.7, 5.1.8, 5.1.11, 16.4 on pages 127, 153, 155, 158, 780, and Annex A.1.1 on page 849. New inquiry functions, `MPI_TYPE_SIZE_X`, `MPI_TYPE_GET_EXTENT_X`, `MPI_TYPE_GET_TRUE_EXTENT_X`, and `MPI_GET_ELEMENTS_X`, return their results as an `MPI_Count` value, which is a new type large enough to represent element counts in memory, file views, etc. A new function, `MPI_STATUS_SET_ELEMENTS_X`, modifies the opaque part of an `MPI_Status` object so that a call to `MPI_GET_ELEMENTS_X` returns the provided `MPI_Count` value (in Fortran, `INTEGER(KIND=MPI_COUNT_KIND)`). The corresponding predefined datatype is `MPI_COUNT`.
 7. Chapter 3 on page 31 through Chapter 19 on page 787. In the C language bindings, the array-arguments' interfaces were modified to consistently use `[]` instead of `*`. Exceptions are `MPI_INIT`, which continues to use `char ***argv` (correct because of subtle rules regarding the use of the `&` operator with `char *argv[]`), and `MPI_INIT_THREAD`, which is changed to be consistent with `MPI_INIT`.
 8. Sections 3.2.5, 5.1.5, 5.1.11, 5.2 on pages 39, 148, 158, 178. The functions `MPI_GET_COUNT` and `MPI_GET_ELEMENTS` were defined to set the count argument to `MPI_UNDEFINED` when that argument would overflow. The functions `MPI_PACK_SIZE` and `MPI_TYPE_SIZE` were defined to set the size argument

1 to MPI_UNDEFINED when that argument would overflow. In all other MPI-2.2 rou-
2 tines, the type and semantics of the count arguments remain unchanged, i.e., int or
3 INTEGER.

- 4
- 5 9. Section 3.2.6 on page 42, and Section 3.8 on page 94.
6 MPI_STATUS_IGNORE can also be used in MPI_IProbe, MPI_Probe, MPI_Improbe,
7 and MPI_Mprobe.
- 8
- 9 10. Section 3.8 on page 94 and Section 3.10 on page 111.
10 The use of MPI_PROC_NULL in probe operations was clarified. A special predefined
11 message MPI_MESSAGE_NO_PROC was defined for the use of matching probe (i.e., the
12 new MPI_Mprobe and MPI_Improbe) with MPI_PROC_NULL.
- 13 11. Sections 3.8.2, 3.8.3, 19.3.4, A.1.1 on pages 97, 100, 836, 849.
14 Like MPI_Probe and MPI_IProbe, the new MPI_Mprobe and
15 MPI_Improbe operations allow incoming messages to be queried without actually
16 receiving them, except that MPI_Mprobe and MPI_Improbe provide a mechanism
17 to receive the specific message with the new routines MPI_Mrecv and
18 MPI_Irecv regardless of other intervening probe or receive operations. The opaque
19 object MPI_Message, the null handle MPI_MESSAGE_NULL, and the conversion functions
20 MPI_Message_c2f and MPI_Message_f2c were defined.
- 21
- 22 12. Section 5.1.2 on page 129 and Section 5.1.13 on page 162.
23 The routine MPI_Type_create_hindexed_block and constant
24 MPI_COMBINER_HINDEXED_BLOCK were added.
- 25
- 26 13. Chapter 6 on page 189 and Section 6.12 on page 250.
27 Added nonblocking interfaces to all collective operations.
- 28 14. Sections 7.4.2, 7.4.4, 12.2.7, on pages 321, 340, 564.
29 The new routines MPI_Comm_dup_with_info, MPI_Comm_set_info,
30 MPI_Comm_get_info, MPI_Win_set_info, and MPI_Win_get_info were
31 added. The routine MPI_Comm_dup must also duplicate info hints.
- 32
- 33 15. Section 7.4.2 on page 321.
34 Added MPI_Comm_idup.
- 35
- 36 16. Section 7.4.2 on page 321.
37 Added the new communicator construction routine MPI_Comm_create_group,
38 which is invoked only by the processes in the group of the new communicator being
39 constructed.
- 40 17. Section 7.4.2 on page 321.
41 Added the MPI_Comm_split_type routine and the communicator split type con-
42 stant MPI_COMM_TYPE_SHARED.
- 43
- 44 18. Section 7.6.2 on page 353.
45 In MPI-2.2, communication involved in an MPI_Intercomm_create operation
46 could interfere with point-to-point communication on the parent communicator with
47 the same tag or MPI_ANY_TAG. This interference has been removed in MPI-3.0.
- 48

19. Section 7.8 on page 376. 1
 Section 6.8 on page 238. The constant `MPI_MAX_OBJECT_NAME` also applies for type 2
 and window names. 3
20. Section 8.5.8 on page 408. 4
`MPI_CART_MAP` can also be used for a zero-dimensional topologies. 5
21. Section 8.6 on page 410 and Section 8.7 on page 422. 6
 The following neighborhood collective communication routines were added to sup- 7
 port sparse communication on virtual topology grids: `MPI_NEIGHBOR_ALLGATHER`, 8
`MPI_NEIGHBOR_ALLGATHERV`, `MPI_NEIGHBOR_ALLTOALL`, 9
`MPI_NEIGHBOR_ALLTOALLV`, `MPI_NEIGHBOR_ALLTOALLW` and the nonblocking 10
 variants `MPI_INEIGHBOR_ALLGATHER`, `MPI_INEIGHBOR_ALLGATHERV`, 11
`MPI_INEIGHBOR_ALLTOALL`, `MPI_INEIGHBOR_ALLTOALLV`, and 12
`MPI_INEIGHBOR_ALLTOALLW`. The displacement arguments in 13
`MPI_NEIGHBOR_ALLTOALLW` and `MPI_INEIGHBOR_ALLTOALLW` were defined as 14
 address size integers. In `MPI_DIST_GRAPH_NEIGHBORS`, an ordering rule was added 15
 for communicators created with `MPI_DIST_GRAPH_CREATE_ADJACENT`. 16
22. Section 11.2.1 on page 482 and Section 11.2.1 on page 485. 17
 The use of `MPI_INIT`, `MPI_INIT_THREAD` and `MPI_FINALIZE` was clarified. After 18
`MPI` is initialized, the application can access information about the execution envi- 19
 ronment by querying the new predefined info object `MPI_INFO_ENV`. 20
23. Section 11.2.1 on page 482. 21
 Allow calls to `MPI_T` routines before `MPI_INIT` and after `MPI_FINALIZE`. 22
24. Chapter 12 on page 547. 23
 Substantial revision of the entire One-sided chapter, with new routines for window 24
 creation, additional synchronization methods in passive target communication, new 25
 one-sided communication routines, a new memory model, and other changes. 26
25. Section 15.3 on page 722. 27
 A new MPI Tool Information Interface was added. 28
 The following changes are related to the Fortran language support. 29
26. Section 2.3 on page 10, and Sections 19.1.1, 19.1.2, 19.1.7 on pages 787, 788, and 803. 30
 The new `mpi_f08` Fortran module was introduced. 31
27. Section 2.5.1 on page 17, and Sections 19.1.2, 19.1.3, 19.1.7 on pages 788, 791, and 803. 32
 Handles to opaque objects were defined as named types within the `mpi_f08` Fortran 33
 module. The operators `.EQ.`, `.NE.`, `==`, and `/=` were overloaded to allow the comparison 34
 of these handles. The handle types and the overloaded operators are also available 35
 through the `mpi` Fortran module. 36
28. Sections 2.5.4, 2.5.5 on pages 19, 20, Sections 19.1.1, 19.1.10, 19.1.11, 19.1.12, 19.1.13 37
 on pages 787, 813, 814, 815, 818, and Sections 19.1.2, 19.1.3, 19.1.7 on pages 788, 791, 38
 803. 39
 Within the `mpi_f08` Fortran module, choice buffers were defined as assumed-type 40
 and assumed-rank according to Fortran 2008 with TS 29113 [47], and the compile- 41
 time constant `MPI_SUBARRAYS_SUPPORTED` was set to `.TRUE.`. With this, Fortran 42

1 subscript triplets can be used in nonblocking MPI operations; vector subscripts are not
2 supported in nonblocking operations. If the compiler does not support this Fortran
3 TS 29113 feature, the constant is set to `.FALSE.`.

- 4
- 5 29. Section 2.6.2 on page 23, Section 19.1.2 on page 788, and Section 19.1.7 on page 803.
6 The `ieror` dummy arguments are `OPTIONAL` within the `mpi_f08` Fortran module.
- 7
- 8 30. Section 3.2.5 on page 39, Sections 19.1.2, 19.1.3, 19.1.7, on pages 788, 791, 803, and
9 Section 19.3.5 on page 838.
10 Within the `mpi_f08` Fortran module, the status was defined as
11 `TYPE(MPI_Status)`. Additionally, within both the `mpi` and the `mpi_f08` modules, the
12 constants `MPI_STATUS_SIZE`, `MPI_SOURCE`, `MPI_TAG`, `MPI_ERROR`, and
13 `TYPE(MPI_Status)` are defined. New conversion routines were added:
14 `MPI_STATUS_F2F08`, `MPI_STATUS_F082F`, `MPI_Status_c2f08`, and
15 `MPI_Status_f082c`. In `mpi.h`, the new type `MPI_F08_status`, and the external variables
16 `MPI_F08_STATUS_IGNORE` and `MPI_F08_STATUSES_IGNORE` were added.
- 17
- 18 31. Section 3.6 on page 58.
19 In Fortran with the `mpi` module or `mpif.h`, the type of the `buffer_addr` argument of
20 `MPI_BUFFER_DETACH` is incorrectly defined and the argument is therefore unused.
- 21
- 22 32. Section 5.1 on page 127, Section 5.1.6 on page 151, and Section 19.1.15 on page 819.
23 The Fortran alignments of basic datatypes within Fortran derived types are imple-
24 mentation dependent; therefore it is recommended to use the `BIND(C)` attribute for
25 derived types in MPI communication buffers. If an array of structures (in C/C++)
26 or derived types (in Fortran) is to be used in MPI communication buffers, it is rec-
27 ommended that the user creates a portable datatype handle and additionally applies
28 `MPI_TYPE_CREATE_RESIZED` to this datatype handle.
- 29
- 30 33. Sections 5.1.10, 6.9.5, 6.9.7, 7.7.4, 7.8, 9.3.1, 9.3.2, 9.3.3, 16.1, 19.1.9 on pages 158,
31 235, 242, 370, 376, 452, 454, 456, 771, and 805. In some routines, the dummy ar-
32 gument names were changed because they were identical to the Fortran keywords
33 `TYPE` and `FUNCTION`. The new dummy argument names must be used because the
34 `mpi` and `mpi_f08` modules guarantee keyword-based actual argument lists. The ar-
35 gument name type was changed in `MPI_TYPE_DUP`, the Fortran `USER_FUNCTION` of
36 `MPI_OP_CREATE`, `MPI_TYPE_SET_ATTR`, `MPI_TYPE_GET_ATTR`,
37 `MPI_TYPE_DELETE_ATTR`, `MPI_TYPE_SET_NAME`, `MPI_TYPE_GET_NAME`,
38 `MPI_TYPE_MATCH_SIZE`, the callback prototype definition
39 `MPI_Type_delete_attr_function`, and the predefined callback function
40 `MPI_TYPE_NULL_DELETE_FN`; function was changed in `MPI_OP_CREATE`,
41 `MPI_COMM_CREATE_ERRHANDLER`, `MPI_WIN_CREATE_ERRHANDLER`,
42 `MPI_FILE_CREATE_ERRHANDLER`, and `MPI_ERRHANDLER_CREATE`. For consis-
43 tency reasons, `INOUBUF` was changed to `INOUTBUF` in `MPI_REDUCE_LOCAL`, and
44 `intracomm` to `newintracomm` in `MPI_INTERCOMM_MERGE`.
- 45
- 46 34. Section 7.7.2 on page 361.
47 It was clarified that in Fortran, the flag values returned by a
48 `comm_copy_attr_fn` callback, including `MPI_COMM_NULL_COPY_FN` and
`MPI_COMM_DUP_FN`, are `.FALSE.` and `.TRUE.`; see `MPI_COMM_CREATE_KEYVAL`.

35. Section 9.2 on page 446. 1
 With the `mpi` and `mpi_f08` Fortran modules, `MPI_ALLOC_MEM` now also supports `TYPE(C_PTR)` C-pointers instead of only returning an address-sized integer that may be usable together with a nonstandard Cray-pointer. 2
36. Section 19.1.15 on page 819, and Section 19.1.7 on page 803. 3
 Fortran `SEQUENCE` and `BIND(C)` derived application types can now be used as buffers in MPI operations. 4
37. Section 19.1.16 on page 821 to Section 19.1.19 on page 831, Section 19.1.7 on page 803, and Section 19.1.8 on page 804. 5
 The sections about Fortran optimization problems and their solutions were partially rewritten and new methods are added, e.g., the use of the `ASYNCHRONOUS` attribute. The constant `MPI_ASYNC_PROTECTS_NONBLOCKING` tells whether the semantics of the `ASYNCHRONOUS` attribute is extended to protect nonblocking operations. The Fortran routine `MPI_F_SYNC_REG` is added. MPI-3.0 compliance for an MPI library together with a Fortran compiler is defined in Section 19.1.7. 6
38. Section 19.1.2 on page 788. 7
 Within the `mpi_f08` Fortran module, dummy arguments are now declared with `INTENT=IN`, `OUT`, or `INOUT` as defined in the `mpi_f08` interfaces. 8
39. Section 19.1.3 on page 791, and Section 19.1.7 on page 803. 9
 The existing `mpi` Fortran module must implement compile-time argument checking. 10
40. Section 19.1.4 on page 793. 11
 The use of the `mpif.h` Fortran include file is now strongly discouraged. 12
41. Section A.1.1, Table **Predefined functions** on page 858, Section A.1.3 on page 866, and Section A.4.5 on page 963. 13
 Within the new `mpi_f08` module, all callback prototype definitions are now defined with explicit interfaces `PROCEDURE(MPI_...)` that have the `BIND(C)` attribute; user-written callbacks must be modified if the `mpi_f08` module is used. 14
42. Section A.1.3 on page 866. 15
 In some routines, the Fortran callback prototype names were changed from `..._FN` to `..._FUNCTION` to be consistent with the other language bindings. 16

B.5 Changes from Version 2.1 to Version 2.2 17

1. Section 2.5.4 on page 19. 18
 It is now guaranteed that predefined named constant handles (as other constants) can be used in initialization expressions or assignments, i.e., also before the call to `MPI_INIT`. 19
2. Section 2.6 on page 22, and Section 17.2 on page 784. 20
 The C++ language bindings have been deprecated and may be removed in a future version of the MPI specification. 21
3. Section 3.2.2 on page 33. 22
`MPI_CHAR` for printable characters is now defined for C type `char` (instead of signed 23

1 char). This change should not have any impact on applications nor on MPI libraries
2 (except some comment lines), because printable characters could and can be stored in
3 any of the C types char, signed char, and unsigned char, and MPI_CHAR is not allowed
4 for predefined reduction operations.

5
6 4. Section 3.2.2 on page 33.

7 MPI_(U)INT{8,16,32,64}_T, MPI_AINT, MPI_OFFSET, MPI_C_BOOL,
8 MPI_C_COMPLEX, MPI_C_FLOAT_COMPLEX, MPI_C_DOUBLE_COMPLEX, and
9 MPI_C_LONG_DOUBLE_COMPLEX are now valid predefined MPI datatypes.

10
11 5. Section 3.4 on page 50, Section 3.7.2 on page 71, Section 3.9 on page 104, and Sec-
12 tion 6.1 on page 189.

13 The read access restriction on the send buffer for blocking, non blocking and collec-
14 tive API has been lifted. It is permitted to access for read the send buffer while the
15 operation is in progress.

16 6. Section 3.7 on page 69.

17 The Advice to users for IBSEND and IRSEND was slightly changed.

18
19 7. Section 3.7.3 on page 78.

20 The advice to free an active request was removed in the Advice to users for
21 MPI_REQUEST_FREE.

22
23 8. Section 3.7.6 on page 90.

24 MPI_REQUEST_GET_STATUS changed to permit inactive or null requests as input.

25
26 9. Section 6.8 on page 218.

27 “In place” option is added to MPI_ALLTOALL, MPI_ALLTOALLV, and
28 MPI_ALLTOALLW for intra-communicators.

29
30 10. Section 6.9.2 on page 227.

31 Predefined parameterized datatypes (e.g., returned by
32 MPI_TYPE_CREATE_F90_REAL) and optional named predefined datatypes (e.g.
33 MPI_REAL8) have been added to the list of valid datatypes in reduction operations.

34
35 11. Section 6.9.2 on page 227.

36 MPI_(U)INT{8,16,32,64}_T are all considered C integer types for the purposes of the
37 predefined reduction operators. MPI_AINT and MPI_OFFSET are considered Fortran
38 integer types. MPI_C_BOOL is considered a Logical type.
39 MPI_C_COMPLEX, MPI_C_FLOAT_COMPLEX, MPI_C_DOUBLE_COMPLEX, and
40 MPI_C_LONG_DOUBLE_COMPLEX are considered Complex types.

41
42 12. Section 6.9.7 on page 242.

43 The local routines MPI_REDUCE_LOCAL and MPI_OP_COMMUTATIVE have been
44 added.

45
46 13. Section 6.10.1 on page 243.

47 The collective function MPI_REDUCE_SCATTER_BLOCK is added to the MPI stan-
48 dard.

49
50 14. Section 6.11.2 on page 248.

51 Added in place argument to MPI_EXSCAN.

15. Section [7.4.2](#) on page [321](#), and Section [7.6](#) on page [349](#).
Implementations that did not implement `MPI_COMM_CREATE` on inter-communicators will need to add that functionality. As the standard described the behavior of this operation on inter-communicators, it is believed that most implementations already provide this functionality. Note also that the C++ binding for both `MPI_COMM_CREATE` and `MPI_COMM_SPLIT` explicitly allow Intercomms.
16. Section [7.4.2](#) on page [321](#).
`MPI_COMM_CREATE` is extended to allow several disjoint subgroups as input if `comm` is an intra-communicator. If `comm` is an inter-communicator it was clarified that all processes in the same local group of `comm` must specify the same value for `group`.
17. Section [8.5.4](#) on page [390](#).
New functions for a scalable distributed graph topology interface has been added. In this section, the functions `MPI_DIST_GRAPH_CREATE_ADJACENT` and `MPI_DIST_GRAPH_CREATE`, the constants `MPI_UNWEIGHTED`, and the derived C++ class `Distgraphcomm` were added.
18. Section [8.5.5](#) on page [397](#).
For the scalable distributed graph topology interface, the functions `MPI_DIST_GRAPH_NEIGHBORS_COUNT` and `MPI_DIST_GRAPH_NEIGHBORS` and the constant `MPI_DIST_GRAPH` were added.
19. Section [8.5.5](#) on page [397](#).
Remove ambiguity regarding duplicated neighbors with `MPI_GRAPH_NEIGHBORS` and `MPI_GRAPH_NEIGHBORS_COUNT`.
20. Section [9.1.1](#) on page [441](#).
The subversion number changed from 1 to 2.
21. Section [9.3](#) on page [449](#), Section [16.2](#) on page [774](#), and Annex [A.1.3](#) on page [866](#).
Changed function pointer typedef names `MPI_{Comm,File,Win}_errhandler_fn` to `MPI_{Comm,File,Win}_errhandler_function`. Deprecated old “_fn” names.
22. Section [11.2.4](#) on page [492](#).
Attribute deletion callbacks on `MPI_COMM_SELF` are now called in LIFO order. Implementors must now also register all implementation-internal attribute deletion callbacks on `MPI_COMM_SELF` before returning from `MPI_INIT/MPI_INIT_THREAD`.
23. Section [12.3.4](#) on page [573](#).
The restriction added in MPI 2.1 that the operation `MPI_REPLACE` in `MPI_ACCUMULATE` can be used only with predefined datatypes has been removed. `MPI_REPLACE` can now be used even with derived datatypes, as it was in MPI 2.0. Also, a clarification has been made that `MPI_REPLACE` can be used only in `MPI_ACCUMULATE`, not in collective operations that do reductions, such as `MPI_REDUCE` and others.
24. Section [13.2](#) on page [627](#).
Add “*” to the `query_fn`, `free_fn`, and `cancel_fn` arguments to the C++ binding for `MPI::Grequest::Start()` for consistency with the rest of MPI functions that take function pointer arguments.

- 1 25. Section [14.5.2](#) on page [693](#), and Table [14.2](#) on page [695](#).
2 MPI_(U)INT{8,16,32,64}_T, MPI_AINT, MPI_OFFSET, MPI_C_COMPLEX,
3 MPI_C_FLOAT_COMPLEX, MPI_C_DOUBLE_COMPLEX,
4 MPI_C_LONG_DOUBLE_COMPLEX, and MPI_C_BOOL are added as predefined datatypes
5 in the "external32" representation.
6
- 7 26. Section [19.3.7](#) on page [843](#).
8 The description was modified that it only describes how an MPI implementation be-
9 haves, but not how MPI stores attributes internally. The erroneous MPI-2.1 Example
10 16.17 was replaced with three new examples [19.25](#), [19.26](#), and [19.27](#) on pages [843–845](#)
11 explicitly detailing cross-language attribute behavior. Implementations that matched
12 the behavior of the old example will need to be updated.
- 13 27. Annex [A.1.1](#) on page [849](#).
14 Removed type MPI::Fint (compare MPI_Fint in Section [A.1.2](#) on page [865](#)).
15
- 16 28. Annex [A.1.1](#) on page [849](#). Table **Named Predefined Datatypes**.
17 Added MPI_(U)INT{8,16,32,64}_T, MPI_AINT, MPI_OFFSET, MPI_C_BOOL,
18 MPI_C_FLOAT_COMPLEX, MPI_C_COMPLEX, MPI_C_DOUBLE_COMPLEX, and
19 MPI_C_LONG_DOUBLE_COMPLEX are added as predefined datatypes.
20

21 B.6 Changes from Version 2.0 to Version 2.1

- 22 1. Section [3.2.2](#) on page [33](#), and Annex [A.1](#) on page [849](#).
23 In addition, the MPI_LONG_LONG should be added as an optional type; it is a synonym
24 for MPI_LONG_LONG_INT.
25
- 26 2. Section [3.2.2](#) on page [33](#), and Annex [A.1](#) on page [849](#).
27 MPI_LONG_LONG_INT, MPI_LONG_LONG (as synonym),
28 MPI_UNSIGNED_LONG_LONG, MPI_SIGNED_CHAR, and MPI_WCHAR are moved from
29 optional to official and they are therefore defined for all three language bindings.
30
- 31 3. Section [3.2.5](#) on page [39](#).
32 MPI_GET_COUNT with zero-length datatypes: The value returned as the
33 count argument of MPI_GET_COUNT for a datatype of length zero where zero bytes
34 have been transferred is zero. If the number of bytes transferred is greater than zero,
35 MPI_UNDEFINED is returned.
36
- 37 4. Section [5.1](#) on page [127](#).
38 General rule about derived datatypes: Most datatype constructors have replication
39 count or block length arguments. Allowed values are nonnegative integers. If the
40 value is zero, no elements are generated in the type map and there is no effect on
41 datatype bounds or extent.
42
- 43 5. Section [5.3](#) on page [185](#).
44 MPI_BYTE should be used to send and receive data that is packed using
45 MPI_PACK_EXTERNAL.
46
- 47 6. Section [6.9.6](#) on page [240](#).
48 If comm is an inter-communicator in MPI_ALLREDUCE, then both groups should

provide count and datatype arguments that specify the same type signature (i.e., it is not necessary that both groups provide the same count value).

7. Section 7.3.1 on page 310.

MPI_GROUP_TRANSLATE_RANKS and MPI_PROC_NULL: MPI_PROC_NULL is a valid rank for input to MPI_GROUP_TRANSLATE_RANKS, which returns MPI_PROC_NULL as the translated rank.

8. Section 7.7 on page 359.

About the attribute caching functions:

Advice to implementors. High-quality implementations should raise an error when a keyval that was created by a call to MPI_XXX_CREATE_KEYVAL is used with an object of the wrong type with a call to MPI_YYY_GET_ATTR, MPI_YYY_SET_ATTR, MPI_YYY_DELETE_ATTR, or MPI_YYY_FREE_KEYVAL. To do so, it is necessary to maintain, with each keyval, information on the type of the associated user function. (*End of advice to implementors.*)

9. Section 7.8 on page 376.

In MPI_COMM_GET_NAME: In C, a null character is additionally stored at name[resultlen]. resultlen cannot be larger than MPI_MAX_OBJECT_NAME-1. In Fortran, name is padded on the right with blank characters. resultlen cannot be larger than MPI_MAX_OBJECT_NAME.

10. Section 8.4 on page 385.

About MPI_GRAPH_CREATE and MPI_CART_CREATE: All input arguments must have identical values on all processes of the group of comm_old.

11. Section 8.5.1 on page 386.

In MPI_CART_CREATE: If ndims is zero then a zero-dimensional Cartesian topology is created. The call is erroneous if it specifies a grid that is larger than the group size or if ndims is negative.

12. Section 8.5.3 on page 388.

In MPI_GRAPH_CREATE: If the graph is empty, i.e., nnodes = 0, then MPI_COMM_NULL is returned in all processes.

13. Section 8.5.3 on page 388.

In MPI_GRAPH_CREATE: A single process is allowed to be defined multiple times in the list of neighbors of a process (i.e., there may be multiple edges between two processes). A process is also allowed to be a neighbor to itself (i.e., a self loop in the graph). The adjacency matrix is allowed to be nonsymmetric.

Advice to users. Performance implications of using multiple edges or a nonsymmetric adjacency matrix are not defined. The definition of a node-neighbor edge does not imply a direction of the communication. (*End of advice to users.*)

14. Section 8.5.5 on page 397.

In MPI_CARTDIM_GET and MPI_CART_GET: If comm is associated with a zero-dimensional Cartesian topology, MPI_CARTDIM_GET returns ndims=0 and MPI_CART_GET will keep all output arguments unchanged.

- 1 15. Section 8.5.5 on page 397.
2 In `MPI_CART_RANK`: If `comm` is associated with a zero-dimensional Cartesian topol-
3 ogy, `coord` is not significant and 0 is returned in `rank`.
4
- 5 16. Section 8.5.5 on page 397.
6 In `MPI_CART_COORDS`: If `comm` is associated with a zero-dimensional Cartesian
7 topology, `coords` will be unchanged.
8
- 9 17. Section 8.5.6 on page 405.
10 In `MPI_CART_SHIFT`: It is erroneous to call `MPI_CART_SHIFT` with a direction that is
11 either negative or greater than or equal to the number of dimensions in the Cartesian
12 communicator. This implies that it is erroneous to call `MPI_CART_SHIFT` with a
13 `comm` that is associated with a zero-dimensional Cartesian topology.
14
- 15 18. Section 8.5.7 on page 407.
16 In `MPI_CART_SUB`: If all entries in `remain_dims` are false or `comm` is already associated
17 with a zero-dimensional Cartesian topology then `newcomm` is associated with a zero-
18 dimensional Cartesian topology.
19
- 20 18.1. Section 9.1.1 on page 441.
21 The subversion number changed from 0 to 1.
22
- 23 19. Section 9.1.2 on page 442.
24 In `MPI_GET_PROCESSOR_NAME`: In C, a null character is additionally stored at
25 `name[resultlen]`. `resultlen` cannot be larger than `MPI_MAX_PROCESSOR_NAME-1`. In
26 Fortran, `name` is padded on the right with blank characters. `resultlen` cannot be larger
27 than `MPI_MAX_PROCESSOR_NAME`.
28
- 29 20. Section 9.3 on page 449.
30 `MPI_{COMM,WIN,FILE}_GET_ERRHANDLER` behave as if a new error handler object
31 is created. That is, once the error handler is no longer needed,
32 `MPI_ERRHANDLER_FREE` should be called with the error handler returned from
33 `MPI_ERRHANDLER_GET` or `MPI_{COMM,WIN,FILE}_GET_ERRHANDLER` to mark
34 the error handler for deallocation. This provides behavior similar to that of
35 `MPI_COMM_GROUP` and `MPI_GROUP_FREE`.
36
- 37 21. Section 11.2.1 on page 482, see explanations to `MPI_FINALIZE`.
38 `MPI_FINALIZE` is collective over all connected processes. If no processes were spawned,
39 accepted or connected then this means over `MPI_COMM_WORLD`; otherwise it is col-
40 lective over the union of all processes that have been and continue to be connected,
41 as explained in Section 11.10.4 on page 542.
42
- 43 22. Section 11.2.1 on page 482.
44 About `MPI_ABORT`:
45 *Advice to users.* Whether the errorcode is returned from the executable or from
46 the MPI process startup mechanism (e.g., `mpiexec`), is an aspect of quality of the
47 MPI library but not mandatory. (*End of advice to users.*)
48 *Advice to implementors.* Where possible, a high-quality implementation will try
to return the errorcode from the MPI process startup mechanism (e.g. `mpiexec`
or `singleton init`). (*End of advice to implementors.*)

23. Section 10 on page 473. 1
 An implementation must support info objects as caches for arbitrary (key, value) 2
 pairs, regardless of whether it recognizes the key. Each function that takes hints in 3
 the form of an MPI_Info must be prepared to ignore any key it does not recognize. This 4
 description of info objects does not attempt to define how a particular function should 5
 react if it recognizes a key but not the associated value. MPI_INFO_GET_NKEYS, 6
 MPI_INFO_GET_NTHKEY, MPI_INFO_GET_VALUELEN, and MPI_INFO_GET must 7
 retain all (key,value) pairs so that layered functionality can also use the Info object. 8
9
24. Section 12.3 on page 566. 10
 MPI_PROC_NULL is a valid target rank in the MPI RMA calls MPI_ACCUMULATE, 11
 MPI_GET, and MPI_PUT. The effect is the same as for MPI_PROC_NULL in MPI point- 12
 to-point communication. See also item 25 in this list. 13
14
25. Section 12.3 on page 566. 15
 After any RMA operation with rank MPI_PROC_NULL, it is still necessary to finish the 16
 RMA epoch with the synchronization method that started the epoch. See also item 24 17
 in this list. 18
26. Section 12.3.4 on page 573. 19
 MPI_REPLACE in MPI_ACCUMULATE, like the other predefined operations, is defined 20
 only for the predefined MPI datatypes. 21
22
27. Section 14.2.8 on page 646. 23
 About MPI_FILE_SET_VIEW and MPI_FILE_SET_INFO: When an info object that 24
 specifies a subset of valid hints is passed to MPI_FILE_SET_VIEW or 25
 MPI_FILE_SET_INFO, there will be no effect on previously set or defaulted hints that 26
 the info does not specify. 27
28. Section 14.2.8 on page 646. 28
 About MPI_FILE_GET_INFO: If no hint exists for the file associated with fh, a handle 29
 to a newly created info object is returned that contains no key/value pair. 30
31
29. Section 14.3 on page 650. 32
 If a file does not have the mode MPI_MODE_SEQUENTIAL, then 33
 MPI_DISPLACEMENT_CURRENT is invalid as disp in MPI_FILE_SET_VIEW. 34
35
30. Section 14.5.2 on page 693. 36
 The bias of 16 byte doubles was defined with 10383. The correct value is 16383. 37
31. MPI-2.2, Section 16.1.4 (Section was removed in MPI-3.0). 38
 In the example in this section, the buffer should be declared as `const void* buf`. 39
40
32. Section 19.1.9 on page 805. 41
 About MPI_TYPE_CREATE_F90_XXX: 42
Advice to implementors. An application may often repeat a call to 43
 MPI_TYPE_CREATE_F90_XXX with the same combination of (XXX,p,r). The 44
 application is not allowed to free the returned predefined, unnamed datatype 45
 handles. To prevent the creation of a potentially huge amount of handles, the 46
 MPI implementation should return the same datatype handle for the same (47
 REAL/COMPLEX/INTEGER,p,r) combination. Checking for the combination (48

1 p,r) in the preceding call to MPI_TYPE_CREATE_F90_XXX and using a hash-
2 table to find formerly generated handles should limit the overhead of finding
3 a previously generated datatype with same combination of (XXX,p,r). (*End of*
4 *advice to implementors.*)

5
6 33. Section [A.1.1](#) on page [849](#).

7 MPI_BOTTOM is defined as `void * const MPI::BOTTOM`.

8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

DRAFT

Bibliography

- [1] MPI Forum Documents. <https://www.mpi-forum.org/docs/>. Citation on page 6.
- [2] Reverse domain name notation convention. <https://docs.oracle.com/javase/tutorial/java/package/namingpkgs.html>. Citation on page 339.
- [3] V. Bala and S. Kipnis. Process groups: a mechanism for the coordination of and communication among processes in the Venus collective communication library. Technical report, IBM T. J. Watson Research Center, October 1992. Preprint. Citation on page 2.
- [4] V. Bala, S. Kipnis, L. Rudolph, and Marc Snir. Designing efficient, scalable, and portable collective communication libraries. Technical report, IBM T. J. Watson Research Center, October 1992. Preprint. Citation on page 2.
- [5] Purushotham V. Bangalore, Nathan E. Doss, and Anthony Skjellum. MPI++: Issues and features. In *OON-SKI '94*, page in press, 1994. Citation on page 305.
- [6] A. Beguelin, J. Dongarra, A. Geist, R. Manchek, and V. Sunderam. Visualization and debugging in a heterogeneous environment. *IEEE Computer*, 26(6):88–95, June 1993. Citation on page 2.
- [7] Luc Bomans and Rolf Hempel. The Argonne/GMD macros in FORTRAN for portable parallel programming and their implementation on the Intel iPSC/2. *Parallel Computing*, 15:119–132, 1990. Citation on page 2.
- [8] Dan Bonachea and Jason Duell. Problems with using MPI 1.1 and 2.0 as compilation targets for parallel language implementations. *IJHPCN*, 1(1/2/3):91–99, 2004. Citation on page 607.
- [9] Rajesh Bordawekar, Juan Miguel del Rosario, and Alok Choudhary. Design and evaluation of primitives for parallel I/O. In *Proceedings of Supercomputing '93*, pages 452–461, 1993. Citation on page 637.
- [10] R. Butler and E. Lusk. User’s guide to the p4 programming system. Technical Report TM-ANL-92/17, Argonne National Laboratory, 1992. Citation on page 2.
- [11] Ralph Butler and Ewing Lusk. Monitors, messages, and clusters: The p4 parallel programming system. *Parallel Computing*, 20(4):547–564, April 1994. Also Argonne National Laboratory Mathematics and Computer Science Division preprint P362-0493. Citation on page 2.
- [12] Robin Calkin, Rolf Hempel, Hans-Christian Hoppe, and Peter Wypior. Portable programming with the PARMACS message-passing library. *Parallel Computing*, 20(4):615–632, April 1994. Citation on page 2.

- [13] S. Chittor and R. J. Enbody. Performance evaluation of mesh-connected wormhole-routed networks for interprocessor communication in multicomputers. In *Proceedings of the 1990 Supercomputing Conference*, pages 647–656, 1990. Citation on page 383.
- [14] S. Chittor and R. J. Enbody. Predicting the effect of mapping on the communication performance of large multicomputers. In *Proceedings of the 1991 International Conference on Parallel Processing, vol. II (Software)*, pages II-1–II-4, 1991. Citation on page 383.
- [15] Parasoftware Corporation. Express version 1.0: A communication environment for parallel computers, 1988. Citation on page 2.
- [16] Yiannis Cotronis, Anthony Danalis, Dimitrios S. Nikolopoulos, and Jack Dongarra, editors. *Recent Advances in the Message Passing Interface - 18th European MPI Users' Group Meeting, EuroMPI 2011, Santorini, Greece, September 18-21, 2011. Proceedings*, volume 6960 of *Lecture Notes in Computer Science*. Springer, 2011. Citations on pages 1070 and 1072.
- [17] Juan Miguel del Rosario, Rajesh Bordawekar, and Alok Choudhary. Improved parallel I/O via a two-phase run-time access strategy. In *IPPS '93 Workshop on Input/Output in Parallel Computer Systems*, pages 56–70, 1993. Also published in *Computer Architecture News* 21(5), December 1993, pages 31–38. Citation on page 637.
- [18] James Dinan, Sriram Krishnamoorthy, Pavan Balaji, Jeff R. Hammond, Manojkumar Krishnan, Vinod Tipparaju, and Abhinav Vishnu. Noncollective communicator creation in MPI. In Cotronis et al. [16], pages 282–291. Citation on page 328.
- [19] J. Dongarra, A. Geist, R. Manchek, and V. Sunderam. Integrated PVM framework supports heterogeneous network computing. *Computers in Physics*, 7(2):166–75, April 1993. Citation on page 2.
- [20] J. J. Dongarra, R. Hempel, A. J. G. Hey, and D. W. Walker. A proposal for a user-level, message passing interface in a distributed memory environment. Technical Report TM-12231, Oak Ridge National Laboratory, February 1993. Citation on page 2.
- [21] Edinburgh Parallel Computing Centre, University of Edinburgh. *CHIMP Concepts*, June 1991. Citation on page 2.
- [22] Edinburgh Parallel Computing Centre, University of Edinburgh. *CHIMP Version 1.0 Interface*, May 1992. Citation on page 2.
- [23] D. Feitelson. Communicators: Object-based multiparty interactions for parallel programming. Technical Report 91-12, Dept. Computer Science, The Hebrew University of Jerusalem, November 1991. Citation on page 306.
- [24] Message Passing Interface Forum. MPI: A message-passing interface standard. *The International Journal of Supercomputer Applications and High Performance Computing*, 8, 1994. Citation on page 2.
- [25] Message Passing Interface Forum. MPI: A message-passing interface standard (version 1.1). Technical report, 1995. <http://www.mpi-forum.org>. Citation on page 2.

- [26] Al Geist, Adam Beguelin, Jack Dongarra, Weicheng Jiang, Bob Manchek, and Vaidy Sunderam. *PVM: Parallel Virtual Machine—A User’s Guide and Tutorial for Network Parallel Computing*. MIT Press, 1994. Citation on page [481](#). 1
2
3
4
- [27] G. A. Geist, M. T. Heath, B. W. Peyton, and P. H. Worley. PICL: A portable instrumented communications library, C reference manual. Technical Report TM-11130, Oak Ridge National Laboratory, Oak Ridge, TN, July 1990. Citation on page [2](#). 5
6
7
- [28] Brice Goglin, Emmanuel Jeannot, Farouk Mansouri, and Guillaume Mercier. Hardware topology management in MPI applications through hierarchical communicators. *Parallel Computing*, 76:70–90, 2018. Citation on page [338](#). 8
9
10
11
- [29] Ryan E. Grant, Matthew G. F. Dosanjh, Michael J. Levenhagen, Ron Brightwell, and Anthony Skjellum. Finepoints: Partitioned multithreaded MPI communication. In *ISC High Performance Conference (ISC)*, 2019. Citation on page [113](#). 12
13
14
15
- [30] Ryan E. Grant, Anthony Skjellum, and Purushotham V. Bangalore. Lightweight threading with MPI using persistent communications semantics. In *Workshop on Exascale MPI (ExaMPI)*. Held in conjunction with the 2015 International Conference for High Performance Computing, Networking, Storage and Analysis (SC15), 2015. Citation on page [113](#). 16
17
18
19
20
- [31] D. Gregor, T. Hoefler, B. Barrett, and A. Lumsdaine. Fixing probe for multi-threaded MPI applications. Technical Report 674, Indiana University, Jan. 2009. Citation on page [97](#). 21
22
23
24
- [32] William D. Gropp and Barry Smith. Chameleon parallel programming tools users manual. Technical Report ANL-93/23, Argonne National Laboratory, March 1993. Citation on page [2](#). 25
26
27
28
- [33] Michael Hennecke. A Fortran 90 interface to MPI version 1.1. Technical Report Internal Report 63/96, Rechenzentrum, Universität Karlsruhe, D-76128 Karlsruhe, Germany, June 1996. Citation on page [792](#). 29
30
31
- [34] T. Hoefler, G. Bronevetsky, B. Barrett, B. R. de Supinski, and A. Lumsdaine. Efficient MPI support for advanced hybrid programming models. In *Recent Advances in the Message Passing Interface (EuroMPI’10)*, volume LNCS 6305, pages 50–61. Springer, Sep. 2010. Citation on page [97](#). 32
33
34
35
36
- [35] T. Hoefler, P. Gottschling, A. Lumsdaine, and W. Rehm. Optimizing a conjugate gradient solver with non-blocking collective operations. *Elsevier Journal of Parallel Computing (PARCO)*, 33(9):624–633, Sep. 2007. Citation on page [251](#). 37
38
39
40
- [36] T. Hoefler, F. Lorenzen, and A. Lumsdaine. Sparse non-blocking collectives in quantum mechanical calculations. In *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 15th European PVM/MPI Users’ Group Meeting*, volume LNCS 5205, pages 55–63. Springer, Sep. 2008. Citation on page [410](#). 41
42
43
44
- [37] T. Hoefler and A. Lumsdaine. Message progression in parallel computing — to thread or not to thread? In *Proceedings of the 2008 IEEE International Conference on Cluster Computing*. IEEE Computer Society, Oct. 2008. Citation on page [251](#). 45
46
47
48

- [38] T. Hoefer, A. Lumsdaine, and W. Rehm. Implementation and performance analysis of non-blocking collective operations for MPI. In *Proceedings of the 2007 International Conference on High Performance Computing, Networking, Storage and Analysis, SC07*. IEEE Computer Society/ACM, Nov. 2007. Citation on page [252](#).
- [39] T. Hoefer, M. Schellmann, S. Gorchach, and A. Lumsdaine. Communication optimization for medical image reconstruction algorithms. In *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 15th European PVM/MPI Users' Group Meeting*, volume LNCS 5205, pages 75–83. Springer, Sep. 2008. Citation on page [251](#).
- [40] T. Hoefer and J. L. Träff. Sparse collective operations for MPI. In *Proceedings of the 23rd IEEE International Parallel & Distributed Processing Symposium, HIPS'09 Workshop*, May 2009. Citation on page [410](#).
- [41] Torsten Hoefer and Marc Snir. Writing parallel libraries with MPI — common practice, issues, and extensions. In Cotronis et al. [[16](#)], pages 345–355. Citation on page [324](#).
- [42] Daniel J. Holmes, Bradley Morgan, Anthony Skjellum, Purushotham V. Bangalore, and Srinivas Sridharan. Planning for performance: Enhancing achievable performance for MPI through persistent collective operations. *Parallel Computing*, 81:32 – 57, 2019. Citation on page [274](#).
- [43] Institute of Electrical and Electronics Engineers, New York. *IEEE Standard for Binary Floating-Point Arithmetic, IEEE Standard 754-2008*, 2008. Citation on page [693](#).
- [44] International Organization for Standardization, Geneva. *ISO 8859-1:1987: Information processing — 8-bit single-byte coded graphic character sets — Part 1: Latin alphabet No. 1*, February 1987. Citation on page [693](#).
- [45] International Organization for Standardization, Geneva. *ISO/IEC 9945-1:1996: Information technology — Portable Operating System Interface (POSIX) — Part 1: System Application Program Interface (API) [C Language]*, December 1996. Citations on pages [514](#) and [641](#).
- [46] International Organization for Standardization, Geneva. *ISO/IEC 1539-1:2010: Information technology — Programming languages — Fortran — Part 1: Base language*, November 2010. Citations on pages [787](#) and [790](#).
- [47] International Organization for Standardization, Geneva. *ISO/IEC TS 29113:2012: Information technology — Further interoperability of Fortran with C*, December 2012. Citations on pages [787](#), [790](#), [803](#), [804](#), and [1059](#).
- [48] International Organization for Standardization, Geneva. *ISO/IEC 1539-1:2018: Information technology — Programming languages — Fortran — Part 1: Base language*, November 2018. Citations on pages [787](#), [788](#), and [790](#).
- [49] Charles H. Koelbel, David B. Loveman, Robert S. Schreiber, Guy L. Steele Jr., and Mary E. Zosel. *The High Performance Fortran Handbook*. MIT Press, 1993. Citation on page [143](#).
- [50] David Kotz. Disk-directed I/O for MIMD multiprocessors. In *Proceedings of the 1994 Symposium on Operating Systems Design and Implementation*, pages 61–74, November

1994. Updated as Dartmouth TR PCS-TR94-226 on November 8, 1994. Citation on page 637.
- [51] O. Krämer and H. Mühlenbein. Mapping strategies in message-based multiprocessor systems. *Parallel Computing*, 9:213–225, 1989. Citation on page 383.
- [52] S. J. Lefflet, R. S. Fabry, W. N. Joy, P. Lapsley, S. Miller, and C. Torek. An advanced 4.4BSD interprocess communication tutorial, Unix programmer’s supplementary documents (PSD) 21. Technical report, Computer Systems Research Group, Department of Electrical Engineering and Computer Science, University of California, Berkeley, 1993. Available online: <https://docs.freebsd.org/44doc/psd/21.ipc/paper.pdf>. Citation on page 545.
- [53] Bradley Morgan, Daniel J. Holmes, Anthony Skjellum, Purushotham Bangalore, and Srinivas Sridharan. Planning for performance: Persistent collective operations for MPI. In *Proceedings of the 24th European MPI Users’ Group Meeting*, EuroMPI ’17, pages 4:1–4:11, New York, NY, USA, 2017. ACM. Citation on page 274.
- [54] nCUBE Corporation. *nCUBE 2 Programmers Guide, r2.0*, December 1990. Citation on page 2.
- [55] Bill Nitzberg. Performance of the iPSC/860 Concurrent File System. Technical Report RND-92-020, NAS Systems Division, NASA Ames, December 1992. Citation on page 637.
- [56] William J. Nitzberg. *Collective Parallel I/O*. PhD thesis, Department of Computer and Information Science, University of Oregon, December 1995. Citation on page 637.
- [57] *4.4BSD Programmer’s Supplementary Documents (PSD)*. O’Reilly and Associates, 1994. Citation on page 545.
- [58] Paul Pierce. The NX/2 operating system. In *Proceedings of the Third Conference on Hypercube Concurrent Computers and Applications*, pages 384–390. ACM Press, 1988. Citation on page 2.
- [59] Martin Schulz and Bronis R. de Supinski. P^NMPI tools: A whole lot greater than the sum of their parts. In *ACM/IEEE Supercomputing Conference (SC)*, pages 1–10. ACM, 2007. Citation on page 722.
- [60] K. E. Seamons, Y. Chen, P. Jones, J. Jozwiak, and M. Winslett. Server-directed collective I/O in Panda. In *Proceedings of Supercomputing ’95*, December 1995. Citation on page 637.
- [61] A. Skjellum and A. Leung. Zipcode: a portable multicomputer communication library atop the reactive kernel. In D. W. Walker and Q. F. Stout, editors, *Proceedings of the Fifth Distributed Memory Concurrent Computing Conference*, pages 767–776. IEEE Press, 1990. Citations on pages 2 and 306.
- [62] A. Skjellum, S. Smith, C. Still, A. Leung, and M. Morari. The Zipcode message passing system. Technical report, Lawrence Livermore National Laboratory, September 1992. Citation on page 2.

- 1 [63] Anthony Skjellum, Nathan E. Doss, and Purushotham V. Bangalore. Writing libraries
2 in MPI. In Anthony Skjellum and Donna S. Reese, editors, *Proceedings of the Scalable*
3 *Parallel Libraries Conference*, pages 166–173. IEEE Computer Society Press, October
4 1993. Citation on page 305.
- 5 [64] Anthony Skjellum, Nathan E. Doss, and Kishore Viswanathan. Inter-communicator ex-
6 tensions to MPI in the MPIX (MPI eXtension) Library. Technical Report MSU-940722,
7 Mississippi State University — Dept. of Computer Science, August 1994. Archived at
8 <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.49.6283>. Citation on
9 page 193.
- 10 [65] Anthony Skjellum, Steven G. Smith, Nathan E. Doss, Alvin P. Leung, and Manfred
11 Morari. The design and evolution of Zipcode. *Parallel Computing*, 20(4):565–596, April
12 1994. Citations on pages 306 and 349.
- 13 [66] The Internet Society. XDR: External data representation standard, May 2006. <http://www.rfc-editor.org/pdf/rfc/rfc4506.txt.pdf>. Citation on page 693.
- 14 [67] Rajeev Thakur and Alok Choudhary. An extended two-phase method for accessing
15 sections of out-of-core arrays. *Scientific Programming*, 5(4):301–317, Winter 1996.
16 Citation on page 637.
- 17 [68] Jesper Larsson Träff. SMP-aware message passing programming. In *Eighth Inter-*
18 *national Workshop on High-level Parallel Programming Models and Supportive Envi-*
19 *ronments (HIPS), 17th International Parallel and Distributed Processing Symposium*
20 *(IPDPS)*, pages 56–65, 2003. Citation on page 338.
- 21 [69] *The Unicode Standard, Version 2.0*. Addison-Wesley, 1996. ISBN 0-201-48345-9. Ci-
22 tation on page 693.
- 23 [70] D. Walker. Standards for message passing in a distributed memory environment. Tech-
24 nical Report TM-12147, Oak Ridge National Laboratory, August 1992. Citation on
25 page 2.
- 26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48

General Index

This index lists mainly terms of the MPI specification. The underlined page numbers refer to the definitions or parts of the definition of the terms. Bold face page numbers are used for entries appearing in a section or subsection title.

- !(_c), [833](#)
- _c, [21](#), [22](#), [235](#), [833](#)
- a priori enabled, [13](#)
- absolute addresses, [21](#), [148](#), [823](#)
- access epoch, [588](#)
- action
 - in function names, [10](#)
- active, [79](#), [80](#), [81](#), [103](#), [109](#), [110](#), [381](#)
 - operation state, [12](#), [1045](#)
- active target communication, [587](#)
- addresses, [161](#)
 - absolute, [21](#), [22](#), [148](#), [823](#)
 - correct use, [161](#)
 - relative displacement, [21](#), [148](#)
- alignment, [447](#), [554](#), [556](#), [1051](#)
- all-reduce, [240](#)
 - nonblocking, [268](#)
 - persistent, [291](#)
- all-to-all, [218](#)
 - nonblocking, [262](#)
 - persistent, [285](#)
- array arguments, [19](#)
- assertions, [603](#)
- ASYNCHRONOUS
 - Fortran attribute, [825](#)
- attribute, [306](#), [359](#), [843](#)
 - caching, [306](#)
- automatic buffering
 - buffer allocation, [58](#)
- barrier synchronization, [196](#)
 - nonblocking, [253](#)
 - persistent, [275](#)
- basic datatypes, [33](#)
 - additional host language, [34](#)
 - byte, [33](#)
 - C, [34](#)
 - C and Fortran, [35](#)
 - C++, [35](#)
 - Fortran, [33](#)
 - packed, [33](#)
- blocked
 - procedure invocation, [27](#)
- blocking, [15](#), [50](#), [54](#), [99](#), [653](#)
 - I/O, [654](#)
 - point-to-point, [32](#)
- blocking I/O, [653](#)
- blocking operation, [12](#)
- blocking procedure, [15](#)
- bounds of datatypes, [153](#)
- broadcast, [196](#)
 - nonblocking, [253](#)
 - persistent, [276](#)
- buffered mode send, [50](#), [54](#), [103](#), [106](#)
 - buffer allocation, [58](#)
 - nonblocking, [70](#), [71](#), [72](#)
- C
 - language binding, [24](#)
- caching, [305](#), [306](#), [359](#)
- callback functions
 - language interoperability, [842](#)
 - prototype definitions, [866](#)
 - deprecated, [872](#)
- cancel, [24](#), [80](#), [94](#), [94](#), [95](#), [102](#), [103](#), [104](#), [490](#), [774](#), [1050](#), [1051](#)
- cancelled, [94](#), [95](#), [102–104](#)
- canonical pack and unpack, [185](#)
- Cartesian
 - topology, [385](#), [386](#)
- Chameleon, [2](#)
- change-log, [1045](#)
- Chimp, [2](#)
- choice, [20](#)
- class
 - in function names, [10](#)
- clock synchronization, [444](#)
- collective, [15](#), [653](#)
 - split, [701](#)
- collective communication, [189](#)
 - correctness, [297](#)
 - file data access operations, [677](#)
 - neighborhood, [410](#)
 - nonblocking, [250](#)
 - progress, [251](#), [252](#), [301](#), [303](#)
- collective operation, [13](#)
- collective procedure, [15](#)

- 1 commit, [156](#)
2 COMMON blocks, [828](#)
3 communication, [547](#)
4 collective, [189](#)
5 modes, [50](#)
6 one-sided, [547](#)
7 overlap with communication, [58](#), [69](#)
8 overlap with computation, [58](#), [69](#)
9 partitioned point-to-point, [113](#)
10 pending, [12](#), [18](#), [56](#), [70](#), [71](#), [83](#), [89](#), [94](#), [102](#),
11 [252](#), [306–308](#), [322](#), [516](#), [826](#), [832](#), [1045](#)
12 point-to-point, [31](#)
13 RMA, [547](#)
14 typed, untyped, and packed, [46](#)
15 communication modes, [50](#)
16 communicator, [36](#), [305](#), [306](#)
17 hidden, [192](#), [300](#), [596](#)
18 complete operation, [11](#)
19 complete operation state, [12](#)
20 completing, [15](#)
21 completing procedure, [15](#)
22 completion, [50](#), [69](#), [78](#), [79](#), [80](#), [83](#), [104](#)
23 multiple, [83](#), [84–88](#), [91](#), [92](#)
24 completion stage, [11](#)
25 concurrent, [701](#), [702](#)
26 conflict, [701](#)
27 connected, [542](#)
28 constants, [19](#), [847](#), [849](#), [1045](#)
29 C integer constant expressions, [20](#)
30 context, [305](#), [306](#), [308](#)
31 control variable
32 tool information interface, [729](#)
33 conversion, [48](#)
34 representation, [49](#)
35 type, [49](#)
36 coordination, [653](#)
37 counts, [21](#)
38 Cray-pointers, [448](#)
39 create
40 in function names, [9](#)
41 data, [33](#)
42 data conversion, [48](#)
43 datatype, [16](#)
44 derived, [16](#)
45 equivalent, [17](#)
46 named, [16](#)
47 portable, [16](#)
48 predefined, [16](#)
49 unnamed, [16](#)
50 datatypes, [127](#), [840](#)
51 deadlock, [28](#), [617](#)
52 deadlock avoidance
53 cyclic shift, [43](#)
54 lack of buffer space, [58](#)
55 nonblocking communication, [69](#), [83](#)
56 send modes, [58](#)
57 decoupled MPI activities, [18](#), [27](#), [56](#), [306–308](#),
58 [322](#), [544](#), [1045](#)
59 default file error, [711](#)
60 delete
61 in function names, [9](#)
62 deprecated interfaces, [22](#), [771](#)
63 derived datatype, [16](#), [127](#), [819](#)
64 disconnected, [543](#)
65 displacement, [637](#), [651](#)
66 distributed graph
67 topology, [385](#), [390](#)
68 Dynamic Process Model, [481](#)
69 dynamically attached memory, [559](#)
70 elementary datatype, [637](#), [651](#)
71 empty, [79](#), [80](#), [84](#)
72 enabled operation state, [13](#), [16](#), [79](#), [1047](#)
73 a priori, [13](#)
74 end of file, [638](#)
75 envelope, [31](#), [32](#), [35](#), [35](#), [36](#), [38](#), [40](#), [49](#), [50](#), [68](#)
76 data representation conversion, [49](#)
77 environmental inquiries, [442](#)
78 equivalent datatypes, [17](#)
79 erroneous program, [38](#)
80 freeing active request, [81](#)
81 invalid matched receive, [101](#), [102](#)
82 lack of buffer space, [56](#)
83 ready mode before receive, [51](#)
84 type matching, [46](#)
85 error classes, [461](#)
86 error handling, [26](#), [449](#)
87 default file error handler, [639](#), [643](#), [711](#)
88 error codes and classes, [460](#), [464](#)
89 error handlers, [452](#), [464](#), [842](#)
90 fatal after request free, [82](#)
91 finalize, [26](#), [490](#), [491](#), [543](#)
92 I/O, [711](#), [712](#)
93 initial error handler, [26](#), [450](#), [451](#), [483](#),
94 [490](#), [527](#)
95 multiple completions, [86](#), [88](#)
96 one-sided communication, [605](#)
97 process failure, [26](#), [543](#)
98 program error, [26](#)
99 resource error, [26](#)
100 startup, [26](#), [483](#), [527](#)
101 transmission failure, [26](#)
102 establishing communication, [530](#)
103 etype, [637](#), [651](#)
104 event
105 tool information interface, [748](#)
106 event types, [748](#)

- events, [736](#), [748](#)
- examples, [30](#)
- exclusive lock, [589](#), [597](#), [610](#), [611](#)
- exclusive scan
 - nonblocking, [272](#)
 - persistent, [296](#)
- explicit offsets, [654](#), [656](#)
- exposure epoch, [588](#)
- extent of datatypes, [128](#), [152](#), [153](#)
 - true extent, [155](#)
- external32
 - file data representation, [691](#)
- extra-state, [846](#)
- fairness
 - not guaranteed, [56](#)
 - requirement, [89](#), [93](#)
- file, [637](#)
 - data access, [653](#)
 - collective operations, [677](#)
 - explicit offsets, [656](#)
 - individual file pointers, [663](#)
 - seek, [679](#)
 - shared file pointers, [673](#)
 - split collective, [681](#)
 - end of file, [638](#)
 - filetype, [637](#)
 - handle, [639](#)
 - interoperability, [689](#)
 - manipulation, [639](#)
 - nonblocking collective
 - progress, [705](#)
 - offset, [21](#), [638](#)
 - pointer, [639](#)
 - progress, [704](#)
 - size, [638](#)
 - view, [637](#), [638](#), [650](#)
- file access
 - concurrent, [701](#)
 - conflict, [701](#)
 - overlap, [701](#)
- file size, [706](#)
- finalize, [488](#)
- finished, [491](#)
- Fortran
 - language binding, [23](#), [787](#)
- Fortran 2018, [20](#), [23](#), [787–792](#), [794](#), [796](#), [797](#), [800](#), [831](#), [1047](#)
- Fortran support, [787](#)
- freed operation state, [12](#)
- freeing procedure, [15](#)
- freeing stage, [11](#), [81](#), [90](#)
- gather, [198](#)
 - nonblocking, [255](#)
 - persistent, [277](#)
- gather-to-all, [214](#)
 - nonblocking, [260](#)
 - persistent, [282](#)
- general datatype, [127](#)
- generalized requests, [627](#), [627](#)
 - progress, [627](#)
- get
 - in function names, [9](#)
- graph
 - topology, [385](#), [388](#)
- group, [36](#), [305](#), [306](#), [308](#), [351](#)
- group objects, [308](#)
- guarantee progress, [28](#)
- handles, [17](#), [836](#)
- hardware resource type, [333](#), [334](#), [336](#)
- host rank, [778](#)
- I/O
 - blocking, [653](#)
 - nonblocking, [653](#)
- immediate, [15](#), [69](#), [71](#), [85](#), [88](#), [89](#), [92](#), [95](#), [102](#)
- inactive, [78](#), [79–81](#), [84](#), [85](#), [103](#), [109](#), [110](#)
 - operation state, [12](#)
- inclusive scan, [247](#)
 - nonblocking, [271](#)
 - persistent, [294](#)
- incomplete, [15](#), [71](#), [99](#), [100](#)
- incomplete procedure, [15](#)
- independent, [543](#)
- individual file pointers, [654](#), [663](#)
- info object, [473](#)
 - file info, [646](#)
 - keys, [873](#)
 - values, [874](#)
- initial error handler, [26](#)
- initialization, [98](#), [104](#), [109](#)
- initialization procedure, [14](#)
- initialization stage, [11](#)
- initialized operation state, [12](#)
- initiation, [11](#), [69](#), [71](#), [78](#), [82](#), [96](#)
- initiation procedure, [15](#)
- inter-communication, [307](#), [350](#)
- inter-communicator, [192](#), [307](#), [350](#)
 - collective operations, [193](#), [194](#)
 - point-to-point, [36](#), [38](#)
- interlanguage communication, [847](#)
- internal
 - file data representation, [690](#)
- interoperability, [689](#)
- intra-communication, [307](#), [349](#)
- intra-communicator, [192](#), [306](#), [349](#)

- 1 collective operations, [192](#)
- 2 intra-communicator objects, [309](#)
- 3 I/O, [637](#)
- 4 IO rank, [443](#)
- 5 is
- 6 in function names, [9](#)
- 7 language binding, [22](#), [787](#)
- 8 interoperability, [834](#)
- 9 summary, [849](#)
- 10 large
- 11 count, [22](#), [833](#)
- 12 displacement, [833](#)
- 13 large count, [22](#)
- 14 lb_marker, [142](#), [143](#), [146](#), [152](#), [156](#)
- 15 erased, [155](#)
- 16 local, [14](#), [50](#), [69](#), [70](#), [79–81](#), [87](#), [89](#), [95](#), [98](#), [102](#),
[104](#), [109](#), [497](#), [595](#)
- 17 local group, [322](#)
- 18 local procedure, [14](#), [27](#)
- 19 call, under constraint, [16](#)
- 20 logical resource type, [334](#)
- 21 logically concurrent, [55](#)
- 22 loosely synchronous model, [380](#)
- 23 lower bound, [152](#)
- 24 lower-bound markers, [151](#)
- 25 macros, [25](#)
- 26 main thread, [516](#)
- 27 matched probe
- 28 progress, [99](#)
- 29 matched receive, [98](#), [100](#), [100](#)
- 30 matching
- 31 type, [159](#), [705](#)
- 32 matching probe, [96](#), [97](#), [97](#), [98–100](#)
- 33 matching receive, [99](#)
- 34 matching rules
- 35 blocking with nonblocking, [70](#)
- 36 cancel, [103](#)
- 37 envelope, [38](#)
- 38 null process, [111](#)
- 39 ordering, [55](#)
- 40 persistent collective, [274](#)
- 41 persistent point-to-point, [110](#)
- 42 probe, [95](#), [97](#)
- 43 send modes, [54](#)
- 44 type, [46](#), [46](#)
- 45 wildcard, [38](#)
- 46 memory
- 47 alignment, [447](#), [554](#), [556](#), [1051](#)
- 48 allocation, [446](#), [552](#), [554](#)
- 49 system, [17](#)
- 50 memory model, [547](#), [586](#)
- 51 separate, [547](#), [558](#)
- 52 unified, [547](#), [558](#)
- 53 message, [31](#), [97](#)
- 54 buffer, [52](#)
- 55 cancel, [80](#), [94](#), [95](#), [102–104](#)
- 56 data, [31](#), [32](#), [33](#), [50](#), [52](#)
- 57 envelope, [31](#), [32](#), [35](#), [35](#), [36](#), [38](#), [40](#), [49](#), [50](#),
[68](#)
- 58 handle, [97](#), [98](#), [101](#), [102](#)
- 59 intermediate buffering, [46](#), [50](#)
- 60 invalid handle, [99](#)
- 61 predefined handle, [99](#)
- 62 wildcard, [38](#)
- 63 message handle, [97](#), [98](#), [101](#), [102](#)
- 64 invalid, [99](#)
- 65 predefined, [99](#)
- 66 modes, [50](#)
- 67 buffered, [50](#), [54](#)
- 68 ready, [51](#), [54](#)
- 69 standard, [50](#), [54](#)
- 70 synchronous, [51](#), [54](#)
- 71 module variables, [828](#)
- 72 MPI datatype, [16](#)
- 73 mpi module
- 74 Fortran support, [791](#)
- 75 MPI operation, [11](#)
- 76 MPI procedure, [14](#)
- 77 MPI process initialization, [481](#)
- 78 Dynamic Process Model, [481](#)
- 79 Sessions Model, [309](#), [312](#), [443](#), [481](#), [516](#),
[640](#)
- 80 World Model, [309](#), [312](#), [442](#), [481](#), [516](#), [639](#),
[711](#)
- 81 MPI Session handle, [493](#)
- 82 mpi_f08 module
- 83 Fortran support, [788](#)
- 84 MPI_SIZEOF and storage_size(), [24](#), [777](#),
[811–813](#)
- 85 mpiexec, [484](#), [508](#), [512](#)
- 86 mpif.h include file
- 87 Fortran support, [793](#)
- 88 mpirun, [512](#)
- 89 multiple completions, [83](#), [84–88](#), [91](#), [92](#)
- 90 error handling, [86](#), [88](#)
- 91 named datatype, [16](#)
- 92 names, [530](#)
- 93 name publishing, [535](#)
- 94 naming objects, [376](#)
- 95 native
- 96 file data representation, [690](#)
- 97 neighborhood collective communication, [410](#)
- 98 nonblocking, [422](#)
- 99 periodic and dims=1 or 2, [1049](#)
- 100 nonblocking, [15](#), [69](#), [81](#), [98](#), [102](#), [566](#), [653](#)

- communication, [69](#)
- completion, [78](#), [80](#)
- Fortran problems, [822](#)
- I/O, [654](#)
- initiation, [71](#)
- persistent
 - partitioned completion, [120](#)
 - request objects, [70](#)
- nonblocking I/O, [653](#)
- nonblocking operation, [12](#)
- nonblocking procedure, [15](#)
- noncollective operation, [13](#)
- noncollective procedure, [16](#)
- nonlocal, [14](#), [50](#), [51](#), [70](#), [79](#), [96](#), [99](#), [100](#)
- nonlocal procedure, [14](#), [27](#)
 - call, under constraint, [16](#)
- nonovertaking, [55](#), [82](#)
- null handle, [78](#), [84–86](#), [92](#)
- null handles, [87](#), [91](#), [92](#)
- null processes, [111](#)
- offset, [21](#), [638](#)
- one-sided communication, [547](#)
 - Fortran problems, [823](#)
 - progress, [615](#), [616](#), [617](#)
- opaque objects, [17](#), [840](#)
- operation, [11](#)
 - blocking, [12](#)
 - collective, [13](#)
 - complete, [11](#)
 - nonblocking, [12](#)
 - noncollective, [13](#)
 - operation-related, [14](#)
 - partitioned receive, [12](#)
 - partitioned send, [12](#)
 - pending, [12](#), [18](#), [56](#), [70](#), [71](#), [83](#), [89](#), [94](#), [102](#), [252](#), [306–308](#), [322](#), [516](#), [705](#), [826](#), [832](#), [1045](#)
 - persistent, [12](#)
 - semantics, [11](#)
 - stage, [11](#)
 - completion, [11](#)
 - freeing, [11](#)
 - initialization, [11](#)
 - starting, [11](#)
 - state, [12](#)
 - active, [12](#)
 - complete, [12](#)
 - enabled, [13](#), [16](#), [79](#), [1047](#)
 - freed, [12](#)
 - inactive, [12](#)
 - initialized, [12](#)
 - started, [12](#)
- ordered, [55](#), [82](#)
- origin, [548](#)
- overlap, [701](#)
- pack, [178](#)
 - canonical, [185](#)
- packing unit, [181](#)
- parallel procedure, [380](#)
- partitioned completion, [120](#)
- partitioned point-to-point communication, [113](#)
- passive target communication, [588](#)
- pending (communication) operation, [12](#), [18](#), [56](#), [70](#), [71](#), [83](#), [89](#), [94](#), [102](#), [252](#), [306–308](#), [322](#), [516](#), [826](#), [832](#), [1045](#)
- pending communication affector (Fortran), [826](#)
- pending I/O operation, [12](#), [516](#), [705](#), [1045](#)
- performance variable
 - tool information interface, [736](#)
- persistent communication request, [78–81](#), [84–88](#), [103](#), [104](#), [104](#), [105–110](#)
 - active, [79](#)
 - completion, [80](#), [104](#)
 - inactive, [78](#)
 - freeing, [110](#)
 - starting, [104](#), [109](#)
- persistent communication requests
 - collective persistent, [273](#), [429](#)
 - Fortran problems, [823](#)
- persistent operation, [12](#)
- PICL, [2](#)
- PMPI_, [717](#)
- point-to-point communication, [31](#)
 - blocking, [32](#)
 - buffer allocation, [58](#)
 - cancel, [102](#)
 - matched receive, [100](#)
 - matching probe, [97](#)
 - nonblocking, [69](#)
 - persistent, [104](#)
 - probe, [94](#)
 - receive operation, [36](#)
 - send modes, [50](#), [54](#)
 - send operation, [32](#)
 - send-receive operation, [43](#)
 - status, [39](#)
- portable datatype, [16](#)
- ports, [530](#)
- positioning, [653](#)
- POSIX
 - environment, [470](#)
 - FORTTRAN, [18](#)
 - I/O, [640](#), [641](#), [653](#), [690](#), [702](#)
 - model, [637](#)
- predefined datatype, [16](#)
- predefined reduction operations, [227](#)

- 1 private window copy, [586](#)
- 2 probe, [94](#), [94](#), [95](#), [98](#)
- 3 matching, [97](#)
- 4 progress, [96](#)
- 5 procedure, [14](#)
- 6 blocking, [15](#)
- 7 collective, [15](#)
- 8 completing, [15](#)
- 9 freeing, [15](#)
- 10 immediate, [15](#)
- 11 incomplete, [15](#)
- 12 initialization, [14](#)
- 13 initiation, [15](#)
- 14 local, [14](#), [27](#)
- 15 call, under constraint, [16](#)
- 16 nonblocking, [15](#)
- 17 noncollective, [16](#)
- 18 nonlocal, [14](#), [27](#)
- 19 call, under constraint, [16](#)
- 20 operation-related, [14](#)
- 21 semantics, [14](#)
- 22 specification, [10](#)
- 23 starting, [14](#)
- 24 synchronizing, [16](#)
- 25 procedure specification, [10](#)
- 26 process creation, [481](#)
- 27 process failures, [26](#)
- 28 process set names, [873](#)
- 29 processes, [25](#)
- 30 processor name, [444](#)
- 31 profiling interface, [717](#)
- 32 program error, [26](#)
- 33 progress, [27](#), [704](#), [1047](#)
- 34 cancellation, [104](#)
- 35 file, [704](#)
- 36 nonblocking collective, [705](#)
- 37 finalize, [490](#)
- 38 buffered data, [490](#)
- 39 session, [497](#)
- 40 generalized requests, [627](#)
- 41 guarantee, [28](#)
- 42 load/store access synchronization, [603](#),
- 43 [1046](#)
- 44 nonblocking collective communication,
- 45 [251](#), [252](#), [301](#), [303](#)
- 46 one-sided communication, [603](#), [615](#), [616](#),
- 47 [617](#)
- 48 partitioned point-to-point communication,
- [121](#)
- point-to-point communication, [55](#), [83](#), [91](#),
- [93](#), [96](#), [99](#), [1047](#)
- strong, [28](#), [1047](#)
- threads, [515](#)
- weak, [28](#), [1047](#)
- prototype definitions, [866](#)
- deprecated, [872](#)
- public window copy, [586](#)
- PVM, [2](#)
- raise an event, [749](#)
- rank, [308](#)
- ready mode send, [51](#), [54](#), [108](#), [109](#)
- as standard mode send, [54](#)
- nonblocking, [69](#), [71](#), [74](#)
- persistent, [108](#)
- receive, [31](#), [32](#)
- blocking, [36](#), [36](#)
- buffer, [32](#), [100](#), [109](#)
- complete, [69](#)
- context, [351](#)
- matched, [98](#), [100](#), [100](#)
- nonblocking, [69](#), [75](#)
- persistent, [109](#)
- start, [69](#), [75](#)
- reduce, [224](#)
- nonblocking, [267](#)
- persistent, [289](#)
- reduce-scatter, [243](#)
- nonblocking, [269](#), [270](#)
- persistent, [292](#), [293](#)
- reduction operations, [224](#), [842](#)
- predefined, [227](#)
- process-local, [242](#)
- scan, [247](#)
- user-defined, [235](#)
- registration handle, [756](#)
- related, [181](#)
- relative displacement, [21](#), [148](#)
- remote group, [322](#)
- Remote Memory Access
- see RMA, [547](#)
- removed interfaces, [22](#), [783](#)
- representation
- conversion, [49](#)
- request complete
- I/O, [654](#)
- request objects, [70](#)
- completion, [83](#)
- freeing, [81](#), [110](#)
- initiation, [82](#)
- multiple completions, [83](#)
- null handle, [78](#), [84](#)–[86](#), [92](#)
- null handles, [87](#), [91](#), [92](#)
- started, [50](#), [83](#), [99](#), [104](#), [109](#), [110](#)
- resource error, [26](#)
- return code
- tool information interface, [768](#)
- RMA, [547](#)

- communication calls, [566](#)
 - request-based, [580](#)
- dynamic window, [559](#)
- memory model, [586](#)
- synchronization calls, [587](#)
- root, [190](#)
- scan, [247](#)
 - inclusive, [247](#)
 - segmented, [249](#)
- scatter, [208](#)
 - nonblocking, [257](#)
 - persistent, [279](#)
- seek, [679](#)
- segmented scan, [249](#)
- semantic changes, [785](#)
- semantics, [11](#), [1050](#)
 - collective communications, [190](#)
 - nonblocking progress, [252](#), [301](#), [303](#)
 - data conversion, [48](#)
 - deadlock
 - buffer space, [57](#)
 - cyclic shift, [43](#)
 - send to self, [38](#)
 - erroneous program
 - freeing active request, [81](#)
 - invalid matched receive, [101](#), [102](#)
 - lack of buffer space, [56](#)
 - ready mode before receive, [51](#)
 - type matching, [46](#)
 - exceptions
 - completing and local, [52](#)
 - incomplete and nonlocal, [100](#)
 - file
 - collective, [705](#)
 - collective access, [677](#)
 - conflicting access, [702](#)
 - consistency, [701](#)
 - explicit offsets, [656](#)
 - nonblocking collective, [705](#)
 - progress, [704](#)
 - shared file pointer, [673](#)
 - split collective, [701](#)
 - finalize
 - buffered data, [490](#)
 - generalized requests
 - progress, [627](#)
 - inter-communicator, [322](#)
 - MPI_COMM_IDUP, [324](#)
 - nonblocking communication operations, [82](#)
 - nonblocking completion, [78](#)
 - nonblocking partitioned communications, [121](#)
 - operation, [11](#)
 - overlap, [69](#)
 - partitioned point-to-point communication, [114](#)
 - progress, [121](#)
 - point-to-point communication, [55](#)
 - buffered mode send, [67](#)
 - cancel, [80](#), [94](#), [94](#), [95](#), [102](#), [103](#), [104](#)
 - deterministic, [55](#)
 - fairness not guaranteed, [56](#)
 - fairness requirement, [89](#), [93](#)
 - logically concurrent, [55](#)
 - matched receive, [100](#)
 - matching probe, [97](#)
 - nonovertaking, [55](#), [82](#)
 - ordered, [55](#), [82](#)
 - overflow error on truncation, [37](#)
 - persistent, [104](#)
 - probe, [94](#)
 - progress, [55](#), [83](#), [91](#), [93](#), [96](#), [99](#), [1047](#)
 - resource limitations, [56](#)
 - send modes, [50](#), [54](#)
 - send-receive concurrency, [44](#)
 - wildcard receive, [38](#)
 - procedure, [14](#)
 - process failure, [543](#)
 - terms, [11](#)
 - threads
 - progress, [515](#)
 - semantics and correctness
 - one-sided communication, [606](#)
 - progress, [615](#), [616](#), [617](#)
 - send, [31](#)
 - blocking, [32](#), [32](#)
 - buffer, [31](#)
 - complete, [69](#)
 - context, [351](#)
 - nonblocking, [69](#)
 - persistent, [105](#)
 - start, [69](#), [72–74](#)
 - send-receive
 - blocking, [43](#), [43](#)
 - nonblocking, [77](#), [78](#)
 - start, [77](#), [78](#)
 - separate memory model, [547](#), [558](#), [586](#)
 - sequential storage, [162](#)
 - serialization, [69](#)
 - lack of buffer space, [58](#)
 - Sessions Model, [309](#), [312](#), [443](#), [481](#), [516](#), [640](#)
 - set
 - in function names, [9](#)
 - shared file pointers, [654](#), [673](#)
 - shared lock, [589](#), [597](#), [610](#)
 - shared memory, [553](#), [558](#), [600](#), [605](#), [609](#), [622](#)

- 1 allocation, [556](#)
- 2 contiguous, [553](#), [555](#), [556](#)
- 3 domain, [25](#), [333](#), [551](#), [553](#), [555](#), [557](#), [600](#)
- 4 loads/stores for synchronizing, [617](#), [1047](#)
- 5 noncontiguous, [556](#)
- 6 segment, [25](#), [333](#), [551](#), [553](#), [555](#), [556](#)
- 7 window, [554](#), [558](#), [623](#)
- 8 shared memory allocation, [554](#)
- 9 signals, [29](#)
- 10 singleton init, [541](#)
- 11 size changing
- 12 I/O, [706](#)
- 13 source, [351](#)
- 14 split collective, [653](#), [681](#)
- 15 stage, [11](#)
- 16 completion, [11](#)
- 17 freeing, [11](#)
- 18 initialization, [11](#)
- 19 starting, [11](#)
- 20 standard mode send, [50](#), [54](#), [105](#)
- 21 as synchronous mode send, [54](#)
- 22 nonblocking, [70](#), [72](#)
- 23 started, [491](#)
- 24 request objects, [50](#), [83](#), [99](#), [104](#), [109](#), [110](#)
- 25 started operation state, [12](#)
- 26 starting procedure, [14](#), [104](#)
- 27 starting processes, [518](#), [519](#)
- 28 starting stage, [11](#)
- 29 startup, [482](#)
- 30 portable, [512](#)
- 31 state
- 32 operation, [12](#)
- 33 active, [12](#)
- 34 complete, [12](#)
- 35 enabled, [13](#), [16](#), [79](#), [1047](#)
- 36 freed, [12](#)
- 37 inactive, [12](#)
- 38 initialized, [12](#)
- 39 started, [12](#)
- 40 type, [19](#)
- 41 status, [39](#), [838](#)
- 42 array in Fortran, [39](#)
- 43 associating information, [633](#)
- 44 derived type in Fortran 2008, [39](#)
- 45 empty, [79](#), [80](#), [84–87](#), [91](#), [92](#)
- 46 error in status, [39](#)
- 47 for send operation, [79](#)
- 48 ignore, [42](#)
- message length, [39](#)
- structure in C, [39](#)
- test, [90](#)
- strong progress, [28](#), [617](#), [1047](#)
- strong synchronization, [589](#)
- synchronism, [653](#)
- synchronization, [547](#), [566](#)
- synchronization calls
- RMA, [587](#)
- synchronizing procedure, [16](#)
- synchronous mode send, [51](#), [54](#), [78](#), [99](#), [107](#)
- nonblocking, [70](#), [71](#), [73](#)
- persistent, [107](#)
- system memory, [17](#)
- tag values, [443](#)
- target, [548](#)
- target nodes, [649](#)
- thread compliant, [486](#), [514](#)
- thread-safe, [515](#)
- threads, [514](#)
- progress, [515](#)
- thread-safe, [2](#), [442](#), [460](#), [461](#), [464](#), [515](#), [752](#)
- timers and synchronization, [470](#)
- tool information interface, [722](#)
- control variable, [729](#)
- event, [748](#)
- performance variable, [736](#)
- return code, [768](#)
- verbosity level, [724](#)
- tool support, [717](#)
- topologies, [383](#)
- topology
- Cartesian, [385](#), [386](#)
- distributed graph, [385](#), [390](#)
- graph, [385](#), [388](#)
- virtual, [383](#), [384](#)
- transmission failures, [26](#)
- true extent of datatypes, [155](#)
- TS 29113, [23](#), [787–791](#), [794](#), [797](#), [800–804](#),
[813–815](#), [820](#), [825](#), [829](#), [831](#), [1047](#),
[1059](#), [1060](#)
- type
- conversion, [49](#)
- matching rules, [46](#)
- type map, [128](#)
- type matching, [159](#)
- type signature, [128](#)
- types, [865](#)
- ub_marker, [142](#), [143](#), [146](#), [147](#), [152](#), [156](#)
- erased, [155](#)
- unified memory model, [547](#), [558](#), [586](#)
- universe size, [540](#)
- unnamed datatype, [16](#)
- unpack, [178](#)
- canonical, [185](#)
- upper bound, [152](#)
- upper-bound markers, [151](#)
- user functions at process termination, [492](#)

user-defined data representations, 694	1
user-defined reduction operations, 235	2
	3
verbosity level	4
tool information interface, 724	5
version inquiries, 441	6
view, 637 , 638 , 650	7
virtual topology, 306 , 383 , 384	8
	9
weak progress, 28 , 1047	10
weak synchronization, 589	11
wildcard, 38	12
window	13
allocation, 552	14
creation, 548	15
dynamically attached memory, 559	16
shared memory	17
window, 554	18
shared memory allocation, 554	19
World Model, 309 , 312 , 442 , 481 , 516 , 639 , 711	20
XDR, 49	21
	22
Zipcode, 2	23
	24
	25
	26
	27
	28
	29
	30
	31
	32
	33
	34
	35
	36
	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48

Examples Index

This index lists code examples throughout the text. Some examples are referred to by content; others are listed by the major MPI function that they are demonstrating. MPI functions listed in all capital letter are Fortran examples; MPI functions listed in mixed case are C examples.

- Accumulate in RMA
 - MPI_ACCUMULATE, 574
 - MPI_TYPE_GET_EXTENT, 574
 - MPI_WIN_CREATE, 574
 - MPI_WIN_FENCE, 574
 - MPI_WIN_FREE, 574
- Actions after Finalize
 - MPI_Finalize, 491
- Active target and local reads in RMA
 - MPI_Barrier, 613
 - MPI_Put, 613
 - MPI_Win_complete, 613
 - MPI_Win_lock, 613
 - MPI_Win_post, 613
 - MPI_Win_start, 613
 - MPI_Win_unlock, 613
 - MPI_Win_wait, 613
- Allgather
 - MPI_Allgather, 217
- Allreduce of a vector
 - MPI_ALLREDUCE, 241
- argv in C and Fortran
 - MPI_COMM_SPAWN, 522
 - MPI_Comm_spawn, 522
- Array of argv in C and Fortran
 - MPI_COMM_SPAWN_MULTIPLE, 527
 - MPI_Comm_spawn_multiple, 527
- ASYNCHRONOUS, 622, 824, 831
- Attach and detach buffer
 - MPI_Buffer_attach, 66
 - MPI_Buffer_detach, 66
- Attach and detach communicator-specific buffer
 - MPI_Comm_buffer_attach, 66
 - MPI_Comm_buffer_detach, 66
- Attributes between languages, 843
 - set in C
 - MPI_Comm_get_attr, 843
 - MPI_Comm_set_attr, 843
 - set in Fortran
 - MPI_ATTR_GET, 844, 845
 - MPI_ATTR_PUT, 844
 - MPI_COMM_GET_ATTR, 844, 845
 - MPI_COMM_SET_ATTR, 845
- Basic usage of performance variables
 - MPI_T_finalize, 746
 - MPI_T_init_thread, 746
 - MPI_T_pvar_get_info, 746
 - MPI_T_pvar_handle_alloc, 746
 - MPI_T_pvar_handle_free, 746
 - MPI_T_pvar_read, 746
 - MPI_T_pvar_session_create, 746
 - MPI_T_pvar_start, 746
- Blocking/Nonblocking collectives do not match, 301
- Broadcast
 - MPI_Bcast, 197
- C/Fortran handle conversion, 837
- C/Fortran handle conversion and absolute addresses
 - MPI_GET_ADDRESS, 841
- Cartesian virtual topologies, 437
- Client server code with waitsome
 - MPI_IRECV, 90
 - MPI_ISEND, 90
 - MPI_WAIT, 90
 - MPI_WAIT SOME, 90
- Client-server
 - MPI_Comm_accept, 538, 539
 - MPI_Comm_connect, 538, 539
 - MPI_IRECV, 89
 - MPI_ISEND, 89
 - MPI_Open_port, 538, 539
 - MPI_WAIT, 89
 - MPI_WAITANY, 89
- Client-server (with error)
 - MPI_PROBE, 96
 - MPI_RECV, 96
 - MPI_SEND, 96
- Client-server code, 89, 90
 - with probe, 96
 - with probe (wrong), 96
- Client-server model
 - MPI_Comm_remote_size, 330
 - MPI_Comm_split, 330

Client-server with probe	MPI_Win_flush_all, 621	1
MPI_PROBE, 96	MPI_Win_sync, 621	2
MPI_RECV, 96		3
MPI_SEND, 96	Datatype	4
Collective communication	3D array, 170	5
MPI_Bcast, 344	absolute addresses, 175	6
Communication safety	array of structures, 172	7
MPI_Comm_create, 345	elaborate example, 183, 184	8
MPI_Group_free, 345	matching type, 159	9
MPI_Group_incl, 345	matrix transpose, 171, 172	10
Consistency by setting atomic mode	union, 176	11
MPI_File_read_at, 707	Datatype matching	12
MPI_File_set_atomicity, 707	MPI_RECV, 47	13
MPI_File_set_view, 707	MPI_SEND, 47	14
MPI_File_write_at, 707	Datatypes	15
Consistency for writing and reading files	matching, 47	16
asynchronously	MPI_BYTE, 47	17
MPI_File, 709	MPI_RECV, 47	18
MPI_File_iread_at, 709	MPI_SEND, 47	19
MPI_File_iwrite_at, 709	not matching, 47	20
MPI_File_read_all_begin, 709	untyped, 47	21
MPI_File_read_all_end, 709	Datatypes for distributed arrays	22
MPI_File_set_atomicity, 709	MPI_TYPE_CREATE_DARRAY, 148	23
MPI_File_write_all_begin, 709	Deadlock	24
MPI_File_write_all_end, 709	if not buffered, 57	25
Consistency using “sync-barrier-sync”	with MPI_Bcast, 297, 298	26
MPI_File_read_at, 708	wrong message exchange, 57	27
MPI_File_set_view, 708	Deadlock due to synchronization through	28
MPI_File_sync, 708	shared memory	29
MPI_File_write_at, 708	MPI_Bsend, 617	30
Counting semaphore (non-scalable)	MPI_Buffer_attach, 617	31
MPI_Accumulate, 621	MPI_Buffer_detach, 617	32
MPI_Barrier, 621	MPI_Recv, 617	33
MPI_Get_accumulate, 621	MPI_Win_fence, 617	34
MPI_Win_flush, 621	MPI_Win_shared_query, 617	35
MPI_Win_sync, 621	Decoding a datatype	36
Creating a communicator using the Sessions	MPI_Type_get_contents, 177	37
Model	MPI_Type_get_envelope, 177	38
MPI_Comm_create_from_group, 501	Decoding amode in Fortran 77	39
MPI_Group_from_session_pset, 501	MPI_FILE_GET_AMODE, 646	40
MPI_Info_create, 501	Defining a user function	41
MPI_Info_set, 501	MPI_OP_CREATE, 240	42
MPI_Session_finalize, 501	MPI_REDUCE, 240	43
MPI_Session_init, 501	MPI_USER_FUNCTION, 240	44
Critical region with Compare-and-Swap	Dims create	45
MPI_Barrier, 622	MPI_DIMS_CREATE, 388	46
MPI_Compare_and_swap, 622	dist graph creation	47
MPI_Win_flush, 622	MPI_DIST_GRAPH_CREATE, 395	48
MPI_Win_sync, 622	MPI_DIST_GRAPH_CREATE_ADJACENT,	
Critical region with RMA	395	
MPI_Accumulate, 621	Dist_graph_create	
MPI_Barrier, 621	MPI_Dist_graph_create, 396	
MPI_Get_accumulate, 621	Double buffer in RMA	
MPI_Win_flush, 621	MPI_Barrier, 620	

- 1 MPI_Get, [620](#)
2 MPI_Win_complete, [620](#)
3 MPI_Win_post, [620](#)
4 MPI_Win_start, [620](#)
5 MPI_Win_wait, [620](#)
- 6 Errant message exchange
7 MPI_RECV, [57](#)
8 MPI_SEND, [57](#)
- 9 Erroneous attempt to achieve consistency
10 MPI_File_read_at, [708](#)
11 MPI_File_set_view, [708](#)
12 MPI_File_sync, [708](#)
13 MPI_File_write_at, [708](#)
- 14 Erroneous example fragment of concurrent
15 split collective access on a file handle
16 MPI_File_read_all, [682](#)
17 MPI_File_read_all_begin, [682](#)
18 MPI_File_read_all_end, [682](#)
- 19 Erroneous matching of blocking and
20 nonblocking collectives
21 MPI_Alltoall, [301](#)
22 MPI_lalltoall, [301](#)
23 MPI_Wait, [301](#)
- 24 Erroneous matching of collectives
25 MPI_Bcast, [300](#)
26 MPI_Ibarrier, [300](#)
27 MPI_Wait, [300](#)
- 28 Erroneous use of Bcast
29 MPI_Bcast, [297](#), [298](#)
- 30 Example 3.1: Hello world, [31](#)
31 Example 3.2: Datatype matching, [47](#)
32 Example 3.3: Datatype matching, [47](#)
33 Example 3.4: Datatypes, [47](#)
34 Example 3.5: Fortran CHARACTER, [48](#)
35 Example 3.6: Nonovertaking messages, [55](#)
36 Example 3.7: Message matching, [56](#)
37 Example 3.8: Message exchange, [57](#)
38 Example 3.9: Errant message exchange, [57](#)
39 Example 3.10: Exchange relies on buffering, [57](#)
40 Example 3.11: Attach and detach buffer, [66](#)
41 Example 3.12: Attach and detach
42 communicator-specific buffer, [66](#)
43 Example 3.13: Nonblocking point-to-point, [81](#)
44 Example 3.14: Nonblocking send and receive
45 with request free, [82](#)
46 Example 3.15: Message ordering for
47 nonblocking operations, [82](#)
48 Example 3.16: Progress semantics, [83](#)
49 Example 3.17: Client-server, [89](#)
50 Example 3.18: Client server code with
51 waitsome, [90](#)
52 Example 3.19: Client-server with probe, [96](#)
53 Example 3.20: Client-server (with error), [96](#)
- 54 Example 4.1: Partitioned communication, [114](#)
55 Example 4.2: Partitioned communication using
56 threads, [122](#)
57 Example 4.3: Partitioned communication with
58 tasks, [123](#)
59 Example 4.4: Partitioned communication with
60 partial completion, [125](#)
61 Example 5.1: Typemap, [129](#)
62 Example 5.2: Typemap for contiguous, [130](#)
63 Example 5.3: Typemap for vector, [131](#)
64 Example 5.4: Typemap for vector, [131](#)
65 Example 5.5: Typemap for indexed, [134](#)
66 Example 5.6: Typemap for create struct, [140](#)
67 Example 5.7: Datatypes for distributed arrays,
68 [148](#)
69 Example 5.8: Get_address, [149](#)
70 Example 5.9: Typemap of nested datatypes,
71 [152](#)
72 Example 5.10: Type_commit, [157](#)
73 Example 5.11: Matching type with datatypes,
74 [159](#)
75 Example 5.12: Using Get_count and
76 Get_elements, [161](#)
77 Example 5.13: Send/receive of a 3D array, [170](#)
78 Example 5.14: Using indexed datatype, [171](#)
79 Example 5.15: Nested vector datatypes, [171](#)
80 Example 5.16: Transpose with datatypes, [172](#)
81 Example 5.17: Using datatypes with array of
82 structures, [172](#)
83 Example 5.18: Using datatypes with array of
84 structures with absolute addresses,
85 [175](#)
86 Example 5.19: Using datatypes with unions,
87 [176](#)
88 Example 5.20: Decoding a datatype, [177](#)
89 Example 5.21: Using Pack, [183](#)
90 Example 5.22: Pack/Unpack with struct
91 datatype, [183](#)
92 Example 5.23: Pack and Pack_size, [184](#)
93 Example 6.1: Broadcast, [197](#)
94 Example 6.2: Gather, [202](#)
95 Example 6.3: Gather with allocation at the
96 root, [202](#)
97 Example 6.4: Gather with datatype, [202](#)
98 Example 6.5: Gatherv, [203](#)
99 Example 6.6: Gatherv with datatype, [203](#)
100 Example 6.7: Gatherv with datatype, [204](#)
101 Example 6.8: Gatherv with struct datatype,
102 [205](#)
103 Example 6.9: Gatherv with vector datatype,
104 [206](#)
105 Example 6.10: Gather and Gatherv, [207](#)
106 Example 6.11: Scatter, [212](#)
107 Example 6.12: Scatterv, [212](#)

- Example 6.13: Scatterv with vector datatype, [213](#)
- Example 6.14: Allgather, [217](#)
- Example 6.15: Reduction, [228](#)
- Example 6.16: Reduction of a vector, [229](#)
- Example 6.17: Retrieving an unnamed predefined value-index handle, [232](#)
- Example 6.18: Reduction with maxloc, [233](#)
- Example 6.19: Reduction with maxloc, [233](#)
- Example 6.20: Reduction with minloc, [234](#)
- Example 6.21: Reduction with user-defined op, [239](#)
- Example 6.22: Defining a user function, [240](#)
- Example 6.23: Allreduce of a vector, [241](#)
- Example 6.24: User-defined operation with Scan, [249](#)
- Example 6.25: Ibcast, [254](#)
- Example 6.26: Erroneous use of Bcast, [297](#)
- Example 6.27: Erroneous use of Bcast, [297](#)
- Example 6.28: Erroneous use of Bcast, [298](#)
- Example 6.29: Nondeterministic use of Bcast, [298](#)
- Example 6.30: Mixing blocking and nonblocking collective operations, [300](#)
- Example 6.31: Erroneous matching of collectives, [300](#)
- Example 6.32: Progress of nonblocking collectives, [301](#)
- Example 6.33: Erroneous matching of blocking and nonblocking collectives, [301](#)
- Example 6.34: Mixing collective and point-to-point, [302](#)
- Example 6.35: Pipelining nonblocking collectives, [302](#)
- Example 6.36: Overlapping communicators and collectives, [303](#)
- Example 6.37: Independence of nonblocking operations, [303](#)
- Example 7.1: Inter-communicator creation, [327](#)
- Example 7.2: Client-server model, [330](#)
- Example 7.3: Splitting into NUMANode subcommunicators, [334](#)
- Example 7.5: Recursive splitting of COMM_WORLD, [336](#)
- Example 7.6: Parallel output of a message, [343](#)
- Example 7.7: Message exchange, [343](#)
- Example 7.8: Collective communication, [344](#)
- Example 7.9: Using Group_excl, [344](#)
- Example 7.10: Communication safety, [345](#)
- Example 7.11: Library Example #1, [346](#)
- Example 7.12: Library Example #2, [347](#)
- Example 7.13: Three-Group “Pipeline”, [356](#)
- Example 7.14: Three-Group “Ring”, [358](#)
- Example 8.1: Dims create, [388](#)
- Example 8.2: Graph create, [389](#)
- Example 8.3: dist graph creation, [395](#)
- Example 8.4: Dist_graph_create, [396](#)
- Example 8.5: Graph creation and neighbors count, [402](#)
- Example 8.6: Graph creation, [402](#)
- Example 8.7: Using Cart_shift, [406](#)
- Example 8.8: Subgroup cart process topology, [407](#)
- Example 8.9: Neighbor allgather, [412](#)
- Example 8.10: Neighbor alltoall, [417](#)
- Example 8.11: Neighborhood collective communication, [437](#)
- Example 9.1: Use of Alloc_mem in Fortran, [448](#)
- Example 9.2: Use of Alloc_mem in Fortran with Cray pointers, [448](#)
- Example 9.3: Using Alloc_mem, [449](#)
- Example 9.4: Using MPI_Wtime, [470](#)
- Example 11.1: Initializing MPI, [483](#)
- Example 11.2: mpiexec and environment variables, [484](#)
- Example 11.3: Rules for finalize, [489](#)
- Example 11.4: Rules for finalize, [489](#)
- Example 11.5: Finalize and request free, [489](#)
- Example 11.6: Finalize and buffer attach, [490](#)
- Example 11.7: Finalize and cancel, [490](#)
- Example 11.8: Actions after Finalize, [491](#)
- Example 11.9: Finalize in the Sessions Model, [497](#)
- Example 11.10: Creating a communicator using the Sessions Model, [501](#)
- Example 11.11: Using process set query in group creation, [503](#)
- Example 11.12: Using process set query in group creation, [505](#)
- Example 11.13: Requesting and querying memory allocation kinds in the Sessions Model, [510](#)
- Example 11.14: Using mpiexec using -n, [513](#)
- Example 11.15: Using mpiexec using -host, [514](#)
- Example 11.16: Using mpiexec starting programs with separate argument lists, [514](#)
- Example 11.17: Using mpiexec using -arch, [514](#)
- Example 11.18: Using mpiexec using -configfile, [514](#)
- Example 11.19: Threads and MPI, [515](#)
- Example 11.20: argv in C and Fortran, [522](#)
- Example 11.21: Array of argv in C and Fortran, [527](#)

- 1 Example 11.22: Manager-worker with
2 Comm_spawn, [528](#)
- 3 Example 11.23: Client-server, [538](#)
- 4 Example 11.24: Name publishing, [538](#)
- 5 Example 11.25: Client-server, [539](#)
- 6 Example 12.1: Using Get with indexed
7 datatype, [570](#)
- 8 Example 12.2: Using Get, [572](#)
- 9 Example 12.3: Accumulate in RMA, [574](#)
- 10 Example 12.4: RMA Start and complete, [594](#)
- 11 Example 12.5: RMA Lock and unlock, [600](#)
- 12 Example 12.6: Update location in separate
13 memory model, [610](#)
- 14 Example 12.7: Update location in unified
15 memory model, [610](#)
- 16 Example 12.8: Read data in RMA, [611](#)
- 17 Example 12.9: Read data in RMA (unsafe), [611](#)
- 18 Example 12.10: Public and private memory in
19 RMA, [612](#)
- 20 Example 12.11: Public and private memory in
21 RMA, [612](#)
- 22 Example 12.12: Active target and local reads
23 in RMA, [613](#)
- 24 Example 12.13: Deadlock due to
25 synchronization through shared
26 memory, [617](#)
- 27 Example 12.14: Register and Compiler
28 Optimization, [618](#)
- 29 Example 12.15: Put with Fence, [618](#)
- 30 Example 12.16: Get with fence, [619](#)
- 31 Example 12.17: Put with PSCW, [619](#)
- 32 Example 12.18: Get with PSCW, [619](#)
- 33 Example 12.19: Double buffer in RMA, [620](#)
- 34 Example 12.20: Counting semaphore
35 (nonscalable), [621](#)
- 36 Example 12.21: Critical region with RMA, [621](#)
- 37 Example 12.22: Critical region with
38 Compare-and-Swap, [622](#)
- 39 Example 12.23: Shared memory windows, [623](#)
- 40 Example 12.24: Requests in RMA, [623](#)
- 41 Example 12.25: Linked list in RMA, [624](#)
- 42 Example 13.1: User-defined reduce, [631](#)
- 43 Example 14.1: Decoding amode in Fortran 77,
44 [646](#)
- 45 Example 14.2: Read to end of file, [665](#)
- 46 Example 14.3: File pointer update semantics,
47 [668](#)
- 48 Example 14.4: Erroneous example fragment of
concurrent split collective access on a
file handle, [682](#)
- Example 14.5: Consistency by setting atomic
mode, [707](#)
- Example 14.6: Consistency using
“sync-barrier-sync”, [708](#)
- Example 14.7: Erroneous attempt to achieve
consistency, [708](#)
- Example 14.8: Consistency for writing and
reading files asynchronously, [709](#)
- Example 14.9: Overlap computation and
output, [712](#)
- Example 14.10: Local subarray for file output,
[714](#)
- Example 14.11: Local subarray for file output
in Fortran 90, [715](#)
- Example 15.1: Measurement wrapper, [720](#)
- Example 15.2: Profiling interface
implementation using weak symbols, [720](#)
- Example 15.3: Profiling interface
implementation using the C macro
preprocessor, [720](#)
- Example 15.5: Listing names of control
variables, [732](#)
- Example 15.6: Reading a control variable, [736](#)
- Example 15.7: Basic usage of performance
variables, [746](#)
- Example 19.1: Fortran 90
selected KIND, [808](#)
- Example 19.2: MPI_TYPE_MATCH_SIZE
implementation, [811](#)
- Example 19.3: Fortran 90
heterogeneous communication (unsafe),
[812](#)
- Example 19.4: Fortran 90
heterogeneous MPI I/O (unsafe), [812](#)
- Example 19.5: Fortran 90
a scalar is not an array, [815](#)
- Example 19.6: Fortran 90
subarray as buffer, [815](#)
- Example 19.7: Fortran 90
subarray as buffer, [816](#)
- Example 19.8: Fortran 90
subarray as buffer, [817](#)
- Example 19.9: Fortran 90
scalars as buffer, [818](#)
- Example 19.10: Fortran 90
vector subscripts, [818](#)
- Example 19.11: Fortran 90
derived types, [819](#)
- Example 19.12: Fortran 90 register
optimization, [822](#)
- Example 19.13: Fortran 90 register
optimization, [822](#)
- Example 19.14: Fortran 90 register
optimization, [823](#)
- Example 19.15: Fortran 90 register
optimization, [823](#)
- Example 19.16: Protecting nonblocking
communication with

- ASYNCHRONOUS, [824](#)
- Example 19.17: Fortran 90 register optimization in RMA, [827](#)
- Example 19.18: Fortran 90 overlapping communication and computation, [829](#)
- Example 19.19: Fortran 90 overlapping communication and computation, [829](#)
- Example 19.20: Fortran 90 overlapping communication and computation, [829](#)
- Example 19.21: Using separated variables, [831](#)
- Example 19.22: Fortran 90 overlapping communication and computation, [832](#)
- Example 19.23: C/Fortran handle conversion, [837](#)
- Example 19.24: C/Fortran handle conversion and absolute addresses, [841](#)
- Example 19.25: Attributes between languages set in C, [843](#)
- Example 19.26: Attributes between languages set in Fortran, [844](#)
- Example 19.27: Attributes between languages set in Fortran, [845](#)
- Example 19.28: Interlanguage communication, [847](#)
- Exchange relies on buffering
 - [MPI_RECV](#), [57](#)
 - [MPI_SEND](#), [57](#)
- False matching of collective operations, [300](#)
- File pointer update semantics
 - [MPI_FILE_CLOSE](#), [668](#)
 - [MPI_FILE_IREAD](#), [668](#)
 - [MPI_FILE_OPEN](#), [668](#)
 - [MPI_FILE_SET_VIEW](#), [668](#)
 - [MPI_WAIT](#), [668](#)
- Finalize and buffer attach
 - [MPI_Buffer_attach](#), [490](#)
 - [MPI_Finalize](#), [490](#)
- Finalize and cancel
 - [MPI_Barrier](#), [490](#)
 - [MPI_Cancel](#), [490](#)
 - [MPI_Finalize](#), [490](#)
 - [MPI_Iprobe](#), [490](#)
 - [MPI_Test_cancelled](#), [490](#)
- Finalize and request free
 - [MPI_Finalize](#), [489](#)
 - [MPI_Request_free](#), [489](#)
- Finalize in the Sessions Model
 - [MPI_SESSION_FINALIZE](#), [497](#)
- Fortran 90
 - a scalar is not an array
 - [MPI_CART_CREATE](#), [815](#)
 - copying and sequence problem, [815–818](#)
 - derived types
 - [MPI_GET_ADDRESS](#), [819](#)
 - [MPI_TYPE_COMMIT](#), [819](#)
 - [MPI_TYPE_CREATE_RESIZED](#), [819](#)
 - [MPI_TYPE_CREATE_STRUCT](#), [819](#)
 - heterogeneous MPI I/O (unsafe)
 - [MPI_FILE_SET_VIEW](#), [812](#)
 - [MPI_TYPE_MATCH_SIZE](#), [812](#)
 - heterogeneous communication (unsafe)
 - [MPI_TYPE_MATCH_SIZE](#), [812](#)
 - overlapping communication and computation, [829](#), [832](#)
 - register optimization, [822](#), [823](#)
 - scalars as buffer
 - [MPI_IRECV](#), [818](#)
 - selected KIND
 - [MPI_TYPE_CREATE_F90_INTEGER](#), [808](#)
 - [MPI_TYPE_CREATE_F90_REAL](#), [808](#)
 - subarray as buffer
 - [MPI_IRECV](#), [816](#), [817](#)
 - [MPI_ISEND](#), [816](#)
 - [MPI_SEND](#), [815](#)
 - vector subscripts
 - [MPI_SEND](#), [818](#)
- Fortran 90 register optimization in RMA
 - [MPI_F_SYNC_REG](#), [827](#)
 - [MPI_PUT](#), [827](#)
 - [MPI_WIN_FENCE](#), [827](#)
- Fortran CHARACTER
 - [MPI_CHARACTER](#), [48](#)
 - [MPI_RECV](#), [48](#)
 - [MPI_SEND](#), [48](#)
- Gather
 - [MPI_Gather](#), [202](#)
- Gather and Gatherv
 - [MPI_Gather](#), [207](#)
 - [MPI_Gatherv](#), [207](#)
 - [MPI_Type_commit](#), [207](#)
 - [MPI_Type_create_struct](#), [207](#)
- Gather with allocation at the root
 - [MPI_Gather](#), [202](#)
- Gather with datatype
 - [MPI_Gather](#), [202](#)
 - [MPI_Type_commit](#), [202](#)
 - [MPI_Type_contiguous](#), [202](#)
- Gatherv
 - [MPI_Gatherv](#), [203](#)
- Gatherv with datatype
 - [MPI_Gatherv](#), [203](#), [204](#)
 - [MPI_Type_commit](#), [203](#), [204](#)
 - [MPI_Type_vector](#), [203](#), [204](#)
- Gatherv with struct datatype
 - [MPI_Gatherv](#), [205](#)

- 1 MPI_Type_commit, 205
- 2 MPI_Type_create_struct, 205
- 3 Gather with vector datatype
- 4 MPI_Gatherv, 206
- 5 MPI_Type_commit, 206
- 6 MPI_Type_vector, 206
- 7 Get with fence
- 8 MPI_Get, 619
- 9 MPI_Win_fence, 619
- 10 Get with PSCW
- 11 MPI_Get, 619
- 12 MPI_Win_complete, 619
- 13 MPI_Win_post, 619
- 14 MPI_Win_start, 619
- 15 MPI_Win_wait, 619
- 16 Get_address
- 17 MPI_GET_ADDRESS, 149
- 18 Graph create
- 19 MPI_GRAPH_CREATE, 389
- 20 Graph creation
- 21 MPI_GRAPH_CREATE, 402
- 22 Graph creation and neighbors count
- 23 MPI_GRAPH_CREATE, 402
- 24 MPI_GRAPH_NEIGHBORS, 402
- 25 MPI_GRAPH_NEIGHBORS_COUNT, 402
- 26 Hello world
- 27 MPI_Comm_rank, 31
- 28 MPI_Init, 31
- 29 MPI_Recv, 31
- 30 MPI_Send, 31
- 31 Ibcast
- 32 MPI_Ibcast, 254
- 33 Independence of nonblocking operations, 303
- 34 MPI_Ibcast, 303
- 35 Initializing MPI
- 36 MPI_Init, 483
- 37 Inter-communicator, 327, 330
- 38 Inter-communicator creation
- 39 MPI_Comm_create, 327
- 40 MPI_Comm_group, 327
- 41 MPI_Group_free, 327
- 42 MPI_Group_incl, 327
- 43 Interlanguage communication, 847
- 44 MPI_GET_ADDRESS, 847
- 45 MPI_TYPE_CREATE_STRUCT, 847
- 46 Intertwined matching pairs, 56
- 47 Library Example #1
- 48 MPI_Comm_dup, 346
- 49 MPI_Reduce, 346
- 50 Library Example #2
- 51 MPI_Comm_create, 347
- 52 MPI_Comm_group, 347
- 53 MPI_Group_free, 347
- 54 MPI_Group_incl, 347
- 55 Linked list in RMA
- 56 MPI_Accumulate, 624
- 57 MPI_Aint_add, 624
- 58 MPI_Alloc_mem, 624
- 59 MPI_Compare_and_swap, 624
- 60 MPI_Free_mem, 624
- 61 MPI_Get_accumulate, 624
- 62 MPI_Win_attach, 624
- 63 MPI_Win_create_dynamic, 624
- 64 MPI_Win_detach, 624
- 65 MPI_Win_flush, 624
- 66 MPI_Win_lock_all, 624
- 67 MPI_Win_unlock_all, 624
- 68 Linking libraries when using the profiling interface, 721
- 69 Listing names of control variables
- 70 MPI_T_cvar_get_info, 732
- 71 Local subarray for file output
- 72 MPI_Type_create_subarray, 714
- 73 Local subarray for file output in Fortran 90
- 74 MPI_TYPE_CREATE_SUBARRAY, 715
- 75 Manager-worker with Comm_spawn
- 76 MPI_Comm_get_parent, 528
- 77 MPI_Comm_spawn, 528
- 78 Matching type with datatypes
- 79 MPI_RECV, 159
- 80 MPI_SEND, 159
- 81 MPI_TYPE_CONTIGUOUS, 159
- 82 Measurement wrapper
- 83 MPI_Send, 720
- 84 MPI_Type_size, 720
- 85 MPI_Wtime, 720
- 86 Message exchange
- 87 MPI_RECV, 57
- 88 MPI_SEND, 57
- 89 Message exchange (ping-pong), 57
- 90 Message matching
- 91 MPI_BSEND, 56
- 92 MPI_RECV, 56
- 93 MPI_SSEND, 56
- 94 Message ordering for nonblocking operations
- 95 MPI_IRECV, 82
- 96 MPI_ISEND, 82
- 97 MPI_WAIT, 82
- 98 Mixing blocking and nonblocking collective operations, 300
- 99 MPI_Bcast, 300
- 100 MPI_lbarrier, 300
- 101 MPI_Wait, 300
- 102 Mixing collective and point-to-point
- 103 MPI_lbarrier, 302

MPI_Irecv, 302	1
MPI_Send, 302	2
MPI_Wait, 302	3
MPI_Waitall, 302	4
Mixing collective and point-to-point requests, 302	5
MPI_ACCUMULATE	6
Accumulate in RMA, 574	7
MPI_Accumulate	8
Counting semaphore (non-scalable), 621	9
Critical region with RMA, 621	10
Linked list in RMA, 624	11
MPI_Aint	12
Using datatypes with array of structures, 172	13
MPI_Aint_add	14
Linked list in RMA, 624	15
MPI_Allgather	16
Allgather, 217	17
MPI_ALLOC_MEM	18
Use of Alloc_mem in Fortran, 448	19
Use of Alloc_mem in Fortran with Cray pointers, 448	20
MPI_Alloc_mem	21
Linked list in RMA, 624	22
Using Alloc_mem, 449	23
MPI_ALLREDUCE	24
Allreduce of a vector, 241	25
MPI_Alltoall	26
Erroneous matching of blocking and nonblocking collectives, 301	27
mpi_assert_memory_alloc_kinds	28
Requesting and querying memory allocation kinds in the Sessions Model, 510	29
MPI_ASYNC_PROTECTS_NONBLOCKING, 622	30
MPI_ATTR_GET	31
Attributes between languages set in Fortran, 844, 845	32
MPI_ATTR_PUT	33
Attributes between languages set in Fortran, 844	34
MPI_BARRIER	35
Using process set query in group creation, 505	36
MPI_Barrier	37
Active target and local reads in RMA, 613	38
Counting semaphore (non-scalable), 621	39
Critical region with Compare-and-Swap, 622	40
Critical region with RMA, 621	41
Double buffer in RMA, 620	42
Finalize and cancel, 490	43
Public and private memory in RMA, 612	44
Read data in RMA, 611	45
Read data in RMA (unsafe), 611	46
Update location in separate memory model, 610	47
Update location in unified memory model, 610	48
MPI_Bcast	1
Broadcast, 197	2
Collective communication, 344	3
Erroneous matching of collectives, 300	4
Erroneous use of Bcast, 297, 298	5
Mixing blocking and nonblocking collective operations, 300	6
Nondeterministic use of Bcast, 298	7
MPI_BSEND	8
Message matching, 56	9
Nonovertaking messages, 55	10
MPI_Bsend	11
Deadlock due to synchronization through shared memory, 617	12
MPI_Buffer_attach	13
Attach and detach buffer, 66	14
Deadlock due to synchronization through shared memory, 617	15
Finalize and buffer attach, 490	16
MPI_Buffer_detach	17
Attach and detach buffer, 66	18
Deadlock due to synchronization through shared memory, 617	19
MPI_BYTE	20
Datatypes, 47	21
MPI_Cancel	22
Finalize and cancel, 490	23
MPI_CART_COORDS	24
Using Cart_shift, 406	25
MPI_CART_CREATE	26
Fortran 90	27
a scalar is not an array, 815	28
MPI_CART_GET	29
Neighborhood collective communication, 437	30
MPI_CART_RANK	31
Using Cart_shift, 406	32
MPI_CART_SHIFT	33
Neighbor alltoall, 417	34
Neighborhood collective communication, 437	35
Using Cart_shift, 406	36
MPI_CART_SUB	37
Subgroup cart process topology, 407	38
MPI_CARTDIM_GET	39
Neighbor alltoall, 417	40
MPI_CHARACTER	41

- 1 Fortran CHARACTER, [48](#)
- 2 MPI_Comm_accept
- 3 Client-server, [538](#), [539](#)
- 4 Name publishing, [538](#)
- 5 MPI_Comm_buffer_attach
- 6 Attach and detach communicator-specific
- 7 buffer, [66](#)
- 8 MPI_Comm_buffer_detach
- 9 Attach and detach communicator-specific
- 10 buffer, [66](#)
- 11 MPI_Comm_connect
- 12 Client-server, [538](#), [539](#)
- 13 Name publishing, [538](#)
- 14 MPI_Comm_create
- 15 Communication safety, [345](#)
- 16 Inter-communicator creation, [327](#)
- 17 Library Example #2, [347](#)
- 18 Using Group_excl, [344](#)
- 19 MPI_COMM_CREATE_FROM_GROUP
- 20 Using process set query in group creation,
- 21 [505](#)
- 22 MPI_Comm_create_from_group
- 23 Creating a communicator using the
- 24 Sessions Model, [501](#)
- 25 MPI_Comm_dup
- 26 Library Example #1, [346](#)
- 27 MPI_COMM_GET_ATTR
- 28 Attributes between languages
- 29 set in Fortran, [844](#), [845](#)
- 30 MPI_Comm_get_attr
- 31 Attributes between languages
- 32 set in C, [843](#)
- 33 MPI_Comm_get_parent
- 34 Manager-worker with Comm_spawn, [528](#)
- 35 MPI_Comm_group
- 36 Inter-communicator creation, [327](#)
- 37 Library Example #2, [347](#)
- 38 MPI_Comm_rank
- 39 Hello world, [31](#)
- 40 MPI_Comm_remote_size
- 41 Client-server model, [330](#)
- 42 MPI_COMM_SET_ATTR
- 43 Attributes between languages
- 44 set in Fortran, [845](#)
- 45 MPI_Comm_set_attr
- 46 Attributes between languages
- 47 set in C, [843](#)
- 48 MPI_COMM_SPAWN
- argv in C and Fortran, [522](#)
- MPI_Comm_spawn
- argv in C and Fortran, [522](#)
- Manager-worker with Comm_spawn, [528](#)
- MPI_COMM_SPAWN_MULTIPLE
- Array of argv in C and Fortran, [527](#)
- MPI_Comm_spawn_multiple
- Array of argv in C and Fortran, [527](#)
- MPI_Comm_split
- Client-server model, [330](#)
- Three-Group “Pipeline”, [356](#)
- Three-Group “Ring”, [358](#)
- MPI_Comm_split_type
- Recursive splitting of COMM_WORLD,
- [336](#)
- Splitting into NUMANode
- subcommunicators, [334](#)
- MPI_Compare_and_swap
- Critical region with Compare-and-Swap,
- [622](#)
- Linked list in RMA, [624](#)
- MPI_DIMS_CREATE
- Dims create, [388](#)
- Neighborhood collective communication,
- [437](#)
- MPI_DIST_GRAPH_CREATE
- dist graph creation, [395](#)
- MPI_Dist_graph_create
- Dist_graph_create, [396](#)
- MPI_DIST_GRAPH_CREATE_ADJACENT
- dist graph creation, [395](#)
- MPI_F_SYNC_REG
- Fortran 90 register optimization in RMA,
- [827](#)
- Shared memory windows, [623](#)
- MPI_File
- Consistency for writing and reading files
- asynchronously, [709](#)
- MPI_FILE_CLOSE
- File pointer update semantics, [668](#)
- Read to end of file, [665](#)
- MPI_FILE_GET_AMODE
- Decoding amode in Fortran 77, [646](#)
- MPI_FILE_IREAD
- File pointer update semantics, [668](#)
- MPI_File_iread_at
- Consistency for writing and reading files
- asynchronously, [709](#)
- MPI_File_iwrite_at
- Consistency for writing and reading files
- asynchronously, [709](#)
- MPI_FILE_OPEN
- File pointer update semantics, [668](#)
- Read to end of file, [665](#)
- MPI_FILE_READ
- Read to end of file, [665](#)
- MPI_File_read_all
- Erroneous example fragment of concurrent
- split collective access on a file handle,

	682		
MPI_File_read_all_begin	Consistency for writing and reading files asynchronously, 709		
	Erroneous example fragment of concurrent split collective access on a file handle, 682		
MPI_File_read_all_end	Consistency for writing and reading files asynchronously, 709		
	Erroneous example fragment of concurrent split collective access on a file handle, 682		
MPI_File_read_at	Consistency by setting atomic mode, 707		
	Consistency using “sync-barrier-sync”, 708		
	Erroneous attempt to achieve consistency, 708		
MPI_File_set_atomics	Consistency by setting atomic mode, 707		
	Consistency for writing and reading files asynchronously, 709		
MPI_FILE_SET_VIEW	File pointer update semantics, 668		
	Fortran 90		
	heterogeneous MPI I/O (unsafe), 812		
	Read to end of file, 665		
MPI_File_set_view	Consistency by setting atomic mode, 707		
	Consistency using “sync-barrier-sync”, 708		
	Erroneous attempt to achieve consistency, 708		
MPI_File_sync	Consistency using “sync-barrier-sync”, 708		
	Erroneous attempt to achieve consistency, 708		
MPI_File_write_all_begin	Consistency for writing and reading files asynchronously, 709		
	Overlap computation and output, 712		
MPI_File_write_all_end	Consistency for writing and reading files asynchronously, 709		
	Overlap computation and output, 712		
MPI_File_write_at	Consistency by setting atomic mode, 707		
	Consistency using “sync-barrier-sync”, 708		
	Erroneous attempt to achieve consistency, 708		
MPI_Finalize	Actions after Finalize, 491		
	Finalize and buffer attach, 490		
	Finalize and cancel, 490		
	Finalize and request free, 489		
	Rules for finalize, 489		1
MPI_FREE_MEM	Use of Alloc_mem in Fortran, 448		2
	Use of Alloc_mem in Fortran with Cray pointers, 448		3
MPI_Free_mem	Linked list in RMA, 624		4
MPI_Gather	Gather, 202		5
	Gather and Gatherv, 207		6
	Gather with allocation at the root, 202		7
	Gather with datatype, 202		8
	Pack and Pack_size, 184		9
MPI_Gatherv	Gather and Gatherv, 207		10
	Gatherv, 203		11
	Gatherv with datatype, 203, 204		12
	Gatherv with struct datatype, 205		13
	Gatherv with vector datatype, 206		14
	Pack and Pack_size, 184		15
MPI_GET	Using Get, 572		16
	Using Get with indexed datatype, 570		17
MPI_Get	Double buffer in RMA, 620		18
	Get with fence, 619		19
	Get with PSCW, 619		20
	Public and private memory in RMA, 612		21
	Update location in separate memory model, 610		22
	Update location in unified memory model, 610		23
MPI_Get_accumulate	Counting semaphore (nonscalable), 621		24
	Critical region with RMA, 621		25
	Linked list in RMA, 624		26
MPI_GET_ADDRESS	C/Fortran handle conversion and absolute addresses, 841		27
	Fortran 90		28
	derived types, 819		29
	Get_address, 149		30
	Interlanguage communication, 847		31
MPI_Get_address	Pack/Unpack with struct datatype, 183		32
	Using datatypes with array of structures, 172		33
	Using datatypes with array of structures with absolute addresses, 175		34
	Using datatypes with unions, 176		35
MPI_GET_COUNT	Using Get_count and Get_elements, 161		36
MPI_GET_ELEMENTS	Using Get_count and Get_elements, 161		37
			38
			39
			40
			41
			42
			43
			44
			45
			46
			47
			48

- 1 MPI_GRAPH_CREATE
 2 Graph create, 389
 3 Graph creation, 402
 4 Graph creation and neighbors count, 402
 5 MPI_GRAPH_NEIGHBORS
 6 Graph creation and neighbors count, 402
 7 MPI_GRAPH_NEIGHBORS_COUNT
 8 Graph creation and neighbors count, 402
 9 MPI_Grequest_complete
 10 User-defined reduce, 631
 11 MPI_Grequest_start
 12 User-defined reduce, 631
 13 MPI_Group_excl
 14 Using Group_excl, 344
 15 MPI_Group_free
 16 Communication safety, 345
 17 Inter-communicator creation, 327
 18 Library Example #2, 347
 19 Using Group_excl, 344
 20 MPI_GROUP_FROM_SESSION_PSET
 21 Using process set query in group creation,
 22 505
 23 MPI_Group_from_session_pset
 24 Creating a communicator using the
 25 Sessions Model, 501
 26 Using process set query in group creation,
 27 503
 28 MPI_Group_incl
 29 Communication safety, 345
 30 Inter-communicator creation, 327
 31 Library Example #2, 347
 32 MPI_lallreduce
 33 Overlapping communicators and
 34 collectives, 303
 35 MPI_lalltoall
 36 Erroneous matching of blocking and
 37 nonblocking collectives, 301
 38 MPI_lbarrier
 39 Erroneous matching of collectives, 300
 40 Mixing blocking and nonblocking
 41 collective operations, 300
 42 Mixing collective and point-to-point, 302
 43 Progress of nonblocking collectives, 301
 44 MPI_lbroadcast
 45 Ibroadcast, 254
 46 Independence of nonblocking operations,
 47 303
 48 Pipelining nonblocking collectives, 302
 49 MPI_Info_create
 50 Creating a communicator using the
 51 Sessions Model, 501
 52 MPI_INFO_ENV
 53 mpixexec and environment variables, 484
 54 MPI_Info_set
 55 Creating a communicator using the
 56 Sessions Model, 501
 57 MPI_Init
 58 Hello world, 31
 59 Initializing MPI, 483
 60 MPI_Intercomm_create
 61 Three-Group “Pipeline”, 356
 62 Three-Group “Ring”, 358
 63 MPI_lprobe
 64 Finalize and cancel, 490
 65 MPI_Irecv
 66 Client server code with waitsome, 90
 67 Client-server, 89
 68 Fortran 90
 69 scalars as buffer, 818
 70 subarray as buffer, 816, 817
 71 Message ordering for nonblocking
 72 operations, 82
 73 Nonblocking point-to-point, 81
 74 Nonblocking send and receive with request
 75 free, 82
 76 Progress semantics, 83
 77 MPI_irecv
 78 Mixing collective and point-to-point, 302
 79 MPI_ISEND
 80 Client server code with waitsome, 90
 81 Client-server, 89
 82 Fortran 90
 83 subarray as buffer, 816
 84 Message ordering for nonblocking
 85 operations, 82
 86 Nonblocking point-to-point, 81
 87 Nonblocking send and receive with request
 88 free, 82
 89 mpi_memory_alloc_kinds
 90 Requesting and querying memory
 91 allocation kinds in the Sessions
 92 Model, 510
 93 MPI_NEIGHBOR_ALLGATHER
 94 Neighbor allgather, 412
 95 MPI_NEIGHBOR_ALLTOALL
 96 Neighbor alltoall, 417
 97 MPI_Neighbor_alltoall, 417
 98 MPI_OP_CREATE
 99 Defining a user function, 240
 100 MPI_Op_create
 101 Reduction with user-defined op, 239
 102 User-defined operation with Scan, 249
 103 MPI_Open_port
 104 Client-server, 538, 539
 105 Name publishing, 538
 106 MPI_Pack
 107 Pack and Pack_size, 184
 108 Pack/Unpack with struct datatype, 183

- Using Pack, [183](#)
- MPI_Pack_size
 - Pack and Pack_size, [184](#)
- MPI_Parrived
 - Partitioned communication with partial completion, [125](#)
- MPI_Pready
 - Partitioned communication, [114](#)
 - Partitioned communication using threads, [122](#)
 - Partitioned communication with partial completion, [125](#)
 - Partitioned communication with tasks, [123](#)
- MPI_Precv_init
 - Partitioned communication, [114](#)
 - Partitioned communication using threads, [122](#)
 - Partitioned communication with partial completion, [125](#)
 - Partitioned communication with tasks, [123](#)
- MPI_PROBE
 - Client-server (with error), [96](#)
 - Client-server with probe, [96](#)
- MPI_Psend_init
 - Partitioned communication, [114](#)
 - Partitioned communication using threads, [122](#)
 - Partitioned communication with partial completion, [125](#)
 - Partitioned communication with tasks, [123](#)
- MPI_Publish_name
 - Name publishing, [538](#)
- MPI_PUT
 - Fortran 90 register optimization in RMA, [827](#)
- MPI_Put
 - Active target and local reads in RMA, [613](#)
 - Put with Fence, [618](#)
 - Put with PSCW, [619](#)
 - Read data in RMA, [611](#)
 - Read data in RMA (unsafe), [611](#)
 - Register and Compiler Optimization, [618](#)
 - RMA Lock and unlock, [600](#)
 - RMA Start and complete, [594](#)
- MPI_RECV
 - Client-server (with error), [96](#)
 - Client-server with probe, [96](#)
 - Datatype matching, [47](#)
 - Datatypes, [47](#)
 - Errant message exchange, [57](#)
 - Exchange relies on buffering, [57](#)
- Fortran CHARACTER, [48](#)
- Matching type with datatypes, [159](#)
- Message exchange, [57](#)
- Message matching, [56](#)
- Nonovertaking messages, [55](#)
- Progress semantics, [83](#)
- MPI_Recv
 - Deadlock due to synchronization through shared memory, [617](#)
 - Hello world, [31](#)
 - Progress of nonblocking collectives, [301](#)
- MPI_REDUCE
 - Defining a user function, [240](#)
 - Reduction, [228](#)
 - Reduction of a vector, [229](#)
 - Reduction with maxloc, [233](#)
- MPI_Reduce
 - Library Example #1, [346](#)
 - Reduction with maxloc, [233](#)
 - Reduction with minloc, [234](#)
 - Reduction with user-defined op, [239](#)
- MPI_REQUEST_FREE
 - Nonblocking send and receive with request free, [82](#)
- MPI_Request_free
 - Finalize and request free, [489](#)
- MPI_Rget
 - Requests in RMA, [623](#)
- MPI_Rput
 - Requests in RMA, [623](#)
- MPI_Scan
 - User-defined operation with Scan, [249](#)
- MPI_Scatter
 - Scatter, [212](#)
- MPI_Scatterv
 - Scatterv, [212](#)
 - Scatterv with vector datatype, [213](#)
- MPI_SEND
 - Client-server (with error), [96](#)
 - Client-server with probe, [96](#)
 - Datatype matching, [47](#)
 - Datatypes, [47](#)
 - Errant message exchange, [57](#)
 - Exchange relies on buffering, [57](#)
 - Fortran 90
 - subarray as buffer, [815](#)
 - vector subscripts, [818](#)
 - Fortran CHARACTER, [48](#)
 - Matching type with datatypes, [159](#)
 - Message exchange, [57](#)
 - Progress semantics, [83](#)
- MPI_Send
 - Hello world, [31](#)
 - Measurement wrapper, [720](#)

- 1 Mixing collective and point-to-point, [302](#)
- 2 Pack/Unpack with struct datatype, [183](#)
- 3 Progress of nonblocking collectives, [301](#)
- 4 Using datatypes with array of structures,
5 [172](#)
- 6 Using datatypes with array of structures
7 with absolute addresses, [175](#)
- 8 Using datatypes with unions, [176](#)
- 9 MPI_SENDRECV
- 10 Neighbor alltoall, [417](#)
- 11 Nested vector datatypes, [171](#)
- 12 Send/receive of a 3D array, [170](#)
- 13 Transpose with datatypes, [172](#)
- 14 Using indexed datatype, [171](#)
- 15 MPI_SENDRECV_REPLACE
- 16 Using Cart_shift, [406](#)
- 17 MPI_SESSION_FINALIZE
- 18 Finalize in the Sessions Model, [497](#)
- 19 Using process set query in group creation,
20 [505](#)
- 21 MPI_Session_finalize
- 22 Creating a communicator using the
23 Sessions Model, [501](#)
- 24 Using process set query in group creation,
25 [503](#)
- 26 MPI_SESSION_GET_NTH_PSET
- 27 Using process set query in group creation,
28 [505](#)
- 29 MPI_Session_get_nth_pset
- 30 Using process set query in group creation,
31 [503](#)
- 32 MPI_SESSION_GET_NUM_PSETS
- 33 Using process set query in group creation,
34 [505](#)
- 35 MPI_Session_get_num_psets
- 36 Using process set query in group creation,
37 [503](#)
- 38 MPI_SESSION_INIT
- 39 Using process set query in group creation,
40 [505](#)
- 41 MPI_Session_init
- 42 Creating a communicator using the
43 Sessions Model, [501](#)
- 44 Using process set query in group creation,
45 [503](#)
- 46 MPI_SSEND
- 47 Message matching, [56](#)
- 48 Progress semantics, [83](#)
- MPI_T_cvar_get_info
- Listing names of control variables, [732](#)
- MPI_T_cvar_handle_alloc
- Reading a control variable, [736](#)
- MPI_T_cvar_handle_free
- Reading a control variable, [736](#)
- MPI_T_cvar_read
- Reading a control variable, [736](#)
- MPI_T_finalize
- Basic usage of performance variables, [746](#)
- MPI_T_init_thread
- Basic usage of performance variables, [746](#)
- MPI_T_pvar_get_info
- Basic usage of performance variables, [746](#)
- MPI_T_pvar_handle_alloc
- Basic usage of performance variables, [746](#)
- MPI_T_pvar_handle_free
- Basic usage of performance variables, [746](#)
- MPI_T_pvar_read
- Basic usage of performance variables, [746](#)
- MPI_T_pvar_session_create
- Basic usage of performance variables, [746](#)
- MPI_T_pvar_start
- Basic usage of performance variables, [746](#)
- MPI_Test_cancelled
- Finalize and cancel, [490](#)
- MPI_TYPE_COMMIT
- Fortran 90
- derived types, [819](#)
- Nested vector datatypes, [171](#)
- Send/receive of a 3D array, [170](#)
- Transpose with datatypes, [172](#)
- Type_commit, [157](#)
- Using Get with indexed datatype, [570](#)
- Using indexed datatype, [171](#)
- MPI_Type_commit
- Gather and Gatherv, [207](#)
- Gather with datatype, [202](#)
- Gatherv with datatype, [203](#), [204](#)
- Gatherv with struct datatype, [205](#)
- Gatherv with vector datatype, [206](#)
- Pack/Unpack with struct datatype, [183](#)
- Scatterv with vector datatype, [213](#)
- User-defined operation with Scan, [249](#)
- Using datatypes with array of structures,
 [172](#)
- Using datatypes with array of structures
 with absolute addresses, [175](#)
- Using datatypes with unions, [176](#)
- MPI_TYPE_CONTIGUOUS
- Matching type with datatypes, [159](#)
- Typemap for contiguous, [130](#)
- Typemap of nested datatypes, [152](#)
- Using Get_count and Get_elements, [161](#)
- MPI_Type_contiguous
- Gather with datatype, [202](#)
- MPI_TYPE_CREATE_DARRAY
- Datatypes for distributed arrays, [148](#)
- MPI_TYPE_CREATE_F90_INTEGER
- Fortran 90

selected KIND, 808	
MPI_TYPE_CREATE_F90_REAL	
Fortran 90	
selected KIND, 808	
MPI_TYPE_CREATE_HVECTOR	
Nested vector datatypes, 171	
Send/receive of a 3D array, 170	
MPI_Type_create_hvector	
Using datatypes with array of structures, 172	
Using datatypes with array of structures with absolute addresses, 175	
MPI_TYPE_CREATE_INDEXED_BLOCK	
Using Get with indexed datatype, 570	
MPI_TYPE_CREATE_RESIZED	
Fortran 90	
derived types, 819	
MPI_Type_create_resized	
Using datatypes with unions, 176	
MPI_TYPE_CREATE_STRUCT	
Fortran 90	
derived types, 819	
Interlanguage communication, 847	
Transpose with datatypes, 172	
Typemap for create struct, 140	
Typemap of nested datatypes, 152	
MPI_Type_create_struct	
Gather and Gatherv, 207	
Gatherv with struct datatype, 205	
Pack/Unpack with struct datatype, 183	
User-defined operation with Scan, 249	
Using datatypes with array of structures, 172	
Using datatypes with array of structures with absolute addresses, 175	
MPI_TYPE_CREATE_SUBARRAY	
Local subarray for file output in Fortran 90, 715	
MPI_Type_create_subarray	
Local subarray for file output, 714	
MPI_TYPE_EXTENT	
Using Get with indexed datatype, 570	
MPI_TYPE_FREE	
Using Get with indexed datatype, 570	
MPI_Type_get_contents	
Decoding a datatype, 177	
MPI_Type_get_envelope	
Decoding a datatype, 177	
MPI_TYPE_GET_EXTENT	
Accumulate in RMA, 574	
Nested vector datatypes, 171	
Send/receive of a 3D array, 170	
Transpose with datatypes, 172	
Using Get, 572	
MPI_Type_get_extent	
Using datatypes with array of structures, 172	
MPI_TYPE_GET_VALUE_INDEX	
Retrieving an unnamed predefined value-index handle, 232	
MPI_TYPE_INDEXED	
Typemap for indexed, 134	
Using indexed datatype, 171	
MPI_Type_indexed	
Using datatypes with array of structures, 172	
Using datatypes with array of structures with absolute addresses, 175	
MPI_TYPE_MATCH_SIZE	
Fortran 90	
heterogeneous MPI I/O (unsafe), 812	
heterogeneous communication (unsafe), 812	
MPI_TYPE_MATCH_SIZE implementation, 811	
MPI_TYPE_MATCH_SIZE implementation	
MPI_TYPE_MATCH_SIZE, 811	
MPI_Type_size	
Measurement wrapper, 720	
MPI_TYPE_VECTOR	
Nested vector datatypes, 171	
Send/receive of a 3D array, 170	
Transpose with datatypes, 172	
Typemap for vector, 131	
MPI_Type_vector	
Gatherv with datatype, 203 , 204	
Gatherv with vector datatype, 206	
Scatterv with vector datatype, 213	
MPI_Unpack	
Pack and Pack_size, 184	
Pack/Unpack with struct datatype, 183	
MPI_Unpublish_name	
Name publishing, 538	
MPI_USER_FUNCTION	
Defining a user function, 240	
MPI_WAIT	
Client server code with waitsome, 90	
Client-server, 89	
File pointer update semantics, 668	
Message ordering for nonblocking operations, 82	
Nonblocking point-to-point, 81	
Nonblocking send and receive with request free, 82	
Progress semantics, 83	
MPI_Wait	
Erroneous matching of blocking and nonblocking collectives, 301	

- 1 Erroneous matching of collectives, [300](#)
- 2 Mixing blocking and nonblocking
- 3 collective operations, [300](#)
- 4 Mixing collective and point-to-point, [302](#)
- 5 Progress of nonblocking collectives, [301](#)
- 6 MPI_Waitall
- 7 Mixing collective and point-to-point, [302](#)
- 8 Overlapping communicators and
- 9 collectives, [303](#)
- 10 Pipelining nonblocking collectives, [302](#)
- 11 Requests in RMA, [623](#)
- 12 MPI_WAITANY
- 13 Client-server, [89](#)
- 14 MPI_Waitany
- 15 Requests in RMA, [623](#)
- 16 MPI_WAITSOME
- 17 Client server code with waitsome, [90](#)
- 18 MPI_Win_attach
- 19 Linked list in RMA, [624](#)
- 20 MPI_Win_complete
- 21 Active target and local reads in RMA, [613](#)
- 22 Double buffer in RMA, [620](#)
- 23 Get with PSCW, [619](#)
- 24 Public and private memory in RMA, [612](#)
- 25 Put with PSCW, [619](#)
- 26 RMA Start and complete, [594](#)
- 27 MPI_WIN_CREATE
- 28 Accumulate in RMA, [574](#)
- 29 Using Get, [572](#)
- 30 Using Get with indexed datatype, [570](#)
- 31 MPI_Win_create_dynamic
- 32 Linked list in RMA, [624](#)
- 33 MPI_Win_detach
- 34 Linked list in RMA, [624](#)
- 35 MPI_WIN_FENCE
- 36 Accumulate in RMA, [574](#)
- 37 Fortran 90 register optimization in RMA,
- 38 [827](#)
- 39 Using Get, [572](#)
- 40 Using Get with indexed datatype, [570](#)
- 41 MPI_Win_fence
- 42 Deadlock due to synchronization through
- 43 shared memory, [617](#)
- 44 Get with fence, [619](#)
- 45 Put with Fence, [618](#)
- 46 Register and Compiler Optimization, [618](#)
- 47 MPI_Win_flush
- 48 Counting semaphore (non-scalable), [621](#)
- Critical region with Compare-and-Swap,
- [622](#)
- Critical region with RMA, [621](#)
- Linked list in RMA, [624](#)
- Read data in RMA (unsafe), [611](#)
- MPI_Win_flush_all
- Critical region with RMA, [621](#)
- MPI_Win_flush_local
- Update location in unified memory model,
- [610](#)
- MPI_WIN_FREE
- Accumulate in RMA, [574](#)
- Using Get, [572](#)
- MPI_Win_lock
- Active target and local reads in RMA, [613](#)
- Public and private memory in RMA, [612](#)
- Read data in RMA, [611](#)
- RMA Lock and unlock, [600](#)
- Update location in separate memory
- model, [610](#)
- MPI_Win_lock_all
- Linked list in RMA, [624](#)
- Read data in RMA (unsafe), [611](#)
- Requests in RMA, [623](#)
- Shared memory windows, [623](#)
- MPI_Win_post
- Active target and local reads in RMA, [613](#)
- Double buffer in RMA, [620](#)
- Get with PSCW, [619](#)
- Public and private memory in RMA, [612](#)
- Put with PSCW, [619](#)
- MPI_Win_shared_query
- Deadlock due to synchronization through
- shared memory, [617](#)
- MPI_Win_start
- Active target and local reads in RMA, [613](#)
- Double buffer in RMA, [620](#)
- Get with PSCW, [619](#)
- Public and private memory in RMA, [612](#)
- Put with PSCW, [619](#)
- RMA Start and complete, [594](#)
- MPI_Win_sync
- shared memory windows, [623](#)
- MPI_Win_sync
- Counting semaphore (non-scalable), [621](#)
- Critical region with Compare-and-Swap,
- [622](#)
- Critical region with RMA, [621](#)
- Update location in unified memory model,
- [610](#)
- MPI_Win_unlock
- Active target and local reads in RMA, [613](#)
- Public and private memory in RMA, [612](#)
- Read data in RMA, [611](#)
- RMA Lock and unlock, [600](#)
- Update location in separate memory
- model, [610](#)
- MPI_Win_unlock_all
- Linked list in RMA, [624](#)
- Read data in RMA (unsafe), [611](#)

Requests in RMA, 623	
MPI_Win_wait	
Active target and local reads in RMA, 613	
Double buffer in RMA, 620	
Get with PSCW, 619	
Public and private memory in RMA, 612	
Put with PSCW, 619	
MPI_WTIME	
Using MPI_Wtime, 470	
MPI_Wtime	
Measurement wrapper, 720	
mpiexec, 514	
mpiexec	
mpiexec and environment variables, 484	
Using mpiexec	
starting programs with separate	
argument lists, 514	
using -arch, 514	
using -configfile, 514	
using -host, 514	
using -n, 513	
mpiexec and environment variables	
MPI_INFO_ENV, 484	
mpiexec, 484	
Name publishing	
MPI_Comm_accept, 538	
MPI_Comm_connect, 538	
MPI_Open_port, 538	
MPI_Publish_name, 538	
MPI_Unpublish_name, 538	
Neighbor allgather	
MPI_NEIGHBOR_ALLGATHER, 412	
Neighbor alltoall	
MPI_CART_SHIFT, 417	
MPI_CARTDIM_GET, 417	
MPI_NEIGHBOR_ALLTOALL, 417	
MPI_SENDRECV, 417	
Neighborhood collective communication, 437	
MPI_CART_GET, 437	
MPI_CART_SHIFT, 437	
MPI_DIMS_CREATE, 437	
Nested vector datatypes	
MPI_SENDRECV, 171	
MPI_TYPE_COMMIT, 171	
MPI_TYPE_CREATE_HVECTOR, 171	
MPI_TYPE_GET_EXTENT, 171	
MPI_TYPE_VECTOR, 171	
Nonblocking operations, 81 , 82	
message ordering, 82	
progress, 83	
Nonblocking point-to-point	
MPI_Irecv, 81	
MPI_Isend, 81	
MPI_Wait, 81	1
Nonblocking send and receive with request free	2
MPI_Irecv, 82	3
MPI_Isend, 82	4
MPI_REQUEST_FREE, 82	5
MPI_Wait, 82	6
Nondeterministic program with MPI_Bcast,	7
298	
Nondeterministic use of Bcast	8
MPI_Bcast, 298	9
Nonovertaking messages, 55	10
MPI_BSEND, 55	11
MPI_RECV, 55	12
Overlap computation and output	13
MPI_File_write_all_begin, 712	14
MPI_File_write_all_end, 712	15
Overlapping communicators, 303	16
Overlapping communicators and collectives	17
MPI_lallreduce, 303	18
MPI_Waitall, 303	19
Pack and Pack_size	20
MPI_Gather, 184	21
MPI_Gatherv, 184	22
MPI_Pack, 184	23
MPI_Pack_size, 184	24
MPI_Unpack, 184	25
Pack/Unpack with struct datatype	26
MPI_Get_address, 183	27
MPI_Pack, 183	28
MPI_Send, 183	28
MPI_Type_commit, 183	29
MPI_Type_create_struct, 183	30
MPI_Unpack, 183	31
Partitioned communication	32
Equal send/rcv partitioning, 122	33
MPI_Pready, 114	34
MPI_Precv_init, 114	35
MPI_Psend_init, 114	35
Partial completion notification, 125	36
Send with tasks, 123	37
Simple example, 114	38
Partitioned communication using threads	39
MPI_Pready, 122	40
MPI_Precv_init, 122	41
MPI_Psend_init, 122	42
Partitioned communication with partial	43
completion	
MPI_Parrived, 125	44
MPI_Pready, 125	45
MPI_Precv_init, 125	46
MPI_Psend_init, 125	47
Partitioned communication with tasks	48

- 1 MPI_Pready, [123](#)
- 2 MPI_Precv_init, [123](#)
- 3 MPI_Psend_init, [123](#)
- 4 Pipelining nonblocking collective operations, [302](#)
- 5
- 6 Pipelining nonblocking collectives
- 7 MPI_Ibcast, [302](#)
- 8 MPI_Waitall, [302](#)
- 9 Point-to-point
- 10 Hello world, [31](#)
- 11 Profiling interface
- 12 implementation using the C macro
- 13 preprocessor, [720](#)
- 14 implementation using weak symbols, [720](#)
- 15 measurement wrapper, [720](#)
- 16 Progress of matching pairs, [56](#)
- 17 Progress of nonblocking collective operations, [301](#)
- 18 Progress of nonblocking collectives
- 19 MPI_Ibarrier, [301](#)
- 20 MPI_Recv, [301](#)
- 21 MPI_Send, [301](#)
- 22 MPI_Wait, [301](#)
- 23 Progress semantics
- 24 MPI_Irecv, [83](#)
- 25 MPI_RECV, [83](#)
- 26 MPI_SEND, [83](#)
- 27 MPI_SSEND, [83](#)
- 28 MPI_WAIT, [83](#)
- 29 Public and private memory in RMA
- 30 MPI_Barrier, [612](#)
- 31 MPI_Get, [612](#)
- 32 MPI_Win_complete, [612](#)
- 33 MPI_Win_lock, [612](#)
- 34 MPI_Win_post, [612](#)
- 35 MPI_Win_start, [612](#)
- 36 MPI_Win_unlock, [612](#)
- 37 MPI_Win_wait, [612](#)
- 38 Put with Fence
- 39 MPI_Put, [618](#)
- 40 MPI_Win_fence, [618](#)
- 41 Put with PSCW
- 42 MPI_Put, [619](#)
- 43 MPI_Win_complete, [619](#)
- 44 MPI_Win_post, [619](#)
- 45 MPI_Win_start, [619](#)
- 46 MPI_Win_wait, [619](#)
- 47
- 48 Read data in RMA
- MPI_Barrier, [611](#)
- MPI_Put, [611](#)
- MPI_Win_lock, [611](#)
- MPI_Win_unlock, [611](#)
- Read data in RMA (unsafe)
- MPI_Barrier, [611](#)
- MPI_Put, [611](#)
- MPI_Win_flush, [611](#)
- MPI_Win_lock_all, [611](#)
- MPI_Win_unlock_all, [611](#)
- Read to end of file
- MPI_FILE_CLOSE, [665](#)
- MPI_FILE_OPEN, [665](#)
- MPI_FILE_READ, [665](#)
- MPI_FILE_SET_VIEW, [665](#)
- Reading a control variable
- MPI_T_cvar_handle_alloc, [736](#)
- MPI_T_cvar_handle_free, [736](#)
- MPI_T_cvar_read, [736](#)
- Recursive splitting of COMM_WORLD
- MPI_Comm_split_type, [336](#)
- Reduction
- MPI_REDUCE, [228](#)
- Reduction of a vector
- MPI_REDUCE, [229](#)
- Reduction with maxloc
- MPI_REDUCE, [233](#)
- MPI_Reduce, [233](#)
- Reduction with minloc
- MPI_Reduce, [234](#)
- Reduction with user-defined op
- MPI_Op_create, [239](#)
- MPI_Reduce, [239](#)
- Register and Compiler Optimization
- MPI_Put, [618](#)
- MPI_Win_fence, [618](#)
- Requesting and querying memory allocation
- kinds in the Sessions Model
- mpi_assert_memory_alloc_kinds, [510](#)
- mpi_memory_alloc_kinds, [510](#)
- Requests in RMA
- MPI_Rget, [623](#)
- MPI_Rput, [623](#)
- MPI_Waitall, [623](#)
- MPI_Waitany, [623](#)
- MPI_Win_lock_all, [623](#)
- MPI_Win_unlock_all, [623](#)
- Retrieving an unnamed predefined value-index
- handle
- MPI_TYPE_GET_VALUE_INDEX, [232](#)
- RMA Lock and unlock
- MPI_Put, [600](#)
- MPI_Win_lock, [600](#)
- MPI_Win_unlock, [600](#)
- RMA Start and complete
- MPI_Put, [594](#)
- MPI_Win_complete, [594](#)
- MPI_Win_start, [594](#)
- Rules for finalize

MPI_Finalize, 489		
Scatter		
MPI_Scatter, 212		
Scatterv		
MPI_Scatterv, 212		
Scatterv with vector datatype		
MPI_Scatterv, 213		
MPI_Type_commit, 213		
MPI_Type_vector, 213		
Send/receive of a 3D array		
MPI_SENDRECV, 170		
MPI_TYPE_COMMIT, 170		
MPI_TYPE_CREATE_HVECTOR, 170		
MPI_TYPE_GET_EXTENT, 170		
MPI_TYPE_VECTOR, 170		
Shared memory windows		
MPI_F_SYNC_REG, 623		
MPI_Win_lock_all, 623		
MPI_Win_sync, 623		
Splitting into NUMANode subcommunicators		
MPI_Comm_split_type, 334		
Subgroup cart process topology		
MPI_CART_SUB, 407		
Threads and MPI, 515		
Three-Group “Pipeline”		
MPI_Comm_split, 356		
MPI_Intercomm_create, 356		
Three-Group “Ring”		
MPI_Comm_split, 358		
MPI_Intercomm_create, 358		
Tool information interface		
basic usage of performance variables, 746		
listing names of all control variables, 732		
reading the value of a control variable, 736		
Topologies, 437		
Transpose with datatypes		
MPI_SENDRECV, 172		
MPI_TYPE_COMMIT, 172		
MPI_TYPE_CREATE_STRUCT, 172		
MPI_TYPE_GET_EXTENT, 172		
MPI_TYPE_VECTOR, 172		
Type_commit		
MPI_TYPE_COMMIT, 157		
Typemap, 129–131 , 134 , 140 , 148		
Typemap for contiguous		
MPI_TYPE_CONTIGUOUS, 130		
Typemap for create struct		
MPI_TYPE_CREATE_STRUCT, 140		
Typemap for indexed		
MPI_TYPE_INDEXED, 134		
Typemap for vector		
MPI_TYPE_VECTOR, 131		
Typemap of nested datatypes		1
MPI_TYPE_CONTIGUOUS, 152		2
MPI_TYPE_CREATE_STRUCT, 152		3
Update location in separate memory model		4
MPI_Barrier, 610		5
MPI_Get, 610		6
MPI_Win_lock, 610		7
MPI_Win_unlock, 610		8
Update location in unified memory model		9
MPI_Barrier, 610		10
MPI_Get, 610		11
MPI_Win_flush_local, 610		12
MPI_Win_sync, 610		13
Use of Alloc_mem in Fortran		14
MPI_ALLOC_MEM, 448		15
MPI_FREE_MEM, 448		16
Use of Alloc_mem in Fortran with Cray pointers		17
MPI_ALLOC_MEM, 448		18
MPI_FREE_MEM, 448		19
User-defined operation with Scan		20
MPI_Op_create, 249		21
MPI_Scan, 249		22
MPI_Type_commit, 249		23
MPI_Type_create_struct, 249		24
User-defined reduce		25
MPI_Grequest_complete, 631		26
MPI_Grequest_start, 631		27
Using Alloc_mem		28
MPI_Alloc_mem, 449		29
Using Cart_shift		30
MPI_CART_COORDS, 406		31
MPI_CART_RANK, 406		32
MPI_CART_SHIFT, 406		33
MPI_SENDRECV_REPLACE, 406		34
Using datatypes with array of structures		35
MPI_Aint, 172		36
MPI_Get_address, 172		37
MPI_Send, 172		38
MPI_Type_commit, 172		39
MPI_Type_create_hvector, 172		40
MPI_Type_create_struct, 172		41
MPI_Type_get_extent, 172		42
MPI_Type_indexed, 172		43
Using datatypes with array of structures with absolute addresses		44
MPI_Get_address, 175		45
MPI_Send, 175		46
MPI_Type_commit, 175		47
MPI_Type_create_hvector, 175		48
MPI_Type_create_struct, 175		
MPI_Type_indexed, 175		
Using datatypes with unions		

- 1 MPI_Get_address, [176](#)
- 2 MPI_Send, [176](#)
- 3 MPI_Type_commit, [176](#)
- 4 MPI_Type_create_resized, [176](#)
- 5 Using Get
- 6 MPI_GET, [572](#)
- 7 MPI_TYPE_GET_EXTENT, [572](#)
- 8 MPI_WIN_CREATE, [572](#)
- 9 MPI_WIN_FENCE, [572](#)
- 10 MPI_WIN_FREE, [572](#)
- 11 Using Get with indexed datatype
- 12 MPI_GET, [570](#)
- 13 MPI_TYPE_COMMIT, [570](#)
- 14 MPI_TYPE_CREATE_INDEXED_BLOCK,
- 15 [570](#)
- 16 MPI_TYPE_EXTENT, [570](#)
- 17 MPI_TYPE_FREE, [570](#)
- 18 MPI_WIN_CREATE, [570](#)
- 19 MPI_WIN_FENCE, [570](#)
- 20 Using Get_count and Get_elements
- 21 MPI_GET_COUNT, [161](#)
- 22 MPI_GET_ELEMENTS, [161](#)
- 23 MPI_TYPE_CONTIGUOUS, [161](#)
- 24 Using Group_excl
- 25 MPI_Comm_create, [344](#)
- 26 MPI_Group_excl, [344](#)
- 27 MPI_Group_free, [344](#)
- 28 Using indexed datatype
- 29 MPI_SENDRECV, [171](#)
- 30 MPI_TYPE_COMMIT, [171](#)
- 31 MPI_TYPE_INDEXED, [171](#)
- 32 Using MPI_Wtime
- 33 MPI_WTIME, [470](#)
- 34 Using mpiexec
- 35 starting programs with separate
- 36 argument lists
- 37 mpiexec, [514](#)
- 38 using -arch
- 39 mpiexec, [514](#)
- 40 using -configfile
- 41 mpiexec, [514](#)
- 42 using -host
- 43 mpiexec, [514](#)
- 44 using -n
- 45 mpiexec, [513](#)
- 46 Using Pack
- 47 MPI_Pack, [183](#)
- 48 Using process set query in group creation
- 49 MPI_BARRIER, [505](#)
- 50 MPI_COMM_CREATE_FROM_GROUP,
- 51 [505](#)
- 52 MPI_GROUP_FROM_SESSION_PSET, [505](#)
- 53 MPI_Group_from_session_pset, [503](#)
- 54 MPI_SESSION_FINALIZE, [505](#)
- 55 MPI_Session_finalize, [503](#)
- 56 MPI_SESSION_GET_NTH_PSET, [505](#)
- 57 MPI_Session_get_nth_pset, [503](#)
- 58 MPI_SESSION_GET_NUM_PSETS, [505](#)
- 59 MPI_Session_get_num_psets, [503](#)
- 60 MPI_SESSION_INIT, [505](#)
- 61 MPI_Session_init, [503](#)
- 62 Virtual topologies, [437](#)

MPI Constant and Predefined Handle Index

This index lists predefined MPI constants and handles, including info keys and values. Underlined page numbers give the location of the primary definition or use of the indexed term.

"access_style", [648](#), [873](#)
"accumulate_ops", [550](#), [873](#)
"accumulate_ordering", [550](#), [614](#), [873](#)
"alloc_mem", [510](#), [874](#)
"alloc_shared_noncontig", [555](#), [556](#), [873](#)
"appnum", [542](#), [873](#)
"arch", [484](#), [527](#), [873](#)
"argv", [484](#), [873](#)

"cb_block_size", [649](#), [873](#)
"cb_buffer_size", [649](#), [873](#)
"cb_nodes", [649](#), [873](#)
"chunked", [649](#), [874](#)
"chunked_item", [649](#), [873](#)
"chunked_size", [649](#), [874](#)
"collective_buffering", [648](#), [874](#)
"command", [484](#), [874](#)

"external32", [185](#), [690–696](#), [700](#), [807–810](#), [812](#),
[873](#), [1056](#), [1064](#)

"false", [340](#), [341](#), [445](#), [474](#), [550](#), [551](#), [648](#), [874](#)
"file", [484](#), [527](#), [874](#)
"file_perm", [648](#), [649](#), [874](#)
"filename", [649](#), [874](#)

"host", [484](#), [527](#), [874](#)
"hwloc://", [445](#)
"hwloc://L3Cache", [497](#)

"internal", [690–692](#), [700](#), [873](#)
"io_node_list", [649](#), [874](#)
"ip_address", [538](#), [874](#)
"ip_port", [538](#), [874](#)

"maxprocs", [484](#), [485](#), [874](#)
"mpi", [510](#), [874](#)
"mpi://", [445](#), [497](#), [873](#)
"mpi://SELF", [497](#), [498](#), [640](#), [873](#)
"mpi://WORLD", [497](#), [498](#), [503](#), [513](#), [541](#), [873](#)
MPI::_LONG_LONG, [1056](#)
MPI::_BOOL, [1056](#)
MPI::_COMPLEX, [1056](#)
MPI::_DOUBLE_COMPLEX, [1056](#)
MPI::_F_COMPLEX16, [1056](#)
MPI::_F_COMPLEX32, [1056](#)
MPI::_F_COMPLEX4, [1056](#)
MPI::_F_COMPLEX8, [1056](#)
MPI::_INTEGER16, [1056](#)
MPI::_LONG_DOUBLE_COMPLEX, [1056](#)
MPI::_LONG_LONG, [1056](#)
MPI::_REAL16, [1056](#)
MPI_2DOUBLE_PRECISION, [231](#), [856](#)
MPI_2INT, [231](#), [856](#)
MPI_2INTEGER, [231](#), [856](#)
MPI_2REAL, [231](#), [856](#)
"mpi_accumulate_granularity", [550](#), [874](#), [1049](#)
MPI_ADDRESS_KIND, [21](#), [35](#), [129](#), [148](#), [167](#),
[360](#), [447](#), [553](#), [555](#), [558](#), [801](#), [814](#), [833](#),
[834](#), [843](#), [852](#), [855](#)
MPI_AINT, [35](#), [129](#), [228](#), [560](#), [695](#), [833](#), [854](#),
[855](#), [1062](#), [1064](#)
MPI_ANY_SOURCE, [37](#), [38](#), [43](#), [45](#), [55](#), [75–77](#),
[79](#), [94](#), [95](#), [97–99](#), [108](#), [340](#), [381](#), [443](#),
[851](#)
MPI_ANY_TAG, [19](#), [37](#), [38](#), [40](#), [44](#), [45](#), [75–77](#),
[79](#), [94](#), [95](#), [97–99](#), [101](#), [102](#), [108](#), [111](#),
[340](#), [851](#), [1058](#)
MPI_APPNUM, [542](#), [860](#)
MPI_ARGV_NULL, [20](#), [521](#), [522](#), [526](#), [814](#), [862](#)
MPI_ARGVS_NULL, [20](#), [526](#), [814](#), [862](#)
"mpi_assert_allow_overtaking", [341](#), [874](#), [1051](#)
"mpi_assert_exact_length", [341](#), [874](#), [1051](#)
"mpi_assert_memory_alloc_kinds", [341](#), [509](#),
[510](#), [551](#), [650](#), [874](#), [1048](#)
"mpi_assert_no_any_source", [340](#), [874](#), [1051](#)
"mpi_assert_no_any_tag", [340](#), [874](#), [1051](#)
"mpi_assert_strict_persistent_collective_ordering",
[341](#), [1048](#)
MPI_ASYNC_PROTECTS_NONBLOCKING,
[622](#), [788](#), [789](#), [791](#), [793](#), [796](#), [804](#), [805](#),
[825](#), [852](#), [1061](#)
MPI_BAND, [227](#), [228](#), [857](#)
MPI_BOR, [227](#), [228](#), [857](#)
MPI_BOTTOM, [10](#), [20](#), [21](#), [42](#), [59](#), [60](#), [121](#), [148](#),
[161](#), [162](#), [193](#), [392](#), [394](#), [523](#), [560](#), [564](#),

1	789, 792, 800, 814, 819, 821, 823, 824,	MPI_COMM_NULL, 309 , 325 , 327–330 ,
2	826–828 , 831 , 841 , 842 , 847 , 851 , 1068	333–336 , 339 , 340 , 355 , 377 , 387 , 389 ,
3	MPI_BSEND_OVERHEAD, 68 , 851	524 , 544 , 545 , 836 , 858 , 1048 , 1065
4	MPI_BUFFER_AUTOMATIC, 20 , 59 , 60 , 63 , 67 ,	"MPI_COMM_NULL", 377 , 873
5	851 , 1048	MPI_COMM_NULL_COPY_FN, 23 , 24 , 362 ,
6	MPI_BXOR, 227 , 228 , 857	363 , 364 , 790 , 842 , 859 , 1055 , 1060
7	MPI_BYTE, 33 , 34 , 46 , 47 , 49 , 185 , 228 , 638 ,	MPI_COMM_NULL_DELETE_FN, 24 , 363 , 364 ,
8	690 , 691 , 695 , 705 , 847 , 854 , 855 , 1064	859 , 1055
9	Fortran example usage, 47	"MPI_COMM_PARENT", 377 , 873
10	MPI_C_BOOL, 34 , 228 , 695 , 854 , 1056 , 1062 ,	MPI_COMM_SELF, 26 , 309 , 328 , 342 , 359 , 377 ,
11	1064	449 , 450 , 482 , 483 , 488 , 490 , 492 , 493 ,
12	MPI_C_COMPLEX, 34 , 228 , 695 , 854 , 1056 ,	544 , 640 , 785 , 856 , 1050 , 1063
13	1062 , 1064	"MPI_COMM_SELF", 377 , 873
14	MPI_C_DOUBLE_COMPLEX, 34 , 228 , 695 ,	MPI_COMM_TYPE_HW_GUIDED, 333 , 334 ,
15	854 , 1062 , 1064	446 , 856 , 1051
16	MPI_C_FLOAT_COMPLEX, 34 , 228 , 695 , 854 ,	MPI_COMM_TYPE_HW_UNGUIDED, 336 ,
17	1062 , 1064	337 , 856 , 1051
18	MPI_C_LONG_DOUBLE_COMPLEX, 34 , 228 ,	MPI_COMM_TYPE_RESOURCE_GUIDED, 334 ,
19	695 , 854 , 1062 , 1064	335 , 446 , 856 , 1048
20	MPI_CART, 397 , 858	MPI_COMM_TYPE_SHARED, 333 , 334 , 335 ,
21	MPI_CHAR, 34 , 49 , 140 , 229 , 695 , 727 , 728 ,	617 , 856 , 1058
22	854 , 1061 , 1062	MPI_COMM_WORLD, 19 , 29 , 31 , 32 , 36 , 307 ,
23	MPI_CHARACTER, 33 , 48 , 49 , 229 , 695 , 855	309 , 311 , 312 , 320 , 321 , 334–337 , 342 ,
24	Fortran example usage, 48	345 , 353 , 377 , 442–444 , 453 , 465 ,
25	MPI_COMBINER_CONTIGUOUS, 164 , 168 , 861	481–483 , 485 , 487–491 , 493 , 508 , 513 ,
26	MPI_COMBINER_DARRAY, 164 , 170 , 861	514 , 518–521 , 524 , 526 , 540–544 , 689 ,
27	MPI_COMBINER_DUP, 164 , 168 , 861	711 , 734 , 742 , 777 , 778 , 785 , 835 , 847 ,
28	MPI_COMBINER_F90_COMPLEX, 164 , 170 ,	856 , 881 , 1050 , 1066
29	861	"MPI_COMM_WORLD", 377 , 873
30	MPI_COMBINER_F90_INTEGER, 164 , 170 , 861	MPI_COMPLEX, 33 , 228 , 693 , 695 , 806 , 855
31	MPI_COMBINER_F90_REAL, 164 , 170 , 861	MPI_COMPLEX16, 35 , 228 , 696 , 855
32	MPI_COMBINER_HINDEXED, 24 , 164 , 169 ,	MPI_COMPLEX32, 35 , 228 , 696 , 855
33	861	MPI_COMPLEX4, 35 , 228 , 696 , 855
34	MPI_COMBINER_HINDEXED_BLOCK, 164 ,	MPI_COMPLEX8, 35 , 228 , 696 , 855
35	169 , 861 , 1058	MPI_CONGRUENT, 321 , 352 , 856
36	MPI_COMBINER_HINDEXED_INTEGER, 24 ,	MPI_CONVERSION_FN_NULL, 700 , 859 , 1055
37	784 , 1057	MPI_CONVERSION_FN_NULL_C, 700 , 833 ,
38	MPI_COMBINER_HVECTOR, 24 , 164 , 168 , 861	859 , 1050
39	MPI_COMBINER_HVECTOR_INTEGER, 24 ,	MPI_COUNT, 35 , 129 , 228 , 237 , 695 , 727 , 833 ,
40	784 , 1057	854 , 855 , 1057
41	MPI_COMBINER_INDEXED, 164 , 169 , 861	MPI_COUNT_KIND, 21 , 22 , 35 , 834 , 852 , 855 ,
42	MPI_COMBINER_INDEXED_BLOCK, 164 , 169 ,	1057
43	861	MPI_CXX_BOOL, 35 , 228 , 695 , 696 , 855 , 1056
44	MPI_COMBINER_NAMED, 164 , 168 , 232 , 861	MPI_CXX_DOUBLE_COMPLEX, 35 , 228 , 695 ,
45	MPI_COMBINER_RESIZED, 164 , 170 , 861	696 , 855 , 1056
46	MPI_COMBINER_STRUCT, 24 , 164 , 169 , 861	MPI_CXX_FLOAT_COMPLEX, 35 , 228 , 695 ,
47	MPI_COMBINER_STRUCT_INTEGER, 24 , 784 ,	696 , 855 , 1056
48	1057	MPI_CXX_LONG_DOUBLE_COMPLEX, 35 ,
	MPI_COMBINER_SUBARRAY, 164 , 169 , 861	228 , 695 , 696 , 855 , 1056
	MPI_COMBINER_VALUE_INDEX, 164 , 170 ,	MPI_DATATYPE_NULL, 157 , 232 , 379 , 858 ,
	232 , 1048	1048
	MPI_COMBINER_VECTOR, 164 , 168 , 861	"MPI_DATATYPE_NULL", 379 , 873
	MPI_COMM_DUP_FN, 24 , 362 , 363 , 364 , 794 ,	MPI_DISPLACEMENT_CURRENT, 651 , 862 ,
	859 , 1055 , 1060	1067

MPI_DIST_GRAPH, 397 , 858 , 1063	MPI_ERR_QUOTA, 463 , 712 , 850	1
MPI_DISTRIBUTE_BLOCK, 145 , 146 , 862	MPI_ERR_RANK, 463 , 605 , 849	2
MPI_DISTRIBUTE_CYCLIC, 145 , 146 , 862	MPI_ERR_READ_ONLY, 463 , 712 , 850	3
MPI_DISTRIBUTE_DFLT_DARG, 145 , 146 , 862	MPI_ERR_REQUEST, 463 , 849	4
MPI_DISTRIBUTE_NONE, 145 , 146 , 862	MPI_ERR_RMA_ATTACH, 463 , 605 , 850	5
MPI_DOUBLE, 34 , 227 , 695 , 727 , 737 , 738 , 805 , 854	MPI_ERR_RMA_CONFLICT, 463 , 605 , 850	6
MPI_DOUBLE_COMPLEX, 34 , 228 , 693 , 695 , 806 , 855	MPI_ERR_RMA_FLAVOR, 463 , 605 , 850	7
MPI_DOUBLE_INT, 231 , 856	MPI_ERR_RMA_RANGE, 463 , 605 , 850	8
MPI_DOUBLE_PRECISION, 33 , 227 , 695 , 806 , 855	MPI_ERR_RMA_SHARED, 463 , 605 , 850	9
MPI_DUP_FN, 24 , 363 , 772 , 860	MPI_ERR_RMA_SYNC, 463 , 605 , 850	10
MPI_ERR_ACCESS, 462 , 643 , 712 , 850	MPI_ERR_ROOT, 463 , 849	11
MPI_ERR_AMODE, 462 , 641 , 712 , 850	MPI_ERR_SERVICE, 463 , 537 , 850	12
MPI_ERR_ARG, 462 , 849	MPI_ERR_SESSION, 463 , 850 , 1052	13
MPI_ERR_ASSERT, 462 , 605 , 850	MPI_ERR_SIZE, 463 , 605 , 850	14
MPI_ERR_BAD_FILE, 462 , 712 , 850	MPI_ERR_SPAWN, 463 , 522 , 523 , 850	15
MPI_ERR_BASE, 448 , 462 , 605 , 850	MPI_ERR_TAG, 463 , 849	16
MPI_ERR_BUFFER, 462 , 849	MPI_ERR_TOPOLOGY, 463 , 849	17
MPI_ERR_COMM, 462 , 849	MPI_ERR_TRUNCATE, 463 , 849	18
MPI_ERR_CONVERSION, 462 , 700 , 712 , 850	MPI_ERR_TYPE, 164 , 166 , 463 , 712 , 849	19
MPI_ERR_COUNT, 462 , 849	MPI_ERR_UNKNOWN, 461 , 463 , 849	20
MPI_ERR_DIMS, 462 , 849	MPI_ERR_UNSUPPORTED_DATAREP, 463 , 712 , 850	21
MPI_ERR_DISP, 462 , 605 , 850	MPI_ERR_UNSUPPORTED_OPERATION, 463 , 712 , 850	22
MPI_ERR_DUP_DATAREP, 462 , 697 , 712 , 850	MPI_ERR_VALUE_TOO_LARGE, 463 , 697 , 850 , 1050	23
MPI_ERR_ERRHANDLER, 462 , 850 , 1049	MPI_ERR_WIN, 463 , 605 , 850	24
MPI_ERR_FILE, 462 , 712 , 850	MPI_ERRCODES_IGNORE, 20 , 523 , 814 , 862	25
MPI_ERR_FILE_EXISTS, 462 , 712 , 850	MPI_ERRHANDLER_NULL, 459 , 858	26
MPI_ERR_FILE_IN_USE, 462 , 643 , 712 , 850	MPI_ERROR, 39 , 41 , 42 , 79 , 251 , 580 , 635 , 636 , 838 , 852 , 1053 , 1060	27
MPI_ERR_GROUP, 462 , 849	MPI_ERRORS_ABORT, 450 , 483 , 527 , 852 , 1050	28
MPI_ERR_IN_STATUS, 39 , 43 , 79 , 86 , 88 , 452 , 460 , 462 , 630 , 656 , 850	"mpi_errors_abort", 527 , 874 , 1052	29
MPI_ERR_INFO, 462 , 850	MPI_ERRORS_ARE_FATAL, 450 , 451 , 469 , 527 , 605 , 711 , 852 , 1050	30
MPI_ERR_INFO_KEY, 462 , 475 , 850	"mpi_errors_are_fatal", 527 , 874 , 1052	31
MPI_ERR_INFO_NOKEY, 462 , 475 , 850	MPI_ERRORS_RETURN, 450 , 451 , 469 , 508 , 527 , 711 , 847 , 852	32
MPI_ERR_INFO_VALUE, 462 , 475 , 850	"mpi_errors_return", 527 , 874 , 1052	33
MPI_ERR_INTERN, 451 , 462 , 849	MPI_F08_STATUS_IGNORE, 839 , 863 , 1060	34
MPI_ERR_IO, 462 , 712 , 850	MPI_F08_STATUSES_IGNORE, 839 , 863 , 1060	35
MPI_ERR_KEYVAL, 374 , 462 , 850	MPI_F_ERROR, 838 , 852 , 1050	36
MPI_ERR_LASTCODE, 461 , 463 , 465 , 467 , 768 , 851	MPI_F_SOURCE, 838 , 852 , 1050	37
MPI_ERR_LOCKTYPE, 462 , 605 , 850	MPI_F_STATUS_IGNORE, 838 , 863	38
MPI_ERR_NAME, 462 , 537 , 850	MPI_F_STATUS_SIZE, 838 , 852 , 1050	39
MPI_ERR_NO_MEM, 447 , 462 , 850	MPI_F_STATUSES_IGNORE, 838 , 863	40
MPI_ERR_NO_SPACE, 462 , 712 , 850	MPI_F_TAG, 838 , 852 , 1050	41
MPI_ERR_NO_SUCH_FILE, 462 , 642 , 712 , 850	MPI_FILE_NULL, 642 , 711 , 858	42
MPI_ERR_NOT_SAME, 462 , 712 , 850	MPI_FLOAT, 34 , 140 , 225 , 227 , 692 , 695 , 854	43
MPI_ERR_OP, 463 , 605 , 849	MPI_FLOAT_INT, 16 , 231 , 856	44
MPI_ERR_OTHER, 461 , 463 , 849	MPI_GRAPH, 397 , 858	45
MPI_ERR_PENDING, 86 , 463 , 849	MPI_GROUP_EMPTY, 308 , 314 , 315 , 325 , 327 , 328 , 339 , 355 , 858	46
MPI_ERR_PORT, 463 , 534 , 535 , 850		47
MPI_ERR_PROC_ABORTED, 463 , 508 , 850 , 1052		48

1	MPI_GROUP_NULL, 308 , 318 , 319 , 858	MPI_MAX_PORT_NAME, 532 , 853
2	MPI_HOST, 5 , 24 , 777 , 778 , 857 , 1047	MPI_MAX_PROCESSOR_NAME, 444 , 445 , 853 ,
3	"mpi_hw_resource_type", 333 – 336 , 446 , 874 ,	1066
4	1051	MPI_MAX_PSET_NAME_LEN, 500 , 853 , 1045 ,
5	MPI_IDENT, 312 , 321 , 856	1052
6	MPI_IN_PLACE, 20 , 192 , 221 , 792 , 814 , 851	MPI_MAX_STRINGTAG_LEN, 339 , 355 , 853 ,
7	MPI_INFO_ENV, 478 , 484 , 484 , 485 , 493 , 509 ,	1045 , 1052
8	856 , 1059	MPI_MAXLOC, 227 , 229 – 232 , 234 , 857 , 1048
9	MPI_INFO_NULL, 333 , 334 , 395 , 447 , 478 , 523 ,	"mpi_memory_alloc_kinds", 484 , 495 , 508 , 509 ,
10	533 , 641 , 642 , 652 , 858	528 , 874 , 1048
11	"mpi_initial_errhandler", 483 , 484 , 527 , 874 ,	MPI_MESSAGE_NO_PROC, 99 , 101 , 102 , 111 ,
12	1052	852 , 1058
13	MPI_INT, 16 , 34 , 128 , 227 , 692 , 693 , 695 , 727 ,	MPI_MESSAGE_NULL, 99 , 101 , 102 , 858 , 1058
14	728 , 731 , 737 , 740 , 805 , 847 , 848 , 854	MPI_MIN, 227 , 228 , 232 , 857
15	MPI_INT16_T, 34 , 227 , 695 , 854 , 1062 , 1064	"mpi_minimum_memory_alignment", 447 , 554 ,
16	MPI_INT32_T, 34 , 227 , 695 , 727 , 854 , 1062 ,	556 , 874 , 1051
17	1064	MPI_MINLOC, 227 , 229 – 232 , 234 , 857 , 1048
18	MPI_INT64_T, 34 , 227 , 695 , 727 , 854 , 1062 ,	MPI_MODE_APPEND, 640 , 641 , 861
19	1064	MPI_MODE_CREATE, 640 , 641 , 649 , 861
20	MPI_INT8_T, 34 , 227 , 695 , 854 , 1062 , 1064	MPI_MODE_DELETE_ON_CLOSE, 640 , 641 ,
21	MPI_INTEGER, 33 , 46 , 227 , 695 , 805 , 806 , 848 ,	642 , 861
22	855	MPI_MODE_EXCL, 640 , 641 , 861
23	MPI_INTEGER1, 35 , 227 , 696 , 855	MPI_MODE_NOCHECK, 598 , 603 , 604 , 621 ,
24	MPI_INTEGER16, 227 , 696 , 855	861
25	MPI_INTEGER2, 35 , 227 , 693 , 696 , 855	MPI_MODE_NOPRECEDE, 592 , 603 , 604 , 861
26	MPI_INTEGER4, 35 , 227 , 696 , 855	MPI_MODE_NOPUT, 603 , 604 , 861
27	MPI_INTEGER8, 35 , 227 , 696 , 810 , 855	MPI_MODE_NOSTORE, 603 , 604 , 861
28	MPI_INTEGER_KIND, 852	MPI_MODE_NOSUCCEED, 603 , 604 , 861
29	MPI_IO, 443 , 857	MPI_MODE_RDONLY, 640 , 641 , 646 , 861
30	MPI_KEYVAL_INVALID, 363 , 364 , 365 , 851	MPI_MODE_RDWR, 640 , 641 , 861
31	MPI_LAND, 227 , 228 , 857	MPI_MODE_SEQUENTIAL, 640 , 641 , 643 , 644 ,
32	MPI_LASTUSED_CODE, 465 , 860 , 1046	651 , 656 , 664 , 679 , 704 , 861 , 1067
33	MPI_LB, 24 , 784 , 1057	MPI_MODE_UNIQUE_OPEN, 640 , 641 , 861
34	MPI_LOCK_EXCLUSIVE, 597 , 851	MPI_MODE_WRONLY, 640 , 641 , 861
35	MPI_LOCK_SHARED, 597 , 598 , 851	MPI_NO_OP, 550 , 577 , 578 , 857 , 1054
36	MPI_LOGICAL, 33 , 228 , 695 , 855	MPI_NULL_COPY_FN, 24 , 363 , 772 , 860
37	MPI_LONG, 34 , 227 , 695 , 854	MPI_NULL_DELETE_FN, 24 , 363 , 772 , 860
38	MPI_LONG_DOUBLE, 34 , 227 , 695 , 854	MPI_OFFSET, 35 , 228 , 695 , 854 , 855 , 1062 ,
39	MPI_LONG_DOUBLE_INT, 231 , 856	1064
40	MPI_LONG_INT, 231 , 856	MPI_OFFSET_KIND, 21 , 35 , 706 , 814 , 834 ,
41	MPI_LONG_LONG, 34 , 227 , 854 , 1064	835 , 852 , 855 , 862
42	MPI_LONG_LONG_INT, 34 , 227 , 695 , 854 , 1064	MPI_OP_NULL, 238 , 858
43	MPI_LOR, 227 , 228 , 857	MPI_ORDER_C, 19 , 142 , 145 , 146 , 862
44	MPI_LXOR, 227 , 228 , 857	MPI_ORDER_FORTRAN, 19 , 142 , 145 , 146 , 862
45	MPI_MAX, 225 , 227 , 228 , 232 , 249 , 857	MPI_PACKED, 16 , 33 , 34 , 46 , 47 , 179 , 181 , 185 ,
46	MPI_MAX_DATAREP_STRING, 653 , 697 , 853	694 , 695 , 847 , 854 , 855
47	MPI_MAX_ERROR_STRING, 460 , 467 , 853	MPI_PROC_NULL, 32 , 36 , 38 , 95 , 99 , 101 , 102 ,
48	MPI_MAX_INFO_KEY, 462 , 473 , 476 , 775 , 776 ,	111 , 194 , 197 , 200 , 202 , 209 , 212 , 226 ,
	853	311 , 406 , 410 , 443 , 557 , 567 , 777 , 778 ,
	MPI_MAX_INFO_VAL, 462 , 473 , 853	851 , 1058 , 1065 , 1067
	MPI_MAX_LIBRARY_VERSION_STRING, 442 ,	MPI_PROD, 227 , 228 , 857
	853 , 1057	"mpi_pset_name", 334 , 335 , 874 , 1048
	MPI_MAX_OBJECT_NAME, 376 – 379 , 853 ,	MPI_REAL, 33 , 46 , 227 , 693 , 695 , 805 , 806 ,
	1059 , 1065	812 , 855

MPI_REAL16, 35 , 228 , 696 , 855	MPI_T_CB_REQUIRE_MPI_RESTRICTED, 751 ,	1
MPI_REAL2, 35 , 228 , 696 , 855	865	2
MPI_REAL4, 35 , 228 , 696 , 805 , 810 , 855	MPI_T_CB_REQUIRE_NONE, 751 , 865	3
MPI_REAL8, 35 , 228 , 696 , 805 , 855 , 1062	MPI_T_CB_REQUIRE_THREAD_SAFE, 751 ,	4
MPI_REPLACE, 574 , 575 , 577 , 578 , 622 , 857 ,	751 , 752 , 865	5
1063 , 1067	MPI_T_CVAR_HANDLE_NULL, 734 , 863	6
MPI_REQUEST_NULL, 78 , 79–81 , 84–88 , 92 ,	MPI_T_ENUM_NULL, 731 , 740 , 754 , 863	7
93 , 630 , 858	MPI_T_ERR_CANNOT_INIT, 769 , 851	8
MPI_ROOT, 194 , 851	MPI_T_ERR_CVAR_SET_NEVER, 735 , 769 , 851	9
MPI_SEEK_CUR, 672 , 680 , 862	MPI_T_ERR_CVAR_SET_NOT_NOW, 735 , 769 ,	10
MPI_SEEK_END, 672 , 680 , 862	851	11
MPI_SEEK_SET, 672 , 680 , 862	MPI_T_ERR_INVALID, 769 , 851 , 1055	12
MPI_SESSION_NULL, 496 , 858 , 1052	MPI_T_ERR_INVALID_HANDLE, 743 , 759 , 769 ,	13
"mpi_shared_memory", 334 , 335 , 874 , 1051	851	14
MPI_SHORT, 34 , 227 , 695 , 854	MPI_T_ERR_INVALID_INDEX, 24 , 751 , 755 ,	15
MPI_SHORT_INT, 231 , 856	769 , 776 , 851 , 1053	16
MPI_SIGNED_CHAR, 34 , 227 , 229 , 695 , 854 ,	MPI_T_ERR_INVALID_ITEM, 24 , 776 , 851 ,	17
1064	1053	18
MPI_SIMILAR, 312 , 321 , 352 , 856	MPI_T_ERR_INVALID_NAME, 732 , 741 , 755 ,	19
"mpi_size", 498 , 501 , 505 , 874 , 1052	766 , 769 , 851 , 1055	20
MPI_SOURCE, 39 , 41 , 251 , 635 , 636 , 838 , 852 ,	MPI_T_ERR_INVALID_SESSION, 769 , 851	21
1053 , 1060	MPI_T_ERR_MEMORY, 769 , 851	22
MPI_STATUS_IGNORE, 10 , 20 , 42 , 43 , 629 ,	MPI_T_ERR_NOT_ACCESSIBLE, 752 , 769 , 851	23
656 , 792 , 814 , 838 , 839 , 847 , 862 , 863 ,	MPI_T_ERR_NOT_INITIALIZED, 769 , 851	24
1058	MPI_T_ERR_NOT_SUPPORTED, 751 , 769 , 851	25
MPI_STATUS_SIZE, 39 , 793 , 852 , 1060	MPI_T_ERR_OUT_OF_HANDLES, 769 , 851	26
MPI_STATUSES_IGNORE, 19 , 20 , 42 , 43 , 629 ,	MPI_T_ERR_OUT_OF_SESSIONS, 769 , 851	27
630 , 814 , 838 , 839 , 862 , 863	MPI_T_ERR_PVAR_NO_ATOMIC, 746 , 769 ,	28
MPI_SUBARRAYS_SUPPORTED, 788 , 789 ,	851	29
792–794 , 796 , 800–803 , 815–817 , 852 ,	MPI_T_ERR_PVAR_NO_STARTSTOP, 743 ,	30
1059	744 , 769 , 851	31
MPI_SUBVERSION, 442 , 863	MPI_T_ERR_PVAR_NO_WRITE, 745 , 769 , 851	32
MPI_SUCCESS, 23 , 25 , 79 , 86 , 88 , 232 , 362 ,	MPI_T_PVAR_ALL_HANDLES, 743 , 744–746 ,	33
363 , 365–368 , 371 , 372 , 388 , 460–462 ,	864	34
468 , 469 , 483 , 523 , 700 , 723 , 732 , 741 ,	MPI_T_PVAR_CLASS_AGGREGATE, 738 , 864	35
743–745 , 751 , 755 , 759 , 766 , 768 , 769 ,	MPI_T_PVAR_CLASS_COUNTER, 738 , 864	36
772 , 849 , 1051 , 1052	MPI_T_PVAR_CLASS_GENERIC, 738 , 864	37
MPI_SUM, 227 , 228 , 574 , 847 , 857	MPI_T_PVAR_CLASS_HIGHWATERMARK,	38
MPI_T_BIND_MPI_COMM, 725 , 864	737 , 864	39
MPI_T_BIND_MPI_DATATYPE, 725 , 864	MPI_T_PVAR_CLASS_LEVEL, 737 , 864	40
MPI_T_BIND_MPI_ERRHANDLER, 725 , 864	MPI_T_PVAR_CLASS_LOWWATERMARK,	41
MPI_T_BIND_MPI_FILE, 725 , 864	738 , 864	42
MPI_T_BIND_MPI_GROUP, 725 , 864	MPI_T_PVAR_CLASS_PERCENTAGE, 737 , 864	43
MPI_T_BIND_MPI_INFO, 725 , 864	MPI_T_PVAR_CLASS_SIZE, 737 , 864	44
MPI_T_BIND_MPI_MESSAGE, 725 , 864	MPI_T_PVAR_CLASS_STATE, 737 , 864	45
MPI_T_BIND_MPI_OP, 725 , 864	MPI_T_PVAR_CLASS_TIMER, 738 , 864	46
MPI_T_BIND_MPI_REQUEST, 725 , 864	MPI_T_PVAR_HANDLE_NULL, 743 , 863	47
MPI_T_BIND_MPI_SESSION, 725 , 864 , 1052	MPI_T_PVAR_SESSION_NULL, 742 , 863	48
MPI_T_BIND_MPI_WIN, 725 , 864	MPI_T_SCOPE_ALL, 732 , 864	
MPI_T_BIND_NO_OBJECT, 725 , 731 , 734 ,	MPI_T_SCOPE_ALL_EQ, 732 , 735 , 864	
740 , 742 , 755 , 756 , 864	MPI_T_SCOPE_CONSTANT, 732 , 864	
MPI_T_CB_REQUIRE_ASYNC_SIGNAL_SAFE,	MPI_T_SCOPE_GROUP, 732 , 864	
751 , 752 , 865	MPI_T_SCOPE_GROUP_EQ, 732 , 735 , 864	
	MPI_T_SCOPE_LOCAL, 732 , 864	

1 MPI_T_SCOPE_READONLY, [732](#), [864](#)
 2 MPI_T_SOURCE_ORDERED, [750](#), [865](#)
 3 MPI_T_SOURCE_UNORDERED, [750](#), [865](#)
 4 MPI_T_VERBOSITY_MPIDEV_ALL, [724](#), [863](#)
 5 MPI_T_VERBOSITY_MPIDEV_BASIC, [724](#),
 6 [863](#)
 7 MPI_T_VERBOSITY_MPIDEV_DETAIL, [724](#),
 8 [863](#)
 9 MPI_T_VERBOSITY_TUNER_ALL, [724](#), [863](#)
 10 MPI_T_VERBOSITY_TUNER_BASIC, [724](#), [863](#)
 11 MPI_T_VERBOSITY_TUNER_DETAIL, [724](#),
 12 [863](#)
 13 MPI_T_VERBOSITY_USER_ALL, [724](#), [863](#)
 14 MPI_T_VERBOSITY_USER_BASIC, [724](#), [863](#)
 15 MPI_T_VERBOSITY_USER_DETAIL, [724](#), [863](#)
 16 MPI_TAG, [39](#), [41](#), [42](#), [251](#), [635](#), [636](#), [838](#), [852](#),
 17 [1053](#), [1060](#)
 18 MPI_TAG_UB, [36](#), [443](#), [843](#), [846](#), [857](#)
 19 MPI_THREAD_FUNNELED, [485](#), [861](#)
 20 "MPI_THREAD_FUNNELED", [495](#), [874](#)
 21 MPI_THREAD_MULTIPLE, [485](#), [486](#), [488](#), [861](#),
 22 [1055](#)
 23 "MPI_THREAD_MULTIPLE", [495](#), [874](#)
 24 MPI_THREAD_SERIALIZED, [485](#), [861](#)
 25 "MPI_THREAD_SERIALIZED", [495](#), [875](#)
 26 MPI_THREAD_SINGLE, [485](#)–[487](#), [495](#), [861](#)
 27 "MPI_THREAD_SINGLE", [495](#), [875](#)
 28 MPI_TYPE_DUP_FN, [371](#), [859](#), [1055](#)
 29 MPI_TYPE_NULL_COPY_FN, [371](#), [859](#), [1055](#)
 30 MPI_TYPE_NULL_DELETE_FN, [371](#), [859](#),
 31 [1055](#), [1060](#)
 32 MPI_TYPECLASS_COMPLEX, [811](#), [862](#)
 33 MPI_TYPECLASS_INTEGER, [811](#), [862](#)
 34 MPI_TYPECLASS_REAL, [811](#), [862](#)
 35 MPI_UB, [4](#), [24](#), [784](#), [1057](#)
 36 MPI_UINT16_T, [34](#), [227](#), [695](#), [854](#), [1062](#), [1064](#)
 37 MPI_UINT32_T, [34](#), [227](#), [695](#), [727](#), [854](#), [1062](#),
 38 [1064](#)
 39 MPI_UINT64_T, [34](#), [227](#), [550](#), [695](#), [727](#), [854](#),
 40 [1062](#), [1064](#)
 41 MPI_UINT8_T, [34](#), [227](#), [695](#), [854](#), [1062](#), [1064](#)
 42 MPI_UNDEFINED, [22](#), [40](#), [63](#), [84](#), [85](#), [88](#), [89](#),
 43 [91](#)–[93](#), [151](#), [154](#), [156](#), [160](#), [161](#), [182](#),
 44 [310](#), [311](#), [329](#), [330](#), [333](#), [397](#), [408](#), [409](#),
 45 [693](#), [697](#), [807](#), [851](#), [1057](#), [1058](#), [1064](#)
 46 MPI_UNEQUAL, [312](#), [321](#), [352](#), [856](#)
 47 MPI_UNIVERSE_SIZE, [519](#), [540](#), [541](#), [860](#)
 48 MPI_UNSIGNED, [34](#), [227](#), [695](#), [727](#), [737](#), [738](#),
 854
 MPI_UNSIGNED_CHAR, [34](#), [227](#), [229](#), [695](#), [854](#)
 MPI_UNSIGNED_LONG, [34](#), [227](#), [695](#), [727](#), [737](#),
 738, [854](#)
 MPI_UNSIGNED_LONG_LONG, [34](#), [227](#), [695](#),
 727, [737](#), [738](#), [854](#), [1064](#)
 MPI_UNSIGNED_SHORT, [34](#), [227](#), [695](#), [854](#)
 MPI_UNWEIGHTED, [20](#), [392](#), [394](#), [396](#), [403](#),
 405, [814](#), [862](#), [1057](#), [1063](#)
 MPI_VAL, [17](#), [836](#)
 MPI_VERSION, [442](#), [863](#)
 MPI_WCHAR, [34](#), [229](#), [379](#), [693](#), [695](#), [854](#), [1064](#)
 MPI_WEIGHTS_EMPTY, [20](#), [392](#), [394](#), [814](#),
 862, [1057](#)
 MPI_WIN_BASE, [563](#), [846](#), [860](#)
 MPI_WIN_CREATE_FLAVOR, [563](#), [860](#)
 MPI_WIN_DISP_UNIT, [563](#), [860](#)
 MPI_WIN_DUP_FN, [367](#), [859](#), [1055](#)
 MPI_WIN_FLAVOR_ALLOCATE, [563](#), [860](#), [1049](#)
 MPI_WIN_FLAVOR_CREATE, [563](#), [860](#), [1049](#)
 MPI_WIN_FLAVOR_DYNAMIC, [564](#), [860](#)
 MPI_WIN_FLAVOR_SHARED, [564](#), [860](#)
 MPI_WIN_MODEL, [563](#), [587](#), [860](#)
 MPI_WIN_NULL, [380](#), [562](#), [858](#), [1048](#)
 "MPI_WIN_NULL", [380](#), [873](#)
 MPI_WIN_NULL_COPY_FN, [367](#), [859](#), [1055](#)
 MPI_WIN_NULL_DELETE_FN, [367](#), [859](#), [1055](#)
 MPI_WIN_SEPARATE, [564](#), [587](#), [607](#), [860](#)
 MPI_WIN_SIZE, [563](#), [860](#)
 MPI_WIN_UNIFIED, [564](#), [587](#), [608](#), [618](#), [860](#)
 MPI_WTIME_IS_GLOBAL, [443](#), [444](#), [471](#), [843](#),
 857
 "mpich://", [445](#)
 "mpix://UNIVERSE", [497](#)
 "native", [690](#)–[692](#), [700](#), [873](#)
 "nb_proc", [649](#), [874](#)
 "no_locks", [550](#), [563](#), [874](#), [1049](#)
 "none", [614](#), [875](#)
 "num_io_nodes", [649](#), [874](#)
 "openmp://", [445](#)
 "openmpi://", [445](#)
 "path", [527](#), [874](#)
 "pmix://", [445](#)
 "random", [648](#), [875](#)
 "rar", [614](#), [875](#)
 "rar,raw,war,waw", [550](#), [614](#)
 "rar,waw", [614](#)
 "raw", [614](#), [875](#)
 "read_mostly", [648](#), [875](#)
 "read_once", [648](#), [875](#)
 "reverse_sequential", [648](#), [875](#)
 "same_disp_unit", [551](#), [874](#), [1055](#)
 "same_op", [550](#), [875](#)
 "same_op_no_op", [550](#), [875](#)
 "same_size", [551](#), [874](#), [1054](#)
 "sequential", [648](#), [875](#)

"slurm://", 445	1
"soft", 484 , 523 , 528 , 874	2
"striping_factor", 649 , 874	3
"striping_unit", 650 , 874	4
"system", 510 , 875	5
	6
"thread_level", 484 , 495 , 874	7
"true", 340 , 341 , 445 , 474 , 563 , 648 , 875 , 1049	8
	9
"war", 614 , 875	10
"waw", 614 , 875	11
"wdir", 484 , 527 , 874	12
"win_allocate", 510 , 875	13
"win_allocate_shared", 510 , 875	14
"write_mostly", 648 , 875	15
"write_once", 648 , 875	16
	17
	18
	19
	20
	21
	22
	23
	24
	25
	26
	27
	28
	29
	30
	31
	32
	33
	34
	35
	36
	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48

MPI Declarations Index

This index refers to declarations needed in C and Fortran, such as address kind integers, handles, etc. The underlined page numbers is the “main” reference (sometimes there are more than one when key concepts are discussed in multiple areas). Fortran defined types are shown as `TYPE(name)`.

- MPI_Aint, [21](#), [21](#), [35](#), [129](#), [132](#), [135](#), [139](#),
[148–150](#), [153–155](#), [164](#), [185–187](#), [548](#),
[552](#), [554](#), [559](#), [560](#), [563](#), [567](#), [569](#), [573](#),
[575](#), [578–581](#), [583](#), [584](#), [692](#), [697](#), [750](#),
[778](#), [779](#), [814](#), [833](#), [843](#), [854](#), [865](#)
C example usage, [172](#)
MPI_Comm, [17](#), [32](#), [312](#), [319–325](#), [327](#), [329](#),
[332](#), [339](#), [341](#), [342](#), [351](#), [352](#), [353](#), [354](#),
[355](#), [361](#), [362](#), [364–366](#), [499](#), [856](#), [858](#),
[865](#)
TYPE(MPI_Comm), [32](#), [327](#), [353](#), [801](#), [856](#),
[858](#), [866](#)
MPI_Count, [21](#), [22](#), [35](#), [749](#), [750](#), [833](#), [854](#), [865](#),
[1057](#)
MPI_Datatype, [129](#), [371](#), [823](#), [854–856](#), [858](#), [865](#)
TYPE(MPI_Datatype), [129](#), [854–856](#), [858](#), [866](#)
MPI_Errhandler, [452](#), [453–459](#), [837](#), [852](#), [858](#),
[865](#)
TYPE(MPI_Errhandler), [452](#), [852](#), [858](#), [866](#)
MPI_F08_status, [838](#), [863](#), [865](#), [1060](#)
MPI_File, [456](#), [457](#), [639](#), [641](#), [643–645](#), [647](#),
[650](#), [652](#), [656–662](#), [664–672](#), [674–680](#),
[682–689](#), [692](#), [702](#), [703](#), [836](#), [858](#), [865](#)
TYPE(MPI_File), [639](#), [858](#), [866](#)
MPI_Fint, [836](#), [863](#), [865](#), [1064](#)
MPI_Group, [310](#), [311–316](#), [318](#), [352](#), [499](#), [564](#),
[593](#), [594](#), [645](#), [836](#), [858](#), [865](#)
TYPE(MPI_Group), [310](#), [858](#), [866](#)
MPI_Info, [446](#), [473](#), [474–477](#), [498](#), [501](#), [520](#),
[523](#), [524](#), [531](#), [535–537](#), [544](#), [565](#), [639](#),
[642](#), [647](#), [650](#), [775](#), [836](#), [837](#), [856](#), [858](#),
[865](#), [1067](#)
TYPE(MPI_Info), [473](#), [856](#), [858](#), [866](#)
MPI_Message, [97](#), [837](#), [852](#), [858](#), [865](#), [1058](#)
TYPE(MPI_Message), [97](#), [852](#), [858](#), [866](#)
MPI_Offset, [21](#), [21](#), [35](#), [643](#), [644](#), [650](#), [652](#),
[656–662](#), [671](#), [672](#), [679](#), [680](#), [682](#), [683](#),
[698](#), [706](#), [750](#), [833](#), [835](#), [854](#), [865](#)
MPI_Op, [224](#), [235](#), [238](#), [240](#), [242](#), [243](#), [245](#),
[247–249](#), [267–272](#), [289](#), [291–294](#), [296](#),
[573](#), [575](#), [578](#), [583](#), [584](#), [836](#), [857](#), [858](#),
[865](#)
TYPE(MPI_Op), [235](#), [857](#), [858](#), [866](#)
MPI_Request, [71–75](#), [77](#), [79](#), [80](#), [81](#), [83–88](#),
[90–93](#), [102–109](#), [116–120](#), [460](#), [627](#),
[631](#), [659–662](#), [668–670](#), [675](#), [676](#), [816](#),
[836](#), [858](#), [865](#)
TYPE(MPI_Request), [79](#), [858](#), [866](#)
MPI_Session, [458](#), [494](#), [837](#), [858](#), [865](#), [1052](#)
TYPE(MPI_Session), [494](#), [858](#), [866](#), [1052](#)
MPI_Status, [36](#), [39](#), [41–44](#), [79](#), [80](#), [83–88](#),
[90–95](#), [97](#), [99](#), [100](#), [103](#), [160](#), [628](#),
[634–636](#), [656–658](#), [664–667](#), [674](#), [677](#),
[678](#), [683–687](#), [689](#), [780](#), [789](#), [791](#),
[838–840](#), [862](#), [865](#), [1053](#), [1057](#), [1060](#)
TYPE(MPI_Status), [36](#), [39](#), [801](#), [838–840](#), [862](#),
[866](#), [1050](#), [1053](#), [1060](#)
MPI_T_cb_safety, [751](#), [757](#), [758](#), [865](#), [1053](#)
MPI_T_cvar_handle, [733](#), [734](#), [735](#), [863](#), [865](#)
MPI_T_enum, [728](#), [729](#), [730](#), [739](#), [753](#), [863](#), [865](#)
MPI_T_event_instance, [758](#), [761–763](#), [865](#), [1053](#)
MPI_T_event_registration, [755](#), [756–758](#), [760](#),
[865](#), [1053](#)
MPI_T_pvar_handle, [742](#), [743–745](#), [863–865](#)
MPI_T_pvar_session, [741](#), [742–745](#), [863](#), [865](#)
MPI_T_source_order, [749](#), [865](#), [1053](#)
MPI_Win, [367–369](#), [454](#), [455](#), [548](#), [552](#), [554](#),
[559](#), [562](#), [564](#), [565](#), [567](#), [569](#), [573](#), [575](#),
[578–581](#), [583](#), [584](#), [590](#), [593–599](#), [601](#),
[602](#), [836](#), [858](#), [865](#)
TYPE(MPI_Win), [548](#), [552](#), [554](#), [559](#), [858](#), [866](#)

MPI Callback Function Prototype Index

This index lists the C typedef names for callback routines, such as those used with attribute caching or user-defined reduction operations. Fortran example prototypes are given near the text of the C name.

COMM_COPY_ATTR_FUNCTION, [23](#), [24](#), [361](#),
[362](#), [790](#), [859](#), [871](#)
COMM_DELETE_ATTR_FUNCTION, [24](#), [361](#),
[362](#), [859](#), [871](#)
COMM_ERRHANDLER_FUNCTION, [452](#), [453](#),
[871](#)
COPY_FUNCTION, [24](#), [771](#), [771](#), [772](#), [860](#), [872](#)
DATAREP_CONVERSION_FUNCTION, [696](#),
[698](#), [859](#), [872](#)
DATAREP_EXTENT_FUNCTION, [696](#), [697](#), [872](#)
DELETE_FUNCTION, [24](#), [771](#), [772](#), [860](#), [872](#)
FILE_ERRHANDLER_FUNCTION, [456](#), [871](#)
GREQUEST_CANCEL_FUNCTION, [628](#), [630](#),
[872](#)
GREQUEST_FREE_FUNCTION, [628](#), [629](#), [872](#)
GREQUEST_QUERY_FUNCTION, [628](#), [629](#),
[872](#)
MPI_Comm_copy_attr_function, [23](#), [24](#), [361](#),
[361](#), [362](#), [790](#), [859](#), [866](#), [868](#)
MPI_Comm_delete_attr_function, [24](#), [361](#), [362](#),
[859](#), [866](#), [868](#)
MPI_Comm_errhandler_fn, [774](#), [1063](#)
MPI_Comm_errhandler_function, [24](#), [452](#), [452](#),
[453](#), [774](#), [784](#), [867](#), [869](#), [1063](#)
MPI_COMM_NULL_COPY_FN, [859](#)
MPI_Copy_function, [24](#), [771](#), [860](#), [872](#)
MPI_Datarep_conversion_function, [694](#),
[697](#), [698](#), [833](#), [859](#), [867](#), [870](#)
MPI_Datarep_conversion_function_c, [694](#),
[697](#), [698](#), [833](#), [859](#), [867](#), [870](#), [1050](#)
MPI_Datarep_extent_function, [694](#), [697](#), [867](#),
[870](#)
MPI_Delete_function, [24](#), [771](#), [772](#), [860](#), [872](#)
MPI_File_errhandler_fn, [774](#), [1063](#)
MPI_File_errhandler_function, [456](#), [774](#), [867](#),
[869](#), [1063](#)
MPI_Grequest_cancel_function, [628](#), [630](#), [867](#),
[870](#)
MPI_Grequest_free_function, [628](#), [629](#), [867](#), [869](#)
MPI_Grequest_query_function, [628](#), [628](#), [629](#),
[867](#), [869](#)
MPI_Handler_function, [24](#), [784](#), [1057](#)
MPI_Session_errhandler_function, [457](#), [458](#), [867](#),
[869](#), [1052](#)
MPI_T_event_cb_function, [757](#), [758](#), [867](#), [1053](#)
MPI_T_event_dropped_cb_function, [760](#), [867](#),
[1053](#)
MPI_T_event_free_cb_function, [759](#), [867](#), [1053](#)
MPI_Type_copy_attr_function, [370](#), [371](#), [859](#),
[866](#), [869](#)
MPI_Type_delete_attr_function, [370](#), [371](#), [859](#),
[867](#), [869](#), [1060](#)
MPI_User_function, [235](#), [236](#), [240](#), [833](#), [866](#), [867](#)
MPI_User_function_c, [235](#), [236](#), [833](#), [866](#), [868](#),
[1050](#)
MPI_Win_copy_attr_function, [367](#), [367](#), [368](#),
[859](#), [866](#), [868](#)
MPI_Win_delete_attr_function, [367](#), [367](#), [368](#),
[859](#), [866](#), [868](#)
MPI_Win_errhandler_fn, [774](#), [1063](#)
MPI_Win_errhandler_function, [454](#), [774](#), [867](#),
[869](#), [1063](#)
SESSION_ERRHANDLER_FUNCTION, [458](#),
[871](#), [1052](#)
TYPE_COPY_ATTR_FUNCTION, [370](#), [371](#),
[859](#), [871](#)
TYPE_DELETE_ATTR_FUNCTION, [370](#), [372](#),
[859](#), [871](#)
USER_FUNCTION, [235](#), [236](#), [870](#)
WIN_COPY_ATTR_FUNCTION, [367](#), [368](#), [859](#),
[871](#)
WIN_DELETE_ATTR_FUNCTION, [367](#), [368](#),
[859](#), [871](#)
WIN_ERRHANDLER_FUNCTION, [454](#), [455](#), [871](#)

MPI Function Index

The underlined page numbers refer to the function definitions.

- MPI_ABORT, [237](#), [450](#), [483](#), [489](#), [506](#), [508](#), [543](#),
[727](#), [836](#), [1066](#)
- MPI_ACCUMULATE, [547](#), [566](#), [573](#), [574](#), [575](#),
[577](#), [584](#), [588](#), [614](#), [621](#), [622](#), [882](#),
[1063](#), [1067](#)
 - C example usage, [624](#)
 - Fortran example usage, [574](#)
 - Language-independent example usage, [621](#)
- MPI_Accumulate_c, [573](#), [1050](#)
- MPI_ADD_ERROR_CLASS, [464](#), [465](#), [507](#)
- MPI_ADD_ERROR_CODE, [466](#), [507](#)
- MPI_ADD_ERROR_STRING, [467](#), [507](#)
- MPI_ADDRESS, [24](#), [783](#), [800](#), [1057](#)
- MPI_AINT_ADD, [25](#), [148](#), [150](#), [560](#), [1055](#)
 - C example usage, [624](#)
- MPI_AINT_DIFF, [25](#), [148](#), [150](#), [560](#), [1055](#)
- MPI_ALLGATHER, [189](#), [193](#), [194](#), [214](#), [215](#),
[217](#), [219](#), [261](#), [877](#), [880](#)
 - C example usage, [217](#)
- MPI_Allgather_c, [214](#), [1050](#)
- MPI_ALLGATHER_INIT, [189](#), [193](#), [194](#), [282](#),
[880](#), [1051](#)
 - MPI_Allgather_init_c, [283](#), [1050](#)
- MPI_ALLGATHERV, [189](#), [193](#), [194](#), [216](#), [217](#),
[262](#), [880](#)
 - MPI_Allgatherv_c, [216](#), [1050](#)
- MPI_ALLGATHERV_INIT, [189](#), [193](#), [194](#), [284](#),
[880](#), [1051](#)
 - MPI_Allgatherv_init_c, [284](#), [1050](#)
- MPI_ALLOC_MEM, [446](#), [447](#), [448](#), [462](#), [508](#),
[510](#), [551–556](#), [558](#), [561](#), [569](#), [600](#),
[800–802](#), [814](#), [1051](#), [1054](#), [1061](#)
 - C example usage, [449](#), [624](#)
 - Fortran example usage, [448](#)
- MPI_ALLOC_MEM_CPTR, [447](#), [1054](#)
- MPI_ALLREDUCE, [189](#), [192–194](#), [227](#), [235](#),
[240](#), [241](#), [269](#), [877](#), [880](#), [1064](#)
 - Fortran example usage, [241](#)
- MPI_Allreduce_c, [240](#), [1050](#)
- MPI_ALLREDUCE_INIT, [189](#), [193](#), [194](#), [291](#),
[880](#), [1051](#)
 - MPI_Allreduce_init_c, [291](#), [1050](#)
- MPI_ALLTOALL, [189](#), [193](#), [194](#), [218](#), [219](#), [221](#),
[222](#), [263](#), [877](#), [880](#), [1062](#)
 - C example usage, [301](#)
- MPI_Alltoall_c, [218](#), [1050](#)
- MPI_ALLTOALL_INIT, [189](#), [193](#), [194](#), [285](#), [880](#),
[1051](#)
 - MPI_Alltoall_init_c, [285](#), [1050](#)
- MPI_ALLTOALLV, [189](#), [193](#), [194](#), [220](#), [221](#), [222](#),
[224](#), [265](#), [880](#), [1062](#)
 - MPI_Alltoallv_c, [220](#), [1050](#)
- MPI_ALLTOALLV_INIT, [189](#), [193](#), [194](#), [286](#),
[880](#), [1051](#)
 - MPI_Alltoallv_init_c, [287](#), [1050](#)
- MPI_ALLTOALLW, [189](#), [193](#), [194](#), [222](#), [223](#),
[224](#), [266](#), [880](#), [1062](#)
 - MPI_Alltoallw_c, [223](#), [1050](#)
- MPI_ALLTOALLW_INIT, [189](#), [193](#), [194](#), [288](#),
[880](#), [1051](#)
 - MPI_Alltoallw_init_c, [288](#), [1050](#)
- MPI_ASSERT_MEMORY_ALLOC_KINDS
 - C example usage, [510](#)
- MPI_ATTR_DELETE, [24](#), [374](#), [772](#), [774](#)
- MPI_ATTR_GET, [24](#), [374](#), [773](#), [843](#)
 - Fortran example usage, [844](#), [845](#)
- MPI_ATTR_PUT, [24](#), [374](#), [773](#), [843](#), [846](#)
 - Fortran example usage, [844](#)
- MPI_BARRIER, [189](#), [193](#), [194](#), [196](#), [253](#), [506](#),
[611](#), [612](#), [708](#), [876](#), [880](#)
 - C example usage, [490](#), [620](#)
 - Fortran example usage, [505](#)
 - Language-independent example usage,
[610–613](#), [621](#), [622](#)
- MPI_BARRIER_INIT, [189](#), [193](#), [194](#), [275](#), [880](#),
[1051](#)
- MPI_BCAST, [15](#), [189](#), [193](#), [194](#), [196](#), [197](#), [226](#),
[254](#), [300](#), [880](#)
 - C example usage, [197](#), [297](#), [298](#), [300](#), [344](#)
- MPI_Bcast_c, [197](#), [1050](#)
- MPI_BCAST_INIT, [15](#), [189](#), [193](#), [194](#), [276](#), [880](#),
[1051](#)
 - MPI_Bcast_init_c, [276](#), [1050](#)
- MPI_BSEND, [15](#), [28](#), [51](#), [52](#), [68](#), [617](#), [876](#), [879](#)
 - C example usage, [617](#)
 - Fortran example usage, [55](#), [56](#)
- MPI_Bsend_c, [51](#), [1050](#)
- MPI_BSEND_INIT, [105](#), [109](#), [879](#)
 - MPI_Bsend_init_c, [106](#), [1050](#)
- MPI_BUFFER_ATTACH, [60](#), [67](#), [80](#), [617](#), [879](#),

- 1048
- C example usage, 66, 490, 617
- MPI_Buffer_attach_c, 60, 1050
- MPI_BUFFER_DETACH, 62, 67, 617, 879, 1048, 1050, 1060
- C example usage, 66, 617
- MPI_Buffer_detach_c, 62, 1050
- MPI_BUFFER_FLUSH, 64, 67, 1048
- MPI_BUFFER_IFLUSH, 65, 1048
- MPI_CANCEL, 24, 55, 80, 94, 102, 103, 104, 251, 275, 490, 580, 627, 630, 631, 774, 879, 1051
- C example usage, 490
- MPI_CART_COORDS, 385, 400, 401, 1046, 1066
- Fortran example usage, 406
- MPI_CART_CREATE, 350, 385, 386, 387–389, 409, 410, 815, 881, 1054, 1065
- Fortran example usage, 815
- MPI_CART_GET, 385, 399, 1046, 1065
- Fortran example usage, 437
- MPI_CART_MAP, 385, 408, 409, 1059
- MPI_CART_RANK, 385, 399, 400, 1066
- Fortran example usage, 406
- MPI_CART_SHIFT, 385, 405, 406, 410, 1066
- Fortran example usage, 406, 437
- Language-independent example usage, 417
- MPI_CART_SUB, 385, 407, 409, 881, 1066
- Language-independent example usage, 407
- MPI_CARTDIM_GET, 385, 398, 399, 1065
- Language-independent example usage, 417
- MPI_CLOSE_PORT, 532, 533, 536
- MPI_COMM_ACCEPT, 531, 533, 534, 535, 541, 542, 881
- C example usage, 538, 539
- MPI_COMM_ATTACH_BUFFER, 58, 67, 80, 1048
- MPI_Comm_attach_buffer_c, 58
- MPI_COMM_BUFFER_ATTACH
- C example usage, 66
- MPI_COMM_BUFFER_DETACH
- C example usage, 66
- MPI_COMM_C2F, 836
- MPI_COMM_CALL_ERRHANDLER, 468, 470
- MPI_COMM_COMPARE, 321, 352
- MPI_COMM_CONNECT, 463, 534, 535, 541, 542, 881
- C example usage, 538, 539
- MPI_COMM_CREATE, 319, 321, 325, 326, 328–330, 385, 880, 1063
- C example usage, 327, 344, 345, 347
- MPI_COMM_CREATE_ERRHANDLER, 24, 451, 452, 453, 783, 869, 871, 1060
- MPI_COMM_CREATE_FROM_GROUP, 319, 321, 338, 339, 355, 443, 1046, 1052
- C example usage, 501
- Fortran example usage, 505
- MPI_COMM_CREATE_GROUP, 319, 321, 327, 328, 330, 339, 880, 1058
- MPI_COMM_CREATE_KEYVAL, 24, 360, 361, 363, 373, 771, 842, 843, 868, 870, 1060, 1065
- C example usage, 374
- MPI_COMM_DELETE_ATTR, 24, 360, 363–365, 366, 374, 774
- MPI_COMM_DETACH_BUFFER, 61, 67, 1048
- MPI_Comm_detach_buffer_c, 61
- MPI_COMM_DISCONNECT, 110, 363, 374, 489, 496, 497, 506, 524, 543, 544, 877, 880, 1045
- MPI_COMM_DUP, 312, 319, 321, 322, 323, 324, 326, 353, 356, 360, 362, 366, 374, 381, 771, 785, 880, 1046, 1051, 1058
- C example usage, 346
- MPI_COMM_DUP_FN, 24, 362, 363, 364, 794, 859, 1055, 1060
- MPI_COMM_DUP_WITH_INFO, 319, 321, 323, 324, 340, 360, 362, 366, 374, 785, 880, 1058
- MPI_COMM_F2C, 836
- MPI_COMM_FLUSH_BUFFER, 63, 67, 1048
- MPI_COMM_FREE, 110, 319, 322, 339, 353, 356, 363, 364, 366, 374, 489, 492, 496, 506, 524, 543, 544, 772, 877, 880, 1048
- MPI_COMM_FREE_KEYVAL, 24, 360, 364, 374, 772
- MPI_COMM_GET_ATTR, 24, 360, 365, 374, 442, 773, 796, 843, 846
- C example usage, 374, 843
- Fortran example usage, 844, 845
- MPI_COMM_GET_ERRHANDLER, 24, 451, 453, 783, 1066
- MPI_COMM_GET_INFO, 340, 342, 509, 510, 1051, 1058
- MPI_COMM_GET_NAME, 377, 378, 1048, 1065
- MPI_COMM_GET_PARENT, 377, 483, 489, 493, 521, 524
- C example usage, 528
- MPI_COMM_GROUP, 19, 309, 312, 313, 319, 320, 352, 452, 1066
- C example usage, 327, 347, 374
- MPI_COMM_IDUP, 319, 321, 323, 324, 350, 360, 362, 366, 374, 785, 880, 1051, 1053, 1058
- MPI_COMM_IDUP_WITH_INFO, 319, 321, 324, 340, 360, 362, 366, 374, 785, 1051
- MPI_COMM_IFLUSH_BUFFER, 65, 1048

- 1 MPI_COMM_JOIN, [493](#), [544](#), [545](#), [546](#)
2 MPI_COMM_MULTIPLE, [881](#)
3 MPI_COMM_NULL_COPY_FN, [23](#), [24](#), [362](#),
4 [363](#), [364](#), [790](#), [842](#), [859](#), [1055](#), [1060](#)
5 MPI_COMM_NULL_DELETE_FN, [24](#), [363](#), [364](#),
6 [859](#), [1055](#)
7 MPI_COMM_RANK, [15](#), [320](#), [352](#), [798](#)
8 C example usage, [31](#)
9 MPI_COMM_RANK_F08, [798](#)
10 MPI_COMM_REMOTE_GROUP, [352](#)
11 MPI_COMM_REMOTE_SIZE, [352](#), [353](#)
12 C example usage, [330](#)
13 MPI_COMM_SET_ATTR, [24](#), [360](#), [363](#), [364](#),
14 [374](#), [773](#), [796](#), [843](#), [846](#)
15 C example usage, [374](#), [843](#)
16 Fortran example usage, [845](#)
17 MPI_COMM_SET_ERRHANDLER, [24](#), [451](#),
18 [453](#), [783](#)
19 MPI_COMM_SET_INFO, [275](#), [340](#), [341](#), [342](#),
20 [880](#), [1051](#), [1054](#), [1058](#)
21 MPI_COMM_SET_NAME, [376](#)
22 MPI_COMM_SIZE, [319](#), [320](#), [352](#), [621](#)
23 MPI_COMM_SPAWN, [483](#), [484](#), [508](#), [509](#), [513](#),
24 [518](#), [519](#), [520](#), [521–526](#), [528](#), [541](#), [542](#),
25 [881](#), [1052](#)
26 C example usage, [522](#), [528](#)
27 Fortran example usage, [522](#)
28 MPI_COMM_SPAWN_MULTIPLE, [483](#), [484](#),
29 [508](#), [509](#), [513](#), [518](#), [519](#), [524](#), [525](#), [526](#),
30 [542](#), [1052](#)
31 C example usage, [527](#)
32 Fortran example usage, [527](#)
33 MPI_COMM_SPLIT, [319](#), [321](#), [325](#), [326](#), [329](#),
34 [330](#), [331](#), [381](#), [385](#), [387](#), [389](#), [407](#), [409](#),
35 [880](#), [1063](#)
36 C example usage, [330](#), [356](#), [358](#)
37 MPI_COMM_SPLIT_TYPE, [319](#), [321](#), [332](#), [335](#),
38 [337](#), [340](#), [446](#), [551](#), [553](#), [558](#), [617](#), [880](#),
39 [1048](#), [1051](#), [1058](#)
40 C example usage, [334](#), [336](#)
41 MPI_COMM_TEST_INTER, [350](#), [351](#)
42 MPI_COMPARE_AND_SWAP, [547](#), [566](#), [579](#),
43 [621](#), [882](#)
44 C example usage, [624](#)
45 Language-independent example usage, [622](#)
46 MPI_CONVERSION_FN_NULL, [700](#), [859](#), [1055](#)
47 MPI_CONVERSION_FN_NULL_C, [700](#), [833](#),
48 [859](#), [1050](#)
MPI_DIMS_CREATE, [385](#), [387](#), [388](#), [1051](#)
Fortran example usage, [437](#)
Language-independent example usage, [388](#)
MPI_DIST_GRAPH_CREATE, [340](#), [385](#), [390](#),
[392](#), [393–396](#), [405](#), [410](#), [881](#), [1063](#)
C example usage, [396](#)
Language-independent example usage, [395](#)
MPI_DIST_GRAPH_CREATE_ADJACENT, [340](#),
[385](#), [390](#), [391](#), [392](#), [395](#), [396](#), [405](#), [410](#),
[881](#), [1059](#), [1063](#)
Language-independent example usage, [395](#)
MPI_DIST_GRAPH_NEIGHBORS, [385](#), [403](#),
[404](#), [405](#), [410](#), [1059](#), [1063](#)
MPI_DIST_GRAPH_NEIGHBORS_COUNT,
[385](#), [403](#), [405](#), [1056](#), [1063](#)
MPI_DUP_FN, [24](#), [363](#), [772](#), [860](#)
MPI_ERRHANDLER_C2F, [507](#), [837](#)
MPI_ERRHANDLER_CREATE, [24](#), [783](#), [1057](#),
[1060](#)
MPI_ERRHANDLER_F2C, [507](#), [837](#)
MPI_ERRHANDLER_FREE, [451](#), [459](#), [489](#), [496](#),
[507](#), [1066](#)
MPI_ERRHANDLER_GET, [24](#), [783](#), [1057](#), [1066](#)
MPI_ERRHANDLER_SET, [24](#), [783](#), [1057](#)
MPI_ERROR_CLASS, [461](#), [464](#), [507](#)
MPI_ERROR_STRING, [460](#), [461](#), [464](#), [467](#), [507](#)
MPI_EXSCAN, [190](#), [193](#), [227](#), [235](#), [248](#), [249](#),
[273](#), [880](#), [1062](#)
MPI_Exscan_c, [248](#), [1050](#)
MPI_EXSCAN_INIT, [190](#), [193](#), [296](#), [880](#), [1051](#)
MPI_Exscan_init_c, [296](#), [1050](#)
MPI_F_SYNC_REG, [149](#), [623](#), [788](#), [804](#), [805](#),
[824–828](#), [830](#), [1061](#)
Fortran example usage, [827](#)
Language-independent example usage, [623](#)
MPI_FENCE, [618](#)
MPI_FETCH_AND_OP, [547](#), [566](#), [574](#), [577](#),
[578](#), [882](#)
MPI_FILE
C example usage, [709](#)
MPI_FILE_C2F, [836](#)
MPI_FILE_CALL_ERRHANDLER, [469](#), [470](#)
MPI_FILE_CLOSE, [489](#), [496](#), [544](#), [639](#), [640](#),
[641](#), [642](#), [883](#)
Fortran example usage, [665](#), [668](#)
MPI_FILE_CREATE_ERRHANDLER, [451](#), [452](#),
[456](#), [457](#), [869](#), [871](#), [1060](#)
MPI_FILE_DELETE, [641](#), [642](#), [647](#), [649](#), [711](#),
[883](#)
MPI_FILE_F2C, [836](#)
MPI_FILE_GET_AMODE, [645](#), [646](#)
Fortran example usage, [646](#)
MPI_FILE_GET_ATOMICITY, [703](#)
MPI_FILE_GET_BYTE_OFFSET, [664](#), [672](#), [673](#),
[680](#)
MPI_FILE_GET_ERRHANDLER, [451](#), [457](#), [711](#),
[1066](#)
MPI_FILE_GET_GROUP, [645](#)
MPI_FILE_GET_INFO, [509](#), [646](#), [647](#), [648](#), [649](#),
[1051](#), [1067](#)

MPI_FILE_GET_POSITION, 672	
MPI_FILE_GET_POSITION_SHARED, 679 , 680 , 704	
MPI_FILE_GET_SIZE, 644 , 645 , 706	
MPI_FILE_GET_TYPE_EXTENT, 691 , 692 , 700 , 1050	
MPI_File_get_type_extent_c, 692 , 1050	
MPI_FILE_GET_VIEW, 652 , 653 , 883	
MPI_FILE_IREAD, 653 , 668 , 681 , 701 , 883	
Fortran example usage, 668	
MPI_File_iread_c, 668 , 1050	
MPI_FILE_IREAD_ALL, 653 , 669 , 670 , 883 , 1055	
MPI_File_iread_all_c, 669 , 1050	
MPI_FILE_IREAD_AT, 653 , 660 , 883	
C example usage, 709	
MPI_File_iread_at_c, 660 , 1050	
MPI_FILE_IREAD_AT_ALL, 653 , 661 , 883 , 1055	
MPI_File_iread_at_all_c, 661 , 1050	
MPI_FILE_IREAD_SHARED, 653 , 675 , 676 , 883	
MPI_File_iread_shared_c, 676 , 1050	
MPI_FILE_IWRITE, 653 , 670 , 883	
MPI_File_iwrite_c, 670 , 1050	
MPI_FILE_IWRITE_ALL, 653 , 671 , 883 , 1055	
MPI_File_iwrite_all_c, 671 , 1050	
MPI_FILE_IWRITE_AT, 653 , 662 , 883	
C example usage, 709	
MPI_File_iwrite_at_c, 662 , 1050	
MPI_FILE_IWRITE_AT_ALL, 653 , 663 , 883 , 1055	
MPI_File_iwrite_at_all_c, 663 , 1050	
MPI_FILE_IWRITE_SHARED, 653 , 676 , 677 , 883	
MPI_File_iwrite_shared_c, 676 , 1050	
MPI_FILE_OPEN, 462 , 517 , 639 , 640 , 641 , 647 , 649 , 651 , 673 , 705 , 706 , 711 , 712 , 883	
Fortran example usage, 665 , 668	
MPI_FILE_PREALLOCATE, 643 , 644 , 701 , 706 , 707 , 883	
MPI_FILE_READ, 653 , 664 , 666 , 668 , 706 , 707 , 883	
Fortran example usage, 665	
MPI_File_read_c, 664 , 1050	
MPI_FILE_READ_ALL, 653 , 665 , 666 , 670 , 681 , 682 , 883	
C example usage, 682	
MPI_File_read_all_c, 665 , 1050	
MPI_FILE_READ_ALL_BEGIN, 15 , 653 , 681 , 682 , 685 , 701 , 830 , 883	
C example usage, 682 , 709	
MPI_File_read_all_begin_c, 685 , 1050	
MPI_FILE_READ_ALL_END, 653 , 681 , 682 , 685 , 701 , 830 , 883	
C example usage, 682 , 709	
MPI_FILE_READ_AT, 653 , 656 , 657 , 658 , 660 , 883	1
C example usage, 707 , 708	2
MPI_File_read_at_c, 656 , 1050	3
MPI_FILE_READ_AT_ALL, 653 , 657 , 658 , 661 , 883	4
MPI_File_read_at_all_c, 657 , 1050	5
MPI_FILE_READ_AT_ALL_BEGIN, 15 , 653 , 682 , 830 , 883	6
MPI_File_read_at_all_begin_c, 682 , 1050	7
MPI_FILE_READ_AT_ALL_END, 653 , 683 , 830 , 883	8
MPI_FILE_READ_ORDERED, 653 , 678 , 883	9
MPI_File_read_ordered_c, 678 , 1050	10
MPI_FILE_READ_ORDERED_BEGIN, 15 , 653 , 687 , 830 , 883	11
MPI_File_read_ordered_begin_c, 687 , 1050	12
MPI_FILE_READ_ORDERED_END, 653 , 688 , 830 , 883	13
MPI_FILE_READ_SHARED, 653 , 674 , 676 , 678 , 883	14
MPI_File_read_shared_c, 674 , 1050	15
MPI_FILE_SEEK, 671 , 672 , 883	16
MPI_FILE_SEEK_SHARED, 679 , 680 , 704 , 883	17
MPI_FILE_SET_ATOMICITY, 641 , 702 , 703 , 883	18
C example usage, 707 , 709	19
MPI_FILE_SET_ERRHANDLER, 451 , 456 , 711	20
MPI_FILE_SET_INFO, 647 , 649 , 883 , 1051 , 1067	21
MPI_FILE_SET_SIZE, 643 , 644 , 701 , 704 , 706 , 707 , 883	22
MPI_FILE_SET_VIEW, 143 , 463 , 640 , 647 , 649 , 650 , 651 , 652 , 672 , 680 , 690 , 697 , 705 , 706 , 712 , 883 , 1067	23
C example usage, 707 , 708	24
Fortran example usage, 665 , 668 , 812	25
MPI_FILE_SYNC, 642 , 654 , 701 , 702 , 703 , 704 , 709 , 883	26
C example usage, 708	27
MPI_FILE_WRITE, 653 , 654 , 666 , 667 , 668 , 670 , 706 , 883	28
MPI_File_write_c, 666 , 1050	29
MPI_FILE_WRITE_ALL, 653 , 667 , 668 , 671 , 883	30
MPI_File_write_all_c, 667 , 1050	31
MPI_FILE_WRITE_ALL_BEGIN, 15 , 653 , 686 , 816 , 830 , 883	32
C example usage, 709 , 712	33
MPI_File_write_all_begin_c, 686 , 1050	34
MPI_FILE_WRITE_ALL_END, 653 , 686 , 830 , 883	35
C example usage, 709 , 712	36
MPI_FILE_WRITE_AT, 653 , 654 , 658 , 659 , 662 , 883	37

- 1 C example usage, [707](#), [708](#)
2 MPI_File_write_at_c, [658](#), [1050](#)
3 MPI_FILE_WRITE_AT_ALL, [653](#), [659](#), [663](#), [883](#)
4 MPI_File_write_at_all_c, [659](#), [1050](#)
5 MPI_FILE_WRITE_AT_ALL_BEGIN, [15](#), [653](#),
6 [683](#), [830](#), [883](#)
7 MPI_File_write_at_all_begin_c, [684](#), [1050](#)
8 MPI_FILE_WRITE_AT_ALL_END, [653](#), [684](#),
9 [830](#), [883](#)
10 MPI_FILE_WRITE_ORDERED, [653](#), [677](#), [678](#),
11 [679](#), [883](#)
12 MPI_File_write_ordered_c, [679](#), [1050](#)
13 MPI_FILE_WRITE_ORDERED_BEGIN, [15](#), [653](#),
14 [688](#), [830](#), [883](#)
15 MPI_File_write_ordered_begin_c, [688](#), [1050](#)
16 MPI_FILE_WRITE_ORDERED_END, [653](#), [689](#),
17 [830](#), [883](#)
18 MPI_FILE_WRITE_SHARED, [653](#), [654](#), [675](#),
19 [677](#), [679](#), [883](#)
20 MPI_File_write_shared_c, [675](#), [1050](#)
21 MPI_FINALIZE, [20](#), [26](#), [27](#), [29](#), [443](#), [450](#), [481](#),
22 [488](#), [489–493](#), [496](#), [516](#), [543](#), [544](#), [640](#),
23 [723](#), [733](#), [746](#), [748](#), [835](#), [838](#), [839](#), [881](#),
24 [1045](#), [1046](#), [1052](#), [1059](#), [1066](#)
25 C example usage, [489–491](#)
26 MPI_FINALIZED, [491](#), [492](#), [507](#), [835](#), [1055](#)
27 MPI_FNIALIZE, [110](#)
28 MPI_FREE_MEM, [447](#), [448](#), [462](#), [553](#), [555](#)
29 C example usage, [624](#)
30 Fortran example usage, [448](#)
31 MPI_GATHER, [189](#), [192–194](#), [198](#), [201](#), [202](#),
32 [209](#), [215](#), [226](#), [256](#), [880](#)
33 C example usage, [184](#), [202](#), [207](#)
34 MPI_Gather_c, [198](#), [1050](#)
35 MPI_GATHER_INIT, [189](#), [193](#), [194](#), [277](#), [880](#),
36 [1051](#)
37 MPI_Gather_init_c, [277](#), [1050](#)
38 MPI_GATHERV, [189](#), [193](#), [194](#), [200](#), [201–203](#),
39 [211](#), [217](#), [257](#), [880](#)
40 C example usage, [184](#), [203–207](#)
41 MPI_Gatherv_c, [200](#), [1050](#)
42 MPI_GATHERV_INIT, [189](#), [193](#), [194](#), [278](#), [880](#),
43 [1051](#)
44 MPI_Gatherv_init_c, [278](#), [1050](#)
45 MPI_GET, [547](#), [566](#), [569](#), [570](#), [575](#), [577](#), [583](#),
46 [588](#), [606](#), [611](#), [622](#), [827](#), [882](#), [1067](#)
47 C example usage, [619](#), [620](#)
48 Fortran example usage, [570](#), [572](#)
Language-independent example usage,
[610](#), [612](#)
MPI_Get_c, [570](#), [1050](#)
MPI_GET_ACCUMULATE, [547](#), [566](#), [574](#), [575](#),
[577](#), [578](#), [586](#), [614](#), [621](#), [882](#), [1054](#)
C example usage, [624](#)
Language-independent example usage, [621](#)
MPI_Get_accumulate_c, [576](#), [1050](#)
MPI_GET_ADDRESS, [24](#), [129](#), [148](#), [149](#), [150](#),
[151](#), [162](#), [560](#), [783](#), [800](#), [818](#), [823](#), [841](#),
[842](#)
C example usage, [172](#), [175](#), [176](#), [183](#)
Fortran example usage, [149](#), [819](#), [841](#), [847](#)
MPI_GET_COUNT, [39](#), [40](#), [41](#), [79](#), [160](#), [161](#),
[580](#), [634](#), [656](#), [1057](#), [1064](#)
Fortran example usage, [161](#)
MPI_Get_count_c, [40](#), [1050](#)
MPI_GET_ELEMENTS, [79](#), [160](#), [161](#), [634](#), [635](#),
[656](#), [780](#), [1057](#)
Fortran example usage, [161](#)
MPI_Get_elements_c, [160](#), [1050](#)
MPI_Get_elements_c, [24](#)
MPI_GET_ELEMENTS_X, [24](#), [780](#), [1047](#), [1057](#)
MPI_GET_HW_RESOURCE_INFO, [445](#), [446](#),
[1048](#)
MPI_GET_LIBRARY_VERSION, [442](#), [507](#), [1054](#),
[1055](#), [1057](#)
MPI_GET_PROCESSOR_NAME, [444](#), [445](#), [1066](#)
MPI_GET_VERSION, [441](#), [442](#), [507](#), [803](#), [1055](#)
MPI_GRAPH_CREATE, [385](#), [388](#), [389](#), [390](#), [395](#),
[398](#), [402](#), [409](#), [410](#), [881](#), [1065](#)
Language-independent example usage,
[389](#), [402](#)
MPI_GRAPH_GET, [385](#), [398](#)
MPI_GRAPH_MAP, [385](#), [409](#)
MPI_GRAPH_NEIGHBORS, [385](#), [401](#), [402](#), [410](#),
[1063](#)
Language-independent example usage, [402](#)
MPI_GRAPH_NEIGHBORS_COUNT, [385](#), [401](#),
[402](#), [1063](#)
Language-independent example usage, [402](#)
MPI_GRAPHDIMS_GET, [385](#), [397](#), [398](#)
MPI_GREQUEST_COMPLETE, [627](#), [629](#), [630](#),
[631](#)
C example usage, [631](#)
MPI_GREQUEST_START, [628](#), [869](#), [872](#), [1063](#)
C example usage, [631](#)
MPI_GROUP_C2F, [836](#)
MPI_GROUP_COMPARE, [311](#), [315](#)
MPI_GROUP_DIFFERENCE, [314](#)
MPI_GROUP_EXCL, [315](#), [317](#)
C example usage, [344](#)
MPI_GROUP_F2C, [836](#)
MPI_GROUP_FREE, [318](#), [319](#), [320](#), [452](#), [489](#),
[496](#), [1066](#)
C example usage, [327](#), [344](#), [345](#), [347](#)
MPI_GROUP_FROM_SESSION_PSET, [318](#),
[1052](#)
C example usage, [501](#), [503](#)
Fortran example usage, [505](#)

MPI_GROUP_INCL, 314 , 315 , 316			
C example usage, 327 , 345 , 347			
MPI_GROUP_INTERSECTION, 313			
MPI_GROUP_RANGE_EXCL, 317			
MPI_GROUP_RANGE_INCL, 316			
MPI_GROUP_RANK, 310 , 320			
MPI_GROUP_SIZE, 310 , 320			
MPI_GROUP_TRANSLATE_RANKS, 311 , 1065			
MPI_GROUP_UNION, 313			
MPI_IALLGATHER, 189 , 193 , 194 , 260 , 880			
MPI_Iallgather_c, 260 , 1050			
MPI_IALLGATHERV, 189 , 193 , 194 , 261 , 880			
MPI_Iallgather_v_c, 262 , 1050			
MPI_IALLREDUCE, 189 , 193 , 194 , 268 , 880			
C example usage, 303			
MPI_Iallreduce_c, 268 , 1050			
MPI_IALLTOALL, 189 , 193 , 194 , 262 , 880			
C example usage, 301			
MPI_Ialltoall_c, 263 , 1050			
MPI_IALLTOALLV, 189 , 193 , 194 , 264 , 880			
MPI_Ialltoall_v_c, 264 , 1050			
MPI_IALLTOALLW, 189 , 193 , 194 , 265 , 880			
MPI_Ialltoall_w_c, 266 , 1050			
MPI_IBARRIER, 189 , 193 , 194 , 251 , 253 , 300 , 876 , 880			
C example usage, 300 – 302			
MPI_IBCAST, 15 , 189 , 193 , 194 , 253 , 254 , 304 , 876 , 880			
C example usage, 254 , 302 , 303			
MPI_Ibcast_c, 254 , 1050			
MPI_IBSEND, 72 , 79 , 109 , 879			
MPI_Ibsend_c, 72 , 1050			
MPI_IXSCAN, 190 , 193 , 272 , 880			
MPI_Ixscan_c, 273 , 1050			
MPI_IGATHER, 189 , 193 , 194 , 255 , 880			
MPI_Igather_c, 255 , 1050			
MPI_IGATHERV, 189 , 193 , 194 , 256 , 880			
MPI_Igather_v_c, 256 , 1050			
MPI_IMPROBE, 15 , 28 , 94 , 97 , 98 , 99 , 102 , 517 , 876 , 879 , 1053 , 1058			
MPI_IMRECV, 97 – 100 , 101 , 102 , 489 , 496 , 876 , 879 , 1058			
MPI_Imrecv_c, 101 , 1050			
MPI_INEIGHBOR_ALLGATHER, 386 , 422 , 881 , 1059			
MPI_Ineighbor_allgather_c, 423 , 1050			
MPI_INEIGHBOR_ALLGATHERV, 386 , 423 , 881 , 1059			
MPI_Ineighbor_allgather_v_c, 424 , 1050			
MPI_INEIGHBOR_ALLTOALL, 386 , 425 , 881 , 1059			
MPI_Ineighbor_alltoall_c, 425 , 1050			
MPI_INEIGHBOR_ALLTOALLV, 386 , 426 , 881 , 1059			
MPI_Ineighbor_alltoall_v_c, 426 , 1050			
MPI_INEIGHBOR_ALLTOALLW, 386 , 427 , 881 , 1059			
MPI_Ineighbor_alltoall_w_c, 428 , 1050			
MPI_INFO_C2F, 507 , 837			
MPI_INFO_CREATE, 474 , 507			
C example usage, 501			
MPI_INFO_CREATE_ENV, 478 , 507 , 1052			
MPI_INFO_DELETE, 462 , 475 , 477 , 507			
MPI_INFO_DUP, 477 , 507			
MPI_INFO_ENV			
Language-independent example usage, 484			
MPI_INFO_F2C, 507 , 837			
MPI_INFO_FREE, 343 , 445 , 478 , 489 , 496 , 501 , 507 , 566 , 648 , 750 , 754 , 757 , 759			
MPI_INFO_GET, 24 , 775 , 1052 , 1067			
MPI_INFO_GET_NKEYS, 473 , 476 , 477 , 507 , 1067			
MPI_INFO_GET_NTHKEY, 473 , 477 , 507 , 1067			
MPI_INFO_GET_STRING, 24 , 473 , 475 , 476 , 507 , 775 , 776 , 1052			
MPI_INFO_GET_VALUELEN, 24 , 776 , 1052 , 1067			
MPI_INFO_SET, 474 , 475 , 477 , 507 , 775			
C example usage, 501			
MPI_INIT, 20 , 26 , 29 , 309 , 443 , 450 , 478 , 481 , 482 , 483 , 485 – 489 , 491 – 493 , 506 , 516 , 518 , 520 – 522 , 524 , 540 , 541 , 723 , 726 , 733 , 746 , 835 , 838 , 839 , 876 , 881 , 1046 , 1052 , 1057 , 1059 , 1061 , 1063			
C example usage, 31 , 483			
MPI_INIT_THREAD, 26 , 309 , 450 , 478 , 481 , 482 , 485 , 486 – 488 , 491 – 493 , 495 , 516 , 518 , 541 , 723 , 726 , 733 , 835 , 881 , 1046 , 1057 , 1059 , 1063			
C example usage, 31 , 483			
MPI_INITIALIZED, 491 , 492 , 507 , 835 , 1055			
MPI_INTERCOMM_CREATE, 319 , 322 , 328 , 353 , 354 , 355 , 356 , 880 , 1058			
C example usage, 356 , 358			
MPI_INTERCOMM_CREATE_FROM_GROUPS, 319 , 321 , 322 , 353 , 354 , 355 , 443 , 1046 , 1052			
MPI_INTERCOMM_MERGE, 321 , 328 , 350 , 353 , 356 , 880 , 1060			
MPI_IPROBE, 15 , 28 , 40 , 94 , 95 – 99 , 102 , 517 , 876 , 879 , 1058			
C example usage, 490			
MPI_IRECV, 15 , 28 , 75 , 102 , 817 , 818 , 821 , 822 , 879			
C example usage, 302			
Fortran example usage, 81 – 83 , 89 , 90 , 816 – 818			
MPI_Irecv_c, 75 , 1050			
MPI_IREDUCE, 189 , 193 , 194 , 267 , 268 , 880			

- 1 MPI_Ireduce_c, [267](#), [1050](#)
2 MPI_IREDUCE_SCATTER, [190](#), [193](#), [194](#), [270](#),
3 [880](#)
4 MPI_Ireduce_scatter_c, [271](#), [1050](#)
5 MPI_IREDUCE_SCATTER_BLOCK, [189](#), [193](#),
6 [194](#), [269](#), [880](#)
7 MPI_Ireduce_scatter_block_c, [269](#), [1050](#)
8 MPI_IRSEND, [74](#), [879](#)
9 MPI_Irsend_c, [74](#), [1050](#)
10 MPI_IS_THREAD_MAIN, [485](#), [487](#), [1055](#)
11 MPI_ISCAN, [190](#), [193](#), [271](#), [880](#)
12 MPI_Iscan_c, [272](#), [1050](#)
13 MPI_ISCATTER, [189](#), [193](#), [194](#), [257](#), [880](#)
14 MPI_Iscatter_c, [258](#), [1050](#)
15 MPI_ISCATTERV, [189](#), [193](#), [194](#), [259](#), [880](#)
16 MPI_Iscatterv_c, [259](#), [1050](#)
17 MPI_ISEND, [15](#), [28](#), [71](#), [109](#), [795](#), [797](#), [799](#),
18 [816](#), [817](#), [822](#), [879](#)
19 Fortran example usage, [81](#), [82](#), [89](#), [90](#), [816](#)
20 MPI_Isend_c, [71](#), [1050](#)
21 MPI_ISENDRECV, [76](#), [879](#), [1050](#)
22 MPI_Isendrecv_c, [76](#), [1050](#)
23 MPI_ISENDRECV_REPLACE, [77](#), [879](#), [1050](#)
24 MPI_Isendrecv_replace_c, [77](#), [1050](#)
25 MPI_ISSEND, [73](#), [879](#)
26 MPI_Issend_c, [73](#), [1050](#)
27 MPI_KEYVAL_CREATE, [24](#), [771](#), [773](#), [872](#)
28 MPI_KEYVAL_FREE, [24](#), [374](#), [772](#)
29 MPI_LOOKUP_NAME, [462](#), [531](#), [536](#), [537](#), [538](#)
30 MPI_MEMORY_ALLOC_KINDS
31 C example usage, [510](#)
32 MPI_MESSAGE_C2F, [837](#), [1058](#)
33 MPI_MESSAGE_F2C, [837](#), [1058](#)
34 MPI_MPROBE, [15](#), [94](#), [97](#), [99](#), [102](#), [517](#), [876](#),
35 [879](#), [1058](#)
36 MPI_MRECV, [15](#), [97–99](#), [100](#), [101](#), [102](#), [489](#),
37 [496](#), [876](#), [879](#), [1058](#)
38 MPI_Mrecv_c, [100](#), [1050](#)
39 MPI_NEIGHBOR_ALLGATHER, [386](#), [411](#), [412](#),
40 [413](#), [415](#), [423](#), [881](#), [1049](#), [1059](#)
41 Language-independent example usage, [412](#)
42 MPI_Neighbor_allgather_c, [411](#), [1050](#)
43 MPI_NEIGHBOR_ALLGATHER_INIT, [429](#), [881](#),
44 [1051](#)
45 MPI_Neighbor_allgather_init_c, [430](#), [1050](#)
46 MPI_NEIGHBOR_ALLGATHERV, [386](#), [413](#), [425](#),
47 [881](#), [1049](#), [1059](#)
48 MPI_Neighbor_allgatherv_c, [414](#), [1050](#)
MPI_NEIGHBOR_ALLGATHERV_INIT, [430](#),
[795](#), [881](#), [1051](#), [1053](#)
MPI_Neighbor_allgatherv_init_c, [431](#), [1050](#)
MPI_NEIGHBOR_ALLTOALL, [386](#), [415](#), [417](#),
[426](#), [881](#), [1049](#), [1059](#)
Language-independent example usage, [417](#)
MPI_Neighbor_alltoall_c, [415](#), [1050](#)
MPI_NEIGHBOR_ALLTOALL_INIT, [432](#), [881](#),
[1051](#)
MPI_Neighbor_alltoall_init_c, [432](#), [1050](#)
MPI_NEIGHBOR_ALLTOALLV, [386](#), [418](#), [427](#),
[881](#), [1049](#), [1059](#)
MPI_Neighbor_alltoallv_c, [418](#), [1050](#)
MPI_NEIGHBOR_ALLTOALLV_INIT, [433](#), [795](#),
[881](#), [1051](#), [1053](#)
MPI_Neighbor_alltoallv_init_c, [434](#), [1050](#)
MPI_NEIGHBOR_ALLTOALLW, [386](#), [420](#), [429](#),
[881](#), [1049](#), [1059](#)
MPI_Neighbor_alltoallw_c, [421](#), [1050](#)
MPI_NEIGHBOR_ALLTOALLW_INIT, [435](#), [795](#),
[881](#), [1051](#), [1053](#)
MPI_Neighbor_alltoallw_init_c, [435](#), [1050](#)
MPI_NULL_COPY_FN, [24](#), [363](#), [772](#), [860](#)
MPI_NULL_DELETE_FN, [24](#), [363](#), [772](#), [860](#)
MPI_OP_C2F, [836](#)
MPI_OP_COMMUTATIVE, [243](#), [1062](#)
MPI_OP_CREATE, [235](#), [238](#), [794](#), [867](#), [870](#),
[1060](#)
C example usage, [239](#), [249](#)
Fortran example usage, [240](#)
MPI_Op_create_c, [235](#), [834](#), [867](#), [1050](#)
MPI_OP_F2C, [836](#)
MPI_OP_FREE, [238](#), [489](#), [496](#)
MPI_OPEN_PORT, [531](#), [532](#), [534–538](#)
C example usage, [538](#), [539](#)
MPI_PACK, [69](#), [178](#), [181–183](#), [185](#), [694](#), [698](#),
[699](#)
C example usage, [183](#), [184](#)
MPI_Pack_c, [179](#), [1050](#)
MPI_PACK_EXTERNAL, [8](#), [185](#), [810](#), [1064](#)
MPI_Pack_external_c, [186](#), [1050](#)
MPI_PACK_EXTERNAL_SIZE, [188](#)
MPI_Pack_external_size_c, [188](#), [1050](#)
MPI_PACK_SIZE, [68](#), [182](#), [1057](#)
C example usage, [184](#)
MPI_Pack_size_c, [182](#), [1050](#)
MPI_PARRIVED, [28](#), [114](#), [120](#), [121](#), [877](#), [879](#),
[1051](#)
C example usage, [125](#)
MPI_PCONTROL, [718](#), [719](#)
MPI_PREADY, [114–116](#), [118](#), [119–121](#), [877](#),
[879](#), [1051](#)
C example usage, [114](#), [122](#), [123](#), [125](#)
MPI_PREADY_LIST, [119](#), [120](#), [1051](#)
MPI_PREADY_RANGE, [119](#), [120](#), [1051](#)
MPI_PRECV_INIT, [114](#), [117](#), [118](#), [122](#), [879](#),
[1051](#)
C example usage, [114](#), [122](#), [123](#), [125](#)
MPI_PROBE, [15](#), [38](#), [40](#), [43](#), [94](#), [95](#), [96](#), [97](#), [99](#),
[101](#), [102](#), [517](#), [876](#), [879](#), [1058](#)

Fortran example usage, 96	712, 870, 872	1
MPI_PSEND_INIT, 114 , 116 , 117–119 , 122 , 879 , 1051	MPI_Register_datarep_c, 694 , 700 , 834 , 870 , 1050	2
C example usage, 114 , 122 , 123 , 125		3
MPI_PUBLISH_NAME, 531 , 535 , 536–538	MPI_REMOVE_ERROR_CLASS, 465 , 507 , 1049	4
C example usage, 538	MPI_REMOVE_ERROR_CODE, 466 , 507 , 1049	5
MPI_PUT, 547 , 566 , 567 , 570 , 574 , 581 , 588 , 605 , 606 , 612 , 622 , 816 , 827 , 882 , 1049 , 1067	MPI_REMOVE_ERROR_STRING, 467 , 507 , 1049	6
C example usage, 594 , 600 , 618 , 619	MPI_REQUEST_C2F, 836	7
Fortran example usage, 827	MPI_REQUEST_F2C, 836	8
Language-independent example usage, 611 , 613 , 618	MPI_REQUEST_FREE, 81 , 82 , 102 , 110 , 251 , 274 , 275 , 489 , 496 , 544 , 580 , 629–631 , 876 , 877 , 879–881 , 1062	9
MPI_Put_c, 567 , 1050	C example usage, 489	12
MPI_QUERY_THREAD, 487 , 488 , 516 , 1055	Fortran example usage, 82	13
MPI_RACCUMULATE, 547 , 566 , 574 , 577 , 583 , 584 , 882	MPI_REQUEST_GET_STATUS, 28 , 43 , 90 , 91 , 92 , 629 , 879 , 1047 , 1062	14
MPI_Raccumulate_c, 583 , 1050	MPI_REQUEST_GET_STATUS_ALL, 92 , 93 , 94 , 1048	15
MPI_RECV, 15 , 28 , 32 , 37 , 39–43 , 95 , 97 , 98 , 101 , 128 , 159 , 160 , 181 , 190 , 199 , 301 , 617 , 635 , 708 , 746 , 823 , 827 , 828 , 879	MPI_REQUEST_GET_STATUS_ANY, 91 , 92 , 94 , 1048	17
C example usage, 31 , 301 , 617	MPI_REQUEST_GET_STATUS_SOME, 93 , 94 , 1048	19
Fortran example usage, 47 , 48 , 55–57 , 83 , 96 , 159	MPI_RGET, 547 , 566 , 581 , 583 , 882	21
MPI_Recv_c, 37 , 1050	C example usage, 623	22
MPI_RECV_INIT, 15 , 108 , 109 , 118 , 879	MPI_Rget_c, 582 , 1050	23
MPI_Recv_init_c, 108 , 1050	MPI_RGET_ACCUMULATE, 547 , 566 , 574 , 577 , 584 , 586 , 882	24
MPI_REDUCE, 189 , 193 , 194 , 224 , 225–227 , 235–238 , 241 , 244 , 246 , 248 , 249 , 268 , 574 , 577 , 578 , 880 , 1063	MPI_Rget_accumulate_c, 585 , 1050	25
C example usage, 233 , 234 , 239 , 346	MPI_RPUT, 547 , 566 , 580 , 581 , 882	26
Fortran example usage, 228 , 229 , 233 , 240	C example usage, 623	27
MPI_Reduce_c, 225 , 1050	MPI_Rput_c, 580 , 1050	28
MPI_REDUCE_INIT, 189 , 193 , 194 , 289 , 880 , 1051	MPI_RSEND, 15 , 28 , 53 , 876 , 879	29
MPI_Reduce_init_c, 290 , 1050	MPI_Rsend_c, 53 , 1050	30
MPI_REDUCE_LOCAL, 226 , 227 , 235 , 242 , 1060 , 1062	MPI_RSEND_INIT, 107 , 879	31
MPI_Reduce_local_c, 242 , 1050	MPI_Rsend_init_c, 107 , 1050	31
MPI_REDUCE_SCATTER, 189 , 193 , 194 , 227 , 235 , 245 , 246 , 271 , 877 , 880	MPI_SCAN, 190 , 193 , 227 , 235 , 247 , 249 , 272 , 880	32
MPI_Reduce_scatter_c, 245 , 1050	C example usage, 249	34
MPI_REDUCE_SCATTER_BLOCK, 189 , 193 , 194 , 227 , 235 , 243 , 244 , 245 , 270 , 877 , 880 , 1062	MPI_Scan_c, 247 , 1050	35
MPI_Reduce_scatter_block_c, 244 , 1050	MPI_SCAN_INIT, 190 , 193 , 294 , 880 , 1051	36
MPI_REDUCE_SCATTER_BLOCK_INIT, 189 , 193 , 194 , 292 , 880 , 1051	MPI_Scan_init_c, 295 , 1050	37
MPI_Reduce_scatter_block_init_c, 292 , 1050	MPI_SCATTER, 189 , 193 , 194 , 208 , 209 , 211 , 212 , 244 , 258 , 880	38
MPI_REDUCE_SCATTER_INIT, 190 , 193 , 194 , 293 , 880 , 1051	C example usage, 212	39
MPI_Reduce_scatter_init_c, 294 , 1050	MPI_Scatter_c, 208 , 1050	40
MPI_REGISTER_DATAREP, 462 , 694 , 697–700 , 712 , 870 , 872	MPI_SCATTER_INIT, 189 , 193 , 194 , 279 , 880 , 1051	41
	MPI_Scatter_init_c, 280 , 1050	42
	MPI_SCATTERV, 189 , 193 , 194 , 210 , 211 , 212 , 246 , 260 , 880	43
	C example usage, 212 , 213	45
	MPI_Scatterv_c, 210 , 1050	46
	MPI_SCATTERV_INIT, 189 , 193 , 194 , 281 , 880 , 1051	47
		48

- 1 MPI_Scatterv_init_c, [281](#), [1050](#)
2 MPI_SEND, [15](#), [28](#), [31](#), [32](#), [33](#), [41](#), [47](#), [128](#), [158](#),
3 [159](#), [179](#), [301](#), [640](#), [708](#), [720](#), [823](#), [824](#),
4 [827](#), [879](#)
5 C example usage, [31](#), [172](#), [175](#), [176](#), [183](#),
6 [301](#), [302](#), [720](#)
7 Fortran example usage, [47](#), [48](#), [57](#), [83](#), [96](#),
8 [159](#), [815](#), [818](#)
9 MPI_Send_c, [32](#), [1050](#)
10 MPI_SEND_INIT, [15](#), [104](#), [109](#), [117](#), [879](#)
11 MPI_Send_init_c, [105](#), [1050](#)
12 MPI_SENDF, [43](#), [111](#), [405](#), [879](#)
13 Fortran example usage, [170](#)–[172](#)
14 Language-independent example usage, [417](#)
15 MPI_Sendrecv_c, [44](#), [1050](#)
16 MPI_SENDF_REPLACE, [45](#), [111](#), [879](#)
17 Fortran example usage, [406](#)
18 MPI_Sendrecv_replace_c, [45](#), [1050](#)
19 MPI_SESSION_ATTACH_BUFFER, [59](#), [67](#), [80](#),
20 [1048](#)
21 MPI_Session_attach_buffer_c, [59](#)
22 MPI_SESSION_C2F, [837](#), [1052](#)
23 MPI_SESSION_CALL_ERRHANDLER, [469](#), [470](#),
24 [507](#), [1052](#)
25 MPI_SESSION_CREATE_ERRHANDLER, [451](#),
26 [452](#), [457](#), [458](#), [495](#), [507](#), [869](#), [871](#), [1052](#)
27 MPI_SESSION_DETACH_BUFFER, [62](#), [67](#), [1048](#)
28 MPI_Session_detach_buffer_c, [62](#)
29 MPI_SESSION_F2C, [837](#), [1052](#)
30 MPI_SESSION_FINALIZE, [27](#), [110](#), [495](#), [496](#),
31 [497](#), [506](#), [835](#), [881](#), [1045](#), [1046](#), [1052](#)
32 C example usage, [501](#), [503](#)
33 Fortran example usage, [505](#)
34 Language-independent example usage, [497](#)
35 MPI_SESSION_FLUSH_BUFFER, [64](#), [67](#), [1048](#)
36 MPI_SESSION_GET_ERRHANDLER, [451](#), [459](#),
37 [1052](#)
38 MPI_SESSION_GET_INFO, [495](#), [500](#), [509](#), [1052](#)
39 MPI_SESSION_GET_NTH_PSET, [318](#), [499](#),
40 [500](#), [1052](#)
41 C example usage, [503](#)
42 Fortran example usage, [505](#)
43 MPI_SESSION_GET_NUM_PSETS, [498](#), [499](#),
44 [1052](#)
45 C example usage, [503](#)
46 Fortran example usage, [505](#)
47 MPI_SESSION_GET_PSET_INFO, [501](#), [1052](#)
48 MPI_SESSION_IFLUSH_BUFFER, [65](#), [1048](#)
MPI_SESSION_INIT, [451](#), [489](#), [494](#), [495](#), [506](#),
[508](#), [509](#), [516](#), [541](#), [726](#), [835](#), [881](#),
[1046](#), [1052](#)
C example usage, [501](#), [503](#)
Fortran example usage, [505](#)
MPI_SESSION_SET_ERRHANDLER, [451](#), [458](#),
[1052](#)
MPI_SIZEOF, [24](#), [777](#), [1053](#)
MPI_SSEND, [28](#), [52](#), [879](#)
Fortran example usage, [56](#), [83](#)
MPI_Ssend_c, [52](#), [1050](#)
MPI_SSEND_INIT, [106](#), [879](#)
MPI_Ssend_init_c, [106](#), [1050](#)
MPI_START, [109](#), [110](#), [113](#), [116](#), [119](#), [120](#), [274](#),
[823](#), [876](#), [877](#), [879](#)–[881](#)
MPI_STARTALL, [109](#), [110](#), [113](#), [116](#), [119](#), [120](#),
[274](#), [823](#), [876](#), [879](#)–[881](#)
MPI_STATUS_C2F, [838](#)
MPI_STATUS_C2F08, [839](#), [1060](#)
MPI_STATUS_F082C, [839](#), [1060](#)
MPI_STATUS_F082F, [840](#), [1050](#), [1060](#)
MPI_STATUS_F2C, [838](#)
MPI_STATUS_F2F08, [839](#), [1050](#), [1060](#)
MPI_STATUS_GET_ERROR, [42](#), [1047](#)
MPI_STATUS_GET_SOURCE, [41](#), [1047](#)
MPI_STATUS_GET_TAG, [41](#), [1047](#)
MPI_STATUS_SET_CANCELLED, [635](#), [1053](#)
MPI_STATUS_SET_ELEMENTS, [634](#), [781](#), [1046](#)
MPI_Status_set_elements_c, [634](#)
MPI_Status_set_elements_c, [24](#)
MPI_STATUS_SET_ELEMENTS_X, [24](#), [780](#),
[1047](#), [1057](#)
MPI_STATUS_SET_ERROR, [636](#), [1048](#)
MPI_STATUS_SET_SOURCE, [635](#), [1047](#)
MPI_STATUS_SET_TAG, [636](#), [1047](#)
MPI_T_CATEGORY_CHANGED, [768](#), [1053](#)
MPI_T_CATEGORY_GET_CATEGORIES, [767](#),
[768](#), [769](#)
MPI_T_CATEGORY_GET_CVARS, [766](#), [768](#),
[769](#)
MPI_T_CATEGORY_GET_EVENTS, [767](#), [768](#),
[1053](#)
MPI_T_CATEGORY_GET_INDEX, [766](#), [769](#),
[1055](#)
MPI_T_CATEGORY_GET_INFO, [765](#), [768](#), [769](#),
[1054](#), [1055](#)
MPI_T_CATEGORY_GET_NUM, [764](#)
MPI_T_CATEGORY_GET_NUM_EVENTS, [765](#),
[766](#), [1053](#)
MPI_T_CATEGORY_GET_PVARS, [767](#), [768](#),
[769](#)
MPI_T_CVAR_GET_INDEX, [732](#), [769](#), [1055](#)
MPI_T_CVAR_GET_INFO, [728](#), [730](#), [731](#), [734](#),
[735](#), [768](#), [769](#), [1054](#), [1055](#)
C example usage, [732](#)
MPI_T_CVAR_GET_NUM, [730](#), [734](#)
MPI_T_CVAR_HANDLE_ALLOC, [728](#), [733](#), [734](#),
[735](#), [769](#)
C example usage, [736](#)

MPI_T_CVAR_HANDLE_FREE, 734 , 769	C example usage, 736		
MPI_T_CVAR_READ, 735 , 769	C example usage, 736		
MPI_T_CVAR_WRITE, 735 , 769 , 1047			
MPI_T_ENUM_GET_INFO, 728 , 769			
MPI_T_ENUM_GET_ITEM, 729 , 769			
MPI_T_EVENT_CALLBACK_GET_INFO, 758 , 759 , 1053			
MPI_T_EVENT_CALLBACK_SET_INFO, 758 , 1053			
MPI_T_EVENT_COPY, 752 , 754 , 762 , 1053			
MPI_T_EVENT_GET_INDEX, 755 , 769 , 1053			
MPI_T_EVENT_GET_INFO, 728 , 753 , 754 , 756 , 762 , 768 , 769 , 1053			
MPI_T_EVENT_GET_NUM, 753 , 1053			
MPI_T_EVENT_GET_SOURCE, 752 , 763 , 1053			
MPI_T_EVENT_GET_TIMESTAMP, 750 , 752 , 763 , 1053			
MPI_T_EVENT_HANDLE_ALLOC, 756 , 758 , 769 , 1053			
MPI_T_EVENT_HANDLE_FREE, 759 , 769 , 1053			
MPI_T_EVENT_HANDLE_GET_INFO, 756 , 757 , 1053			
MPI_T_EVENT_HANDLE_SET_INFO, 756 , 1053			
MPI_T_EVENT_READ, 752 , 754 , 762 , 1053			
MPI_T_EVENT_REGISTER_CALLBACK, 757 , 758 , 1053			
MPI_T_EVENT_SET_DROPPED_HANDLER, 760 , 761 , 1053			
MPI_T_FINALIZE, 727	C example usage, 746		
MPI_T_INIT_THREAD, 726 , 727	C example usage, 746		
MPI_T_PVAR_GET_INDEX, 741 , 769 , 1055			
MPI_T_PVAR_GET_INFO, 728 , 738 , 739 , 740 , 742–746 , 768 , 769 , 1054 , 1055	C example usage, 746		
MPI_T_PVAR_GET_NUM, 739 , 743			
MPI_T_PVAR_HANDLE_ALLOC, 728 , 742 , 744 , 745 , 769 , 1047	C example usage, 746		
MPI_T_PVAR_HANDLE_FREE, 743 , 769 , 1047 , 1055	C example usage, 746		
MPI_T_PVAR_READ, 744 , 745 , 752 , 769 , 1055	C example usage, 746		
MPI_T_PVAR_READRESET, 740 , 745 , 746 , 752 , 769 , 1055			
MPI_T_PVAR_RESET, 745 , 752 , 769 , 1047 , 1054 , 1055			
MPI_T_PVAR_SESSION_CREATE, 741 , 769	C example usage, 746		1
	MPI_T_PVAR_SESSION_FREE, 742 , 769		2
	MPI_T_PVAR_START, 743 , 752 , 769 , 1047 , 1054 , 1055		3
	C example usage, 746		4
	MPI_T_PVAR_STOP, 744 , 752 , 769 , 1047 , 1054 , 1055		5
	MPI_T_PVAR_WRITE, 744 , 745 , 752 , 769 , 1047 , 1055		6
	MPI_T_SOURCE_GET_INFO, 749 , 750 , 764 , 769 , 1053		7
	MPI_T_SOURCE_GET_NUM, 749 , 769 , 1053		8
	MPI_T_SOURCE_GET_TIMESTAMP, 751 , 752 , 769 , 1053		9
	MPI_TEST, 28 , 43 , 78 , 79 , 80 , 81–83 , 85 , 90 , 91 , 93 , 102 , 103 , 110 , 113 , 114 , 120 , 121 , 275 , 489 , 496 , 544 , 630 , 631 , 654 , 655 , 876 , 877 , 879–882		10
	MPI_TEST_CANCELLED, 79 , 80 , 103 , 104 , 629 , 635 , 656 , 877 , 879		11
	C example usage, 490		12
	MPI_TESTALL, 28 , 83 , 86 , 87 , 489 , 496 , 516 , 544 , 629 , 630 , 633		13
	MPI_TESTANY, 28 , 83 , 84 , 85 , 89 , 489 , 496 , 516 , 544 , 629 , 630 , 633		14
	MPI_TESTSOME, 28 , 83 , 88 , 89 , 93 , 94 , 489 , 496 , 516 , 544 , 629 , 630 , 633		15
	MPI_TOPO_TEST, 385 , 397		16
	MPI_TYPE_C2F, 836		17
	MPI_TYPE_COMMIT, 156 , 157 , 837 , 1046		18
	C example usage, 172 , 175 , 176 , 183 , 202–207 , 213 , 249		19
	Fortran example usage, 157 , 170–172 , 570 , 819		20
	MPI_TYPE_CONTIGUOUS, 16 , 129 , 132 , 152 , 164 , 638 , 692		21
	C example usage, 202		22
	Fortran example usage, 159 , 161		23
	Language-independent example usage, 130 , 152		24
	MPI_Type_contiguous_c, 129 , 1050		25
	MPI_TYPE_CREATE_DARRAY, 16 , 41 , 144 , 145 , 164		26
	Fortran example usage, 148		27
	MPI_Type_create_darray_c, 144 , 1050		28
	MPI_TYPE_CREATE_F90_COMPLEX, 16 , 164 , 167 , 228 , 694 , 788 , 807 , 809		29
	MPI_TYPE_CREATE_F90_INTEGER, 16 , 164 , 167 , 227 , 694 , 788 , 808 , 809		30
	Fortran example usage, 808		31
	MPI_TYPE_CREATE_F90_REAL, 16 , 164 , 167 , 227 , 694 , 788 , 806 , 807–809 , 1062		32
	Fortran example usage, 808		33
	MPI_TYPE_CREATE_HINDEXED, 17 , 24 , 129 ,		34
			35
			36
			37
			38
			39
			40
			41
			42
			43
			44
			45
			46
			47
			48

- 1 135, 136, 139, 140, 164, 783, 833
2 MPI_Type_create_hindexed_c, 136, 1050
3 MPI_TYPE_CREATE_HINDEXED_BLOCK, 17,
4 129, 138, 164, 833, 1058
5 MPI_Type_create_hindexed_block_c, 138,
6 1050
7 MPI_TYPE_CREATE_HVECTOR, 17, 24, 129,
8 132, 164, 783
9 C example usage, 172, 175
10 Fortran example usage, 170, 171
11 MPI_Type_create_hvector_c, 132, 1050
12 MPI_TYPE_CREATE_INDEXED_BLOCK, 16,
13 137, 138, 164
14 Fortran example usage, 570
15 MPI_Type_create_indexed_block_c, 137,
16 1050
17 MPI_TYPE_CREATE_KEYVAL, 360, 370, 373,
18 843, 868, 871, 1065
19 MPI_TYPE_CREATE_RESIZED, 24, 129, 152,
20 154, 155, 164, 692, 784, 1060
21 C example usage, 176
22 Fortran example usage, 819
23 MPI_Type_create_resized_c, 154, 1050
24 MPI_TYPE_CREATE_STRUCT, 17, 24, 129,
25 139, 140, 152, 164, 223, 783, 833
26 C example usage, 172, 175, 183, 205, 207,
27 249
28 Fortran example usage, 172, 819, 847
29 Language-independent example usage,
30 140, 152
31 MPI_Type_create_struct_c, 139, 1050
32 MPI_TYPE_CREATE_SUBARRAY, 16, 19, 141,
33 143, 145, 164
34 C example usage, 714
35 Fortran example usage, 715
36 MPI_Type_create_subarray_c, 141, 1050
37 MPI_TYPE_DELETE_ATTR, 360, 373, 374,
38 1060
39 MPI_TYPE_DUP, 16, 158, 164, 1060
40 MPI_TYPE_DUP_FN, 371, 859, 1055
41 MPI_TYPE_EXTENT, 24, 783, 1057
42 Fortran example usage, 570
43 MPI_TYPE_F2C, 836
44 MPI_TYPE_FREE, 157, 166, 372, 489, 496
45 Fortran example usage, 570
46 MPI_TYPE_FREE_KEYVAL, 360, 372, 374
47 MPI_TYPE_GET_ATTR, 360, 373, 374, 796,
48 843, 1060
49 MPI_TYPE_GET_CONTENTS, 163, 164, 165,
50 166–168, 833
51 C example usage, 177
52 MPI_Type_get_contents_c, 165, 1050
53 MPI_TYPE_GET_ENVELOPE, 163, 164, 166,
54 167, 232, 808, 833
55 C example usage, 177
56 MPI_Type_get_envelope_c, 163, 1050
57 MPI_TYPE_GET_EXTENT, 24, 153, 156, 779,
58 783, 811, 840
59 C example usage, 172
60 Fortran example usage, 170–172, 572, 574
61 MPI_Type_get_extent_c, 154, 1050
62 MPI_Type_get_extent_c, 24
63 MPI_TYPE_GET_EXTENT_X, 24, 779, 1047,
64 1057
65 MPI_TYPE_GET_NAME, 379, 1048, 1060
66 MPI_TYPE_GET_TRUE_EXTENT, 155, 780
67 MPI_Type_get_true_extent_c, 155, 1050
68 MPI_Type_get_true_extent_c, 24
69 MPI_TYPE_GET_TRUE_EXTENT_X, 24, 779,
70 1047, 1057
71 MPI_TYPE_GET_VALUE_INDEX, 164, 231,
72 232, 1048
73 C example usage, 232
74 MPI_TYPE_HINDEXED, 24, 783, 1057
75 MPI_TYPE_HVECTOR, 24, 783, 1057
76 MPI_TYPE_INDEXED, 16, 133, 134, 135, 137,
77 164
78 C example usage, 172, 175
79 Fortran example usage, 171
80 Language-independent example usage, 134
81 MPI_Type_indexed_c, 134, 1050
82 MPI_TYPE_LB, 24, 783, 1057
83 MPI_TYPE_MATCH_SIZE, 788, 811, 1060
84 C example usage, 811
85 Fortran example usage, 812
86 MPI_TYPE_NULL_COPY_FN, 371, 859, 1055
87 MPI_TYPE_NULL_DELETE_FN, 371, 859,
88 1055, 1060
89 MPI_TYPE_SET_ATTR, 360, 372, 374, 796,
90 843, 846, 1060
91 MPI_TYPE_SET_NAME, 378, 1060
92 MPI_TYPE_SIZE, 151, 778, 1057
93 C example usage, 720
94 MPI_Type_size_c, 151, 1050
95 MPI_Type_size_c, 24
96 MPI_TYPE_SIZE_X, 24, 778, 1047, 1057
97 MPI_TYPE_STRUCT, 24, 783, 1057
98 MPI_TYPE_UB, 24, 783, 1057
99 MPI_TYPE_VECTOR, 16, 130, 131, 132, 135,
100 164
101 C example usage, 203, 204, 206, 213
102 Fortran example usage, 170–172
103 Language-independent example usage, 131
104 MPI_Type_vector_c, 131, 1050
105 MPI_UNPACK, 180, 181, 185, 698
106 C example usage, 183, 184
107 MPI_Unpack_c, 180, 1050
108 MPI_UNPACK_EXTERNAL, 8, 187, 810

- MPI_Unpack_external_c, [187](#), [1050](#)
- MPI_UNPUBLISH_NAME, [463](#), [536](#), [537](#)
 - C example usage, [538](#)
- MPI_USER_FUNCTION
 - Fortran example usage, [240](#)
- MPI_WAIT, [39](#), [43](#), [78](#), [79](#), [80–84](#), [86](#), [102](#), [103](#), [110](#), [113](#), [114](#), [120](#), [121](#), [251](#), [275](#), [301](#), [489](#), [496](#), [517](#), [544](#), [627](#), [630](#), [631](#), [654](#), [655](#), [681](#), [701](#), [703](#), [816](#), [822](#), [823](#), [826](#), [827](#), [875–877](#), [879–882](#)
 - C example usage, [300–302](#)
 - Fortran example usage, [81–83](#), [89](#), [90](#), [668](#)
- MPI_WAITALL, [83](#), [85](#), [86](#), [87](#), [251](#), [302](#), [489](#), [496](#), [516](#), [544](#), [580](#), [629](#), [630](#), [633](#)
 - C example usage, [302](#), [303](#), [623](#)
- MPI_WAITANY, [55](#), [83](#), [84](#), [89](#), [489](#), [496](#), [516](#), [544](#), [629](#), [630](#), [633](#)
 - C example usage, [623](#)
 - Fortran example usage, [89](#)
- MPI_WAITSOME, [83](#), [87](#), [88–90](#), [93](#), [94](#), [489](#), [496](#), [516](#), [544](#), [629](#), [630](#), [633](#)
 - Fortran example usage, [90](#)
- MPI_WIN_ALLOCATE, [508](#), [510](#), [548](#), [551](#), [552](#), [553](#), [556–558](#), [562](#), [563](#), [569](#), [600](#), [801](#), [802](#), [882](#), [1051](#), [1054](#)
 - MPI_Win_allocate_c, [552](#), [1050](#)
- MPI_WIN_ALLOCATE_CPTR, [554](#), [1054](#)
- MPI_WIN_ALLOCATE_SHARED, [333](#), [510](#), [548](#), [551](#), [554](#), [555](#), [557](#), [558](#), [562](#), [564](#), [600](#), [802](#), [882](#), [1051](#), [1052](#), [1054](#)
 - MPI_Win_allocate_shared_c, [554](#), [1050](#)
- MPI_WIN_ALLOCATE_SHARED_CPTR, [556](#), [1054](#)
- MPI_WIN_ATTACH, [551](#), [559](#), [560](#), [561](#), [562](#), [600](#)
 - C example usage, [624](#)
- MPI_WIN_C2F, [836](#)
- MPI_WIN_CALL_ERRHANDLER, [468](#), [470](#)
- MPI_WIN_COMPLETE, [562](#), [588](#), [589](#), [593](#), [594–596](#), [606](#), [612](#), [882](#)
 - C example usage, [594](#), [619](#), [620](#)
 - Language-independent example usage, [612](#), [613](#)
- MPI_WIN_CREATE, [517](#), [548](#), [551–554](#), [556–559](#), [561–563](#), [605](#), [882](#)
 - Fortran example usage, [570](#), [572](#), [574](#)
 - MPI_Win_create_c, [549](#), [1050](#)
- MPI_WIN_CREATE_DYNAMIC, [463](#), [548](#), [558](#), [559](#), [560](#), [561](#), [563](#), [564](#), [605](#), [882](#)
 - C example usage, [624](#)
- MPI_WIN_CREATE_ERRHANDLER, [451](#), [452](#), [454](#), [455](#), [869](#), [871](#), [1060](#)
- MPI_WIN_CREATE_KEYVAL, [360](#), [367](#), [373](#), [843](#), [868](#), [871](#), [1065](#)
- MPI_WIN_DELETE_ATTR, [360](#), [370](#), [374](#)
- MPI_WIN_DETACH, [559](#), [561](#), [562](#), [563](#)
 - C example usage, [624](#)
- MPI_WIN_DUP_FN, [367](#), [859](#), [1055](#)
- MPI_WIN_F2C, [836](#)
- MPI_WIN_FENCE, [562](#), [570](#), [588](#), [590](#), [591](#), [592](#), [603](#), [604](#), [606](#), [607](#), [610](#), [615](#), [827](#), [882](#)
 - C example usage, [617–619](#)
 - Fortran example usage, [570](#), [572](#), [574](#), [827](#)
 - Language-independent example usage, [618](#)
- MPI_WIN_FLUSH, [558](#), [580](#), [581](#), [601](#), [606](#), [621](#), [622](#), [882](#)
 - C example usage, [624](#)
 - Language-independent example usage, [611](#), [621](#), [622](#)
- MPI_WIN_FLUSH_ALL, [580](#), [581](#), [601](#), [606](#), [882](#)
 - Language-independent example usage, [621](#)
- MPI_WIN_FLUSH_LOCAL, [580](#), [602](#), [606](#), [882](#)
 - Language-independent example usage, [610](#)
- MPI_WIN_FLUSH_LOCAL_ALL, [580](#), [602](#), [606](#), [882](#)
- MPI_WIN_FREE, [368](#), [489](#), [496](#), [544](#), [562](#), [563](#), [882](#), [1049](#)
 - Fortran example usage, [572](#), [574](#)
- MPI_WIN_FREE_KEYVAL, [360](#), [368](#), [374](#)
- MPI_WIN_GET_ATTR, [360](#), [369](#), [374](#), [563](#), [796](#), [843](#), [846](#)
- MPI_WIN_GET_ERRHANDLER, [451](#), [455](#), [1066](#)
- MPI_WIN_GET_GROUP, [564](#)
- MPI_WIN_GET_INFO, [509](#), [564](#), [565](#), [566](#), [1051](#), [1058](#)
- MPI_WIN_GET_NAME, [380](#), [1048](#)
- MPI_WIN_LOCK, [550](#), [562](#), [589](#), [597](#), [598–600](#), [603](#), [604](#), [606](#), [610–612](#), [877](#), [882](#)
 - C example usage, [600](#)
 - Language-independent example usage, [610–613](#)
- MPI_WIN_LOCK_ALL, [550](#), [589](#), [598](#), [599](#), [603](#), [604](#), [606](#), [612](#), [621](#), [882](#)
 - C example usage, [623](#), [624](#)
 - Language-independent example usage, [611](#), [623](#)
- MPI_WIN_NULL_COPY_FN, [367](#), [859](#), [1055](#)
- MPI_WIN_NULL_DELETE_FN, [367](#), [859](#), [1055](#)
- MPI_WIN_POST, [562](#), [588](#), [589](#), [593](#), [594](#), [595–597](#), [599](#), [603](#), [604](#), [606](#), [613](#), [615](#), [877](#), [882](#), [1049](#)
 - C example usage, [619](#), [620](#)
 - Language-independent example usage, [612](#), [613](#)
- MPI_WIN_SET_ATTR, [360](#), [369](#), [374](#), [563](#), [796](#), [843](#), [846](#)
- MPI_WIN_SET_ERRHANDLER, [451](#), [455](#)

1 MPI_WIN_SET_INFO, [564](#), [565](#), [882](#), [1051](#),
2 [1058](#)
3 MPI_WIN_SET_NAME, [379](#)
4 MPI_WIN_SHARED_QUERY, [551](#), [553](#), [555](#),
5 [556](#), [557](#), [558](#), [600](#), [802](#), [1049](#), [1054](#)
6 C example usage, [617](#)
7 MPI_Win_shared_query_c, [557](#), [1050](#)
8 MPI_WIN_SHARED_QUERY_CPTR, [559](#), [1054](#)
9 MPI_WIN_START, [562](#), [588](#), [589](#), [593](#), [594](#),
10 [595](#), [597](#), [603](#), [604](#), [613](#), [621](#), [877](#), [882](#)
11 C example usage, [594](#), [619](#), [620](#)
12 Language-independent example usage,
13 [612](#), [613](#)
14 MPI_WIN_SYNC, [602](#), [603](#), [606](#), [608](#), [609](#), [612](#),
15 [621–623](#), [1046](#), [1049](#)
16 Language-independent example usage,
17 [610](#), [621](#), [622](#)
18 MPI_WIN_TEST, [28](#), [596](#), [882](#), [1046](#), [1047](#)
19 MPI_WIN_UNLOCK, [562](#), [581](#), [589](#), [599](#), [600](#),
20 [606](#), [610](#), [611](#), [882](#)
21 C example usage, [600](#)
22 Language-independent example usage,
23 [610–613](#)
24 MPI_WIN_UNLOCK_ALL, [581](#), [589](#), [598](#), [599](#),
25 [606](#), [610](#), [621](#), [882](#)
26 C example usage, [623](#), [624](#)
27 Language-independent example usage, [611](#)
28 MPI_WIN_WAIT, [562](#), [588](#), [589](#), [595](#), [596](#), [599](#),
29 [606](#), [610](#), [612](#), [613](#), [620](#), [882](#)
30 C example usage, [619](#), [620](#)
31 Language-independent example usage,
32 [612](#), [613](#)
33 MPI_WTICK, [471](#), [751](#), [785](#), [1047](#)
34 MPI_WTIME, [15](#), [444](#), [470](#), [471](#), [738](#), [751](#), [785](#),
35 [1047](#)
36 C example usage, [470](#), [720](#)
37 MPIEXEC
38 Language-independent example usage,
39 [484](#), [513](#), [514](#)
40 mpiexec, [483](#), [484](#), [486](#), [508](#), [512](#), [1052](#)
41 mpirun, [512](#)
42
43 PMPI_, [717](#), [796](#)
44 PMPI_AINT_ADD, [25](#)
45 PMPI_AINT_DIFF, [25](#)
46 PMPI_ISEND, [796](#), [799](#)
47 PMPI_WTICK, [785](#), [1047](#)
48 PMPI_WTIME, [785](#), [1047](#)