

MPI: A Message-Passing Interface Standard

Message Passing Interface Forum

Version 1.3

March 14, 2008

This work 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.

Version 1.3: XXX, 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 1.2: July, 18 1997. The MPI-2 Forum introduced MPI 1.2 as Chap.3 in the standard "MPI-2: Extensions to the Message-Passing Interface", July 18, 1997. This section contains clarifications and minor corrections to Version 1.1 of the MPI Standard. The only new function in MPI-1.2 is one for identifying to which version of the MPI Standard the implementation conforms. There are small differences between MPI-1 and MPI-1.1. There are very few differences between MPI-1.1 and MPI-1.2, but large differences between MPI-1.2 and MPI-2.

Version 1.1: June, 1995. Beginning in March, 1995, the Message Passing Interface Forum reconvened to correct errors and make clarifications in the MPI document of May 5, 1994, referred to below as Version 1.0. These discussions resulted in Version 1.1, which is this document. The changes from Version 1.0 are minor. A version of this document with all changes marked is available. This paragraph is an example of a change.

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

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.

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

Please send comments on MPI to mpi-comments@cs.utk.edu. Your comment will be forwarded to MPIF committee members who will attempt to respond.

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

Contents

Acknowledgments	vii
1 Introduction to MPI	1
1.1 Overview and Goals	1
1.2 Who Should Use This Standard?	2
1.3 What Platforms Are Targets For Implementation?	3
1.4 What Is Included In The Standard?	3
1.5 What Is Not Included In The Standard?	3
1.6 Organization of this Document	4
2 MPI Terms and Conventions	6
2.1 Document Notation	6
2.2 Procedure Specification	6
2.3 Semantic Terms	7
2.4 Data Types	8
2.4.1 Opaque objects	8
2.4.2 Array arguments	9
2.4.3 State	9
2.4.4 Named constants	9
2.4.5 Choice	10
2.4.6 Addresses	10
2.5 Language Binding	10
2.5.1 Fortran 77 Binding Issues	11
2.5.2 C Binding Issues	12
2.6 Processes	12
2.7 Error Handling	13
2.8 Implementation issues	14
2.8.1 Independence of Basic Runtime Routines	14
2.8.2 Interaction with signals in POSIX	15
2.9 Examples	15
3 Point-to-Point Communication	16
3.1 Introduction	16
3.2 Blocking Send and Receive Operations	17
3.2.1 Blocking send	17
3.2.2 Message data	17
3.2.3 Message envelope	19

3.2.4	Blocking receive	20	1
3.2.5	Return status	21	2
3.3	Data type matching and data conversion	23	3
3.3.1	Type matching rules	23	4
3.3.2	Data conversion	26	5
3.4	Communication Modes	27	6
3.5	Semantics of point-to-point communication	31	7
3.6	Buffer allocation and usage	34	8
3.6.1	Model implementation of buffered mode	36	9
3.7	Nonblocking communication	37	10
3.7.1	Communication Objects	38	11
3.7.2	Communication initiation	38	12
3.7.3	Communication Completion	41	13
3.7.4	Semantics of Nonblocking Communications	45	14
3.7.5	Multiple Completions	46	15
3.8	Probe and Cancel	52	16
3.9	Persistent communication requests	56	17
3.10	Send-receive	60	18
3.11	Null processes	62	19
3.12	Derived datatypes	62	20
3.12.1	Datatype constructors	64	21
3.12.2	Address and extent functions	70	22
3.12.3	Lower-bound and upper-bound markers	72	23
3.12.4	Commit and free	73	24
3.12.5	Use of general datatypes in communication	75	25
3.12.6	Correct use of addresses	77	26
3.12.7	Examples	78	27
3.13	Pack and unpack	86	28
			29
4	Collective Communication	93	30
4.1	Introduction and Overview	93	31
4.2	Communicator argument	96	32
4.3	Barrier synchronization	96	33
4.4	Broadcast	96	34
4.4.1	Example using MPI_BCAST	97	35
4.5	Gather	97	36
4.5.1	Examples using MPI_GATHER, MPI_GATHERV	100	37
4.6	Scatter	106	38
4.6.1	Examples using MPI_SCATTER, MPI_SCATTERV	108	39
4.7	Gather-to-all	111	40
4.7.1	Examples using MPI_ALLGATHER, MPI_ALLGATHERV	112	41
4.8	All-to-All Scatter/Gather	113	42
4.9	Global Reduction Operations	115	43
4.9.1	Reduce	115	44
4.9.2	Predefined reduce operations	116	45
4.9.3	MINLOC and MAXLOC	118	46
4.9.4	User-Defined Operations	122	47
4.9.5	All-Reduce	126	48

1	4.10 Reduce-Scatter	127
2	4.11 Scan	128
3	4.11.1 Example using MPI_SCAN	128
4	4.12 Correctness	130
5		
6	5 Groups, Contexts, and Communicators	134
7	5.1 Introduction	134
8	5.1.1 Features Needed to Support Libraries	134
9	5.1.2 MPI's Support for Libraries	135
10	5.2 Basic Concepts	137
11	5.2.1 Groups	137
12	5.2.2 Contexts	137
13	5.2.3 Intra-Communicators	138
14	5.2.4 Predefined Intra-Communicators	138
15	5.3 Group Management	139
16	5.3.1 Group Accessors	139
17	5.3.2 Group Constructors	140
18	5.3.3 Group Destructors	144
19	5.4 Communicator Management	145
20	5.4.1 Communicator Accessors	145
21	5.4.2 Communicator Constructors	147
22	5.4.3 Communicator Destructors	150
23	5.5 Motivating Examples	150
24	5.5.1 Current Practice #1	150
25	5.5.2 Current Practice #2	151
26	5.5.3 (Approximate) Current Practice #3	152
27	5.5.4 Example #4	152
28	5.5.5 Library Example #1	153
29	5.5.6 Library Example #2	155
30	5.6 Inter-Communication	157
31	5.6.1 Inter-communicator Accessors	159
32	5.6.2 Inter-communicator Operations	160
33	5.6.3 Inter-Communication Examples	162
34	5.7 Caching	169
35	5.7.1 Functionality	170
36	5.7.2 Attributes Example	174
37	5.8 Formalizing the Loosely Synchronous Model	176
38	5.8.1 Basic Statements	176
39	5.8.2 Models of Execution	176
40		
41	6 Process Topologies	178
42	6.1 Introduction	178
43	6.2 Virtual Topologies	179
44	6.3 Embedding in MPI	179
45	6.4 Overview of the Functions	180
46	6.5 Topology Constructors	181
47	6.5.1 Cartesian Constructor	181
48	6.5.2 Cartesian Convenience Function: MPI_DIMS_CREATE	181

6.5.3	General (Graph) Constructor	182	1
6.5.4	Topology inquiry functions	184	2
6.5.5	Cartesian Shift Coordinates	188	3
6.5.6	Partitioning of Cartesian structures	189	4
6.5.7	Low-level topology functions	190	5
6.6	An Application Example	191	6
			7
7	MPI Environmental Management	193	8
7.1	Implementation information	193	9
7.1.1	Version Inquiries	193	10
7.1.2	Environmental Inquiries	194	11
7.2	Error handling	196	12
7.3	Error codes and classes	199	13
7.4	Timers and synchronization	201	14
7.5	Startup	201	15
			16
8	Profiling Interface	207	17
8.1	Requirements	207	18
8.2	Discussion	207	19
8.3	Logic of the design	208	20
8.3.1	Miscellaneous control of profiling	208	21
8.4	Examples	209	22
8.4.1	Profiler implementation	209	23
8.4.2	MPI library implementation	210	24
8.4.3	Complications	211	25
8.5	Multiple levels of interception	212	26
			27
	Bibliography	213	28
			29
	Language Binding	216	30
A.1	Introduction	216	31
A.2	Defined Constants for C and Fortran	216	32
A.3	C bindings for Point-to-Point Communication	220	33
A.4	C Bindings for Collective Communication	223	34
A.5	C Bindings for Groups, Contexts, and Communicators	224	35
A.6	C Bindings for Process Topologies	225	36
A.7	C bindings for Environmental Inquiry	226	37
A.8	C Bindings for Profiling	226	38
A.9	Fortran Bindings for Point-to-Point Communication	226	39
A.10	Fortran Bindings for Collective Communication	230	40
A.11	Fortran Bindings for Groups, Contexts, etc.	232	41
A.12	Fortran Bindings for Process Topologies	234	42
A.13	Fortran Bindings for Environmental Inquiry	235	43
A.14	Fortran Bindings for Profiling	235	44
			45
	MPI Function Index	236	46
			47
			48

Acknowledgments

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 served in positions of responsibility and are listed below.

- Jack Dongarra, David Walker, Conveners and Meeting Chairs
- Ewing Lusk, Bob Knighten, Minutes
- Marc Snir, William Gropp, Ewing Lusk, Point-to-Point Communications
- Al Geist, Marc Snir, Steve Otto, Collective Communications
- 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 process not mentioned above.

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

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.

MPI operated on a very tight budget (in reality, it had no budget when the first meeting was announced). ARPA and NSF have supported research at various institutions that have

made a contribution towards travel for the U.S. academics. Support for several European participants was provided by ESPRIT.

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

Chapter 1

Introduction to MPI

1.1 Overview and Goals

Message passing is a paradigm used widely on certain classes of parallel machines, especially those with distributed memory. Although there are many variations, the basic concept of processes communicating through messages is well understood. Over the last ten years, substantial progress has been made in casting significant applications in this paradigm. Each vendor has implemented its own variant. More recently, several systems have demonstrated that a message passing system can be efficiently and portably implemented. It is thus an appropriate time to try to define both the syntax and semantics of a core of library routines that will be useful to a wide range of users and efficiently implementable on a wide range of computers.

In designing MPI we have sought to make use of the most attractive features of a number of existing message passing systems, rather than selecting one of them and adopting it as the standard. Thus, MPI has been strongly influenced by work at the IBM T. J. Watson Research Center [1, 2], Intel's NX/2 [23], Express [22], nCUBE's Vertex [21], p4 [7, 6], and PARMACS [5, 8]. Other important contributions have come from Zipcode [24, 25], Chimp [14, 15], PVM [4, 11], Chameleon [19], and PICL [18].

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

A preliminary draft proposal, known as MPI1, was put forward by Dongarra, Hempel, Hey, and Walker in November 1992, and a revised version was completed in February 1993 [12]. MPI1 embodied the main features that were identified at the Williamsburg workshop as being necessary in a message passing standard. Since MPI1 was primarily intended to promote discussion and "get the ball rolling," it focused mainly on point-to-point communications. MPI1 brought to the forefront a number of important standardization issues, but did not include any collective communication routines and was not thread-safe.

In November 1992, a meeting of the MPI working group was held in Minneapolis, at which it was decided to place the standardization process on a more formal footing, and to

generally adopt the procedures and organization of the High Performance Fortran Forum. Subcommittees were formed for the major component areas of the standard, and an email discussion service established for each. In addition, the goal of producing a draft MPI standard by the Fall of 1993 was set. To achieve this goal the MPI working group met every 6 weeks for two days throughout the first 9 months of 1993, and presented the draft MPI standard at the Supercomputing 93 conference in November 1993. These meetings and the email discussion together constituted the MPI Forum, membership of which has been open to all members of the high performance computing community.

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 build 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 provide hardware support for, 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 and allow overlap of computation and communication and offload to communication co-processor, where available.
- Allow for implementations that can be used in a heterogeneous environment.
- Allow convenient C and Fortran 77 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.
- Define an interface that is not too different from current practice, such as PVM, NX, Express, p4, etc., and provides extensions that allow greater flexibility.
- Define an interface that can be implemented on many vendor's platforms, with no significant changes in the underlying communication and system software.
- Semantics of the interface should be language independent.
- The interface should be designed to allow for thread-safety.

1.2 Who Should Use This Standard?

This standard is intended for use by all those who want to write portable message-passing programs in Fortran 77 and 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.

1.3 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 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 programs, as well as those written in the more restricted style of SPMD. Although no explicit support for threads is provided, the interface has been designed so as not to prejudice their use. With this version of MPI no support is provided for dynamic spawning of tasks.

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. Several proprietary, native implementations of MPI, and a public domain, portable implementation of MPI are in progress at the time of this writing [17, 13].

1.4 What Is Included In The Standard?

The standard includes:

- Point-to-point communication
- Collective operations
- Process groups
- Communication contexts
- Process topologies
- Bindings for Fortran 77 and C
- Environmental Management and inquiry
- Profiling interface

1.5 What Is Not Included In The Standard?

The standard does not specify:

- Explicit shared-memory operations
- Operations that require more operating system support than is currently standard; for example, interrupt-driven receives, remote execution, or active messages

- Program construction tools
- Debugging facilities
- Explicit support for threads
- Support for task management
- I/O functions

There are many features that have been considered and not included in this standard. This happened for a number of reasons, one of which is the time constraint that was self-imposed in finishing the standard. Features that are not included can always be offered as extensions by specific implementations. Perhaps future versions of MPI will address some of these issues.

1.6 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, Collective Communications, 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 5, Groups, Contexts, and Communicators, 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 6, Process Topologies, explains a set of utility functions meant to assist in the mapping of process groups (a linearly ordered set) to richer topological structures such as multi-dimensional grids.
- Chapter 7, 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 8, Profiling Interface, explains a simple name-shifting convention that any MPI implementation must support. One motivation for this is the ability to put performance profiling calls into MPI without the need for access to the MPI source code. The name shift is merely an interface, it says nothing about how the actual profiling should be done and in fact, the name shift can be useful for other purposes.

- 1 • Annex [8.5](#), Language Bindings, gives specific syntax in Fortran 77 and C, for all MPI
2 functions, constants, and types.
- 3
- 4 • The MPI Function Index is a simple index showing the location of the precise definition
5 of each MPI function, together with both C and Fortran bindings.
- 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

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 that speaks to users and 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 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:

- the call uses but does not update an argument marked IN,
- the call may update an argument marked OUT,
- the call both uses and updates an argument marked INOUT.

There is one special case — if an argument is a handle to an opaque object (these terms are defined in Section 2.4.1), and the object is updated by the procedure call, then the argument is marked OUT. It is marked this way even though the handle itself is not modified — we use the OUT attribute to denote that what the handle *references* is updated.

1 The definition of MPI tries to avoid, to the largest possible extent, the use of INOUT
2 arguments, because such use is error-prone, especially for scalar arguments.

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

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

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

```
13 void copyIntBuffer( int *pin, int *pout, int len )
14 { int i;
15   for (i=0; i<len; ++i) *pout++ = *pin++;
16 }
17
```

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

```
19
20 int a[10];
21 copyIntBuffer( a, a+3, 7);
22
```

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

25 All MPI functions are first specified in the language-independent notation. Immediately
26 below this, the ANSI C version of the function is shown, and below this, a version of the
27 same function in Fortran 77.

29 2.3 Semantic Terms

30
31 When discussing MPI procedures the following semantic terms are used. The first two are
32 usually applied to communication operations.

33
34 **nonblocking** If the procedure may return before the operation completes, and before the
35 user is allowed to re-use resources (such as buffers) specified in the call.

36
37 **blocking** If return from the procedure indicates the user is allowed to re-use resources
38 specified in the call.

39
40 **local** If completion of the procedure depends only on the local executing process. Such an
41 operation does not require communication with another user process.

42
43 **non-local** If completion of the operation may require the execution of some MPI procedure
44 on another process. Such an operation may require communication occurring with
45 another user process.

46
47 **collective** If all processes in a process group need to invoke the procedure.
48

2.4 Data Types

2.4.1 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 assignment and comparisons.

In Fortran, all handles have type `INTEGER`. In C, a different handle type is defined for each category of objects. These should be types that support assignment and equality operators.

In Fortran, the handle can be an index to a table of opaque objects in system table; in C it can be such index or a pointer to the object. More bizarre possibilities exist.

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

A call to deallocate invalidates the handle and marks the object for deallocation. The object is not accessible to the user after the call. However, MPI need not deallocate the object immediately. Any operation pending (at the time of the deallocate) that involves this object will complete normally; the object will be deallocated afterwards.

MPI calls do not change the value of handles, with the exception of calls that allocate and deallocate objects, and of the call `MPI_TYPE_COMMIT`, in Section 3.12.4.

A null handle argument is an erroneous IN argument in MPI calls, unless an exception is explicitly stated in the text that defines the function. Such exception is allowed for handles to request objects in Wait and Test calls (sections 3.7.3 and 3.7.5). Otherwise, a null handle can only be passed to a function that allocates a new object and returns a reference to it in the handle.

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

MPI provides certain predefined opaque objects and predefined, static handles to these objects. Such objects may not be destroyed.

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

The explicit separating of handles in user space, objects in system space, allows space-reclaiming, deallocation calls to be made at appropriate points in the user program. If the opaque objects were in user space, one would have to be very careful not to go out of scope before any pending operation requiring that object completed. The specified design allows an object to be marked for deallocation, the user program can

1 then go out of scope, and the object itself still persists until any pending operations
2 are complete.

3 The requirement that handles support assignment/comparison is made since such
4 operations are common. This restricts the domain of possible implementations. The
5 alternative would have been to allow handles to have been an arbitrary, opaque type.
6 This would force the introduction of routines to do assignment and comparison, adding
7 complexity, and was therefore ruled out. (*End of rationale.*)
8

9 *Advice to users.* A user may accidentally create a dangling reference by assigning to a
10 handle the value of another handle, and then deallocating the object associated with
11 these handles. Conversely, if a handle variable is deallocated before the associated
12 object is freed, then the object becomes inaccessible (this may occur, for example, if
13 the handle is a local variable within a subroutine, and the subroutine is exited before
14 the associated object is deallocated). It is the user's responsibility to avoid adding
15 or deleting references to opaque objects, except as a result of calls that allocate or
16 deallocate such objects. (*End of advice to users.*)
17

18 *Advice to implementors.* The intended semantics of opaque objects is that each
19 opaque object is separate from each other; each call to allocate such an object copies
20 all the information required for the object. Implementations may avoid excessive
21 copying by substituting referencing for copying. For example, a derived datatype
22 may contain references to its components, rather than copies of its components; a
23 call to MPI_COMM_GROUP may return a reference to the group associated with the
24 communicator, rather than a copy of this group. In such cases, the implementation
25 must maintain reference counts, and allocate and deallocate objects such that the
26 visible effect is as if the objects were copied. (*End of advice to implementors.*)
27

28 2.4.2 Array arguments

29 An MPI call may need an argument that is an array of opaque objects, or an array of
30 handles. The array-of-handles is a regular array with entries that are handles to objects
31 of the same type in consecutive locations in the array. Whenever such an array is used,
32 an additional len argument is required to indicate the number of valid entries (unless this
33 number can be derived otherwise). The valid entries are at the beginning of the array; len
34 indicates how many of them there are, and need not be the entire size of the array. The
35 same approach is followed for other array arguments.
36

37 2.4.3 State

38 MPI procedures use at various places arguments with *state* types. The values of such data
39 type are all identified by names, and no operation is defined on them. For example, the
40 MPI_ERRHANDLER_SET routine has a state type argument with values MPI_ERRORS_ARE_FATAL,
41 MPI_ERRORS_RETURN, etc.
42
43

44 2.4.4 Named constants

45 MPI procedures sometimes assign a special meaning to a special value of a basic type argu-
46 ment; e.g. tag is an integer-valued argument of point-to-point communication operations,
47 with a special wild-card value, MPI_ANY_TAG. Such arguments will have a range of regular
48

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 can be queried using environmental inquiry functions (Chapter 7).

MPI also provides predefined named constant handles, such as `MPI_COMM_WORLD` which is a handle to an object that represents all processes available at start-up time and allowed to communicate with any of them.

All named constants, with the exception of `MPI_BOTTOM` in Fortran, can be used in initialization expressions or assignments. These constants do not change values during execution. Opaque objects accessed by constant handles are defined and do not change value between MPI initialization (`MPI_INIT()` call) and MPI completion (`MPI_FINALIZE()` call).

2.4.5 Choice

MPI functions sometimes use arguments with a *choice* (or union) data type. Distinct calls to the same routine may pass by reference actual arguments of different types. The mechanism for providing such arguments will differ from language to language. For Fortran, the document uses `<type>` to represent a choice variable, for C, we use `(void *)`.

2.4.6 Addresses

Some MPI procedures use *address* arguments that represent an absolute address in the calling program. The datatype of such an argument is an integer of the size needed to hold any valid address in the execution environment.

2.5 Language Binding

This section defines the rules for MPI language binding in general and for Fortran 77 and ANSI C in particular. Defined here are various object representations, as well as the naming conventions used for expressing this standard. The actual calling sequences are defined elsewhere.

It is expected that any Fortran 90 and C++ implementations use the Fortran 77 and ANSI C bindings, respectively. Although we consider it premature to define other bindings to Fortran 90 and C++, the current bindings are designed to encourage, rather than discourage, experimentation with better bindings that might be adopted later.

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

There are several important language binding issues not addressed by this standard. This standard does not discuss the interoperability of message passing between languages. It is fully expected that many implementations will have such features, and that such features are a sign of the quality of the implementation.

```

1      double precision a
2      integer b
3      ...
4      call MPI_send(a,...)
5      call MPI_send(b,...)
6

```

Figure 2.1: An example of calling a routine with mismatched formal and actual arguments.

2.5.1 Fortran 77 Binding Issues

All MPI names have an `MPI_` prefix, and all characters are capitals. Programs must not declare variables or functions with names beginning with the prefix, `MPI_`. This is mandated to avoid possible name collisions.

All MPI Fortran subroutines have a return code in the last argument. A few MPI operations are functions, which do not have the return code argument. The return code value for successful completion is `MPI_SUCCESS`. Other error codes are implementation dependent; see Chapter 7.

Handles are represented in Fortran as `INTEGERS`. Binary-valued variables are of type `LOGICAL`.

Array arguments are indexed from one.

Unless explicitly stated, the MPI F77 binding is consistent with ANSI standard Fortran 77. There are several points where this standard diverges from the ANSI Fortran 77 standard. These exceptions are consistent with common practice in the Fortran community. In particular:

- MPI identifiers are limited to thirty, not six, significant characters.
- MPI identifiers may contain underscores after the first character.
- An MPI subroutine with a choice argument may be called with different argument types. An example is shown in Figure 2.1. This violates the letter of the Fortran standard, but such a violation is common practice. An alternative would be to have a separate version of `MPI_SEND` for each data type.
- Although not required, it is strongly suggested that named MPI constants (`PARAMETERS`) be provided in an include file, called `mpif.h`. On systems that do not support include files, the implementation should specify the values of named constants.
- Vendors are encouraged to provide type declarations in the `mpif.h` file on Fortran systems that support user-defined types. One should define, if possible, the type `MPI_ADDRESS_TYPE`, which is an `INTEGER` of the size needed to hold an address in the execution environment. On systems where type definition is not supported, it is up to the user to use an `INTEGER` of the right kind to represent addresses (i.e., `INTEGER*4` on a 32 bit machine, `INTEGER*8` on a 64 bit machine, etc.).

All MPI named constants can be used wherever an entity declared with the `PARAMETER` attribute can be used in Fortran. There is one exception to this rule: the MPI constant `MPI_BOTTOM` (section 3.12.2) can only be used as a buffer argument.

2.5.2 C Binding Issues

We use the ANSI 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 variables or functions with names beginning with the prefix, `MPI_`. This is mandated to avoid possible name collisions.

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

Almost all C functions return an error code. The successful return code will be `MPI_SUCCESS`, but failure return codes are implementation dependent. A few C functions do not return values, so that they can be implemented as macros.

Type declarations are provided for handles to each category of opaque objects. Either a pointer or an integer type is used.

Array arguments are indexed from zero.

Logical flags are integers with value 0 meaning “false” and a non-zero value meaning “true.”

Choice arguments are pointers of type `void*`.

Address arguments are of MPI defined type `MPI_Aint`. This is defined to be an int of the size needed to hold any valid address on the target architecture.

All named MPI constants can be used in initialization expressions or assignments like C constants.

2.6 Processes

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. This document specifies the behavior of a parallel program assuming that only MPI calls are used for communication. The interaction of an MPI program with other possible means of communication (e.g., shared memory) is not specified.

MPI does not specify the execution model for each process. A process can be sequential, or can be multi-threaded, with threads possibly executing concurrently. Care has been taken to make MPI “thread-safe,” by avoiding the use of implicit state. The desired interaction of MPI with threads is that concurrent threads be all allowed to execute MPI calls, and calls be reentrant; a blocking MPI call blocks only the invoking thread, allowing the scheduling of another thread.

MPI does not provide mechanisms to specify the initial allocation of processes to an MPI computation and their binding to physical processors. It is expected that vendors will provide mechanisms to do so either at load time or at run time. Such mechanisms will allow the specification of the initial number of required processes, the code to be executed by each initial process, and the allocation of processes to processors. Also, the current proposal does not provide for dynamic creation or deletion of processes during program execution (the total number of processes is fixed), although it is intended to be consistent with such extensions. Finally, we always identify processes according to their relative rank in a group, that is, consecutive integers in the range `0..groupsize-1`.

2.7 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 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, or to reflect unrecoverable errors as failures. Whenever possible, such failures will be reflected as errors in the relevant communication call. Similarly, MPI itself provides no mechanisms for handling processor failures. The error handling facilities described in section 7.2 can be used to restrict the scope of an unrecoverable error, or design error recovery at the application level.

Of course, MPI programs may still be erroneous. A **program error** can occur when an MPI call is called with an incorrect argument (non-existing 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.

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. In certain circumstances, when the MPI function may complete several distinct operations, and therefore may generate several independent errors, the MPI function may return multiple error codes. By default, an error detected during the execution of the MPI library causes the parallel computation to abort. 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 himself or herself. Also, the user may provide his or her own error-handling routines, which will be invoked whenever an MPI call returns abnormally. The MPI error handling facilities are described in section 7.2.

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; finally some errors may be “catastrophic” and may prevent MPI from returning control to the caller in a consistent state.

Another subtle issue 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 exception 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 completed, so that no error value can be used to indicate the nature of the error (e.g., an error in a send with the ready mode). Such an error must be treated as fatal, since information cannot be returned for the user to recover from it.

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.

2.8 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 I/O or 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.

2.8.1 Independence of Basic Runtime Routines

MPI programs require that library routines that are part of the basic language environment (such as `date` and `write` in Fortran and `printf` and `malloc` in ANSI 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 ANSI C environment regardless of the size of `MPI_COMM_WORLD` (assuming that I/O is available at the executing nodes).

```
int rank;
MPI_Init( argc, argv );
MPI_Comm_rank( MPI_COMM_WORLD, &rank );
if (rank == 0) printf( "Starting program\n" );
MPI_Finalize();
```

The corresponding Fortran 77 program is 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 tasks. 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 nodes).

```
MPI_Comm_rank( MPI_COMM_WORLD, &rank );
printf( "Output from task rank %d\n", 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).

2.8.2 Interaction with signals in POSIX

MPI does not specify either the interaction of processes with signals, in a UNIX environment, or with other events that do not relate to MPI communication. That is, signals are not significant from the view point of MPI, and implementors should attempt to implement MPI so that signals are transparent: an MPI call suspended by a signal should resume and complete after the signal is handled. Generally, the state of a computation that is visible or significant from the view-point of MPI should only be affected by MPI calls.

The intent of MPI to be thread and signal safe has a number of subtle effects. For example, on Unix systems, a catchable signal such as SIGALRM (an alarm signal) must not cause an MPI routine to behave differently than it would have in the absence of the signal. Of course, if the signal handler issues MPI calls or changes the environment in which the MPI routine is operating (for example, consuming all available memory space), the MPI routine should behave as appropriate for that situation (in particular, in this case, the behavior should be the same as for a multithreaded MPI implementation).

A second effect is that a signal handler that performs MPI calls must not interfere with the operation of MPI. For example, an MPI receive of any type that occurs within a signal handler must not cause erroneous behavior by the MPI implementation. Note that an implementation is permitted to prohibit the use of MPI calls from within a signal handler, and is not required to detect such use.

It is highly desirable that MPI not use SIGALRM, SIGFPE, or SIGIO. An implementation is *required* to clearly document all of the signals that the MPI implementation uses; a good place for this information is a Unix ‘man’ page on MPI.

2.9 Examples

The examples in this document are for illustration purposes only. They are not intended to specify the standard. Furthermore, the examples have not been carefully checked or verified.

Chapter 3

Point-to-Point Communication

3.1 Introduction

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

```
#include "mpi.h"
main( argc, argv )
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 /* code for process one */
    {
        MPI_Recv(message, 20, MPI_CHAR, 0, 99, MPI_COMM_WORLD, &status);
        printf("received :%s:\n", message);
    }
    MPI_Finalize();
}
```

In this example, process zero (`myrank = 0`) 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 specify the envelope for the message sent.

Process one (`myrank = 1`) 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.

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 channel-like constructs and send-receive operations. We then 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.

3.2 Blocking Send and Receive Operations

3.2.1 Blocking send

The syntax of the blocking send operation 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 (nonnegative 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)

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

```
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 Sec. 3.4.

3.2.2 Message data

The send buffer specified by the `MPI_SEND` operation 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 below.

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	

Possible values for this argument for C and the corresponding C types are listed below.

MPI datatype	C datatype
MPI_CHAR	signed char
MPI_SHORT	signed short int
MPI_INT	signed int
MPI_LONG	signed long int
MPI_UNSIGNED_CHAR	unsigned char
MPI_UNSIGNED_SHORT	unsigned short int
MPI_UNSIGNED	unsigned int
MPI_UNSIGNED_LONG	unsigned long int
MPI_FLOAT	float
MPI_DOUBLE	double
MPI_LONG_DOUBLE	long double
MPI_BYTE	
MPI_PACKED	

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 hand, a byte has the same binary value on all machines. The use of the type `MPI_PACKED` is explained in Section 3.13.

1 MPI requires support of the datatypes listed above, which match the basic datatypes
 2 of Fortran 77 and ANSI C. Additional MPI datatypes should be provided if the host lan-
 3 guage has additional data types: `MPI_LONG_LONG_INT`, for C integers declared to be of
 4 type `long long`; `MPI_DOUBLE_COMPLEX` for double precision complex in Fortran declared
 5 to be of type `DOUBLE COMPLEX`; `MPI_REAL2`, `MPI_REAL4` and `MPI_REAL8` for Fortran
 6 reals, declared to be of type `REAL*2`, `REAL*4` and `REAL*8`, respectively; `MPI_INTEGER1`
 7 `MPI_INTEGER2` and `MPI_INTEGER4` for Fortran integers, declared to be of type `INTEGER*1`,
 8 `INTEGER*2` and `INTEGER*4`, respectively; etc.

9
 10 *Rationale.* One goal of the design is to allow for MPI to be implemented as a
 11 library, with no need for additional preprocessing or compilation. Thus, one cannot
 12 assume that a communication call has information on the datatype of variables in the
 13 communication buffer; this information must be supplied by an explicit argument.
 14 The need for such datatype information will become clear in Section 3.3.2. (*End of*
 15 *rationale.*)

16 3.2.3 Message envelope

17
 18 In addition to the data part, messages carry information that can be used to distinguish
 19 messages and selectively receive them. This information consists of a fixed number of fields,
 20 which we collectively call the **message envelope**. These fields are

21
 22 source
 23 destination
 24 tag
 25 communicator
 26

27 The message source is implicitly determined by the identity of the message sender. The
 28 other fields are specified by arguments in the send operation.

29 The message destination is specified by the `dest` argument.

30 The integer-valued message tag is specified by the `tag` argument. This integer can be
 31 used by the program to distinguish different types of messages. The range of valid tag
 32 values is `0,...,UB`, where the value of `UB` is implementation dependent. It can be found by
 33 querying the value of the attribute `MPI_TAG_UB`, as described in Chapter 7. MPI requires
 34 that `UB` be no less than 32767.

35 The `comm` argument specifies the **communicator** that is used for the send operation.
 36 Communicators are explained in Chapter 5; below is a brief summary of their usage.

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

41 The communicator also specifies the set of processes that share this communication
 42 context. This **process group** is ordered and processes are identified by their rank within
 43 this group. Thus, the range of valid values for `dest` is `0, ... , n-1`, where `n` is the number of
 44 processes in the group. (If the communicator is an inter-communicator, then destinations
 45 are identified by their rank in the remote group. See Chapter 5.)

46 A predefined communicator `MPI_COMM_WORLD` is provided by MPI. It allows commu-
 47 nication with all processes that are accessible after MPI initialization and processes are
 48 identified by their rank in the group of `MPI_COMM_WORLD`.

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

Users may define new communicators, as explained in Chapter 5. Communicators provide an important encapsulation mechanism for libraries and modules. They allow modules to have their own disjoint communication universe and their own process numbering scheme. (*End of advice to users.*)

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

3.2.4 Blocking receive

The syntax of the blocking receive operation is given below.

`MPI_RECV (buf, count, datatype, source, tag, comm, status)`

OUT	<code>buf</code>	initial address of receive buffer (choice)
IN	<code>count</code>	number of elements in receive buffer (integer)
IN	<code>datatype</code>	datatype of each receive buffer element (handle)
IN	<code>source</code>	rank of source (integer)
IN	<code>tag</code>	message tag (integer)
IN	<code>comm</code>	communicator (handle)
OUT	<code>status</code>	status object (Status)

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

```
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 Sec. 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.

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

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

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

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

21 The message tag is specified by the **tag** argument of the receive operation. The
22 argument **source**, if different from `MPI_ANY_SOURCE`, is specified as a rank within the
23 process group associated with that same communicator (remote process group, for in-
24 tercommunicators). Thus, the range of valid values for the **source** argument is $\{0, \dots, n-1\} \cup \{\text{MPI_ANY_SOURCE}\}$, where n is the number of processes in this group.

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

31 Source = destination is allowed, that is, a process can send a message to itself. (How-
32 ever, it is unsafe to do so with the blocking send and receive operations described above,
33 since this may lead to deadlock. See Sec. 3.5.)
34

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

41 3.2.5 Return status

42 The source or tag of a received message may not be known if wildcard values were used
43 in the receive operation. Also, if multiple requests are completed by a single MPI function
44 (see Section 3.7.5), a distinct error code may need to be returned for each request. The
45 information is returned by the **status** argument of `MPI_RECV`. The type of **status** is MPI-
46 defined. Status variables need to be explicitly allocated by the user, that is, they are not
47 system objects.
48

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, `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.

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 which 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.

`MPI_GET_COUNT(status, datatype, count)`

IN	<code>status</code>	return status of receive operation (Status)
IN	<code>datatype</code>	datatype of each receive buffer entry (handle)
OUT	<code>count</code>	number of received entries (integer)

`int MPI_Get_count(MPI_Status *status, MPI_Datatype datatype, int *count)`

`MPI_GET_COUNT(STATUS, DATATYPE, COUNT, IERROR)`

`INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR`

Returns the number of entries received. (Again, we count *entries*, each of type *datatype*, not *bytes*.) The `datatype` argument should match the argument provided by the receive call that set the `status` variable. (We shall later see, in Section 3.12.5, that `MPI_GET_COUNT` may return, in certain situations, the value `MPI_UNDEFINED`.)

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

The `datatype` argument is passed to `MPI_GET_COUNT` so as to improve performance. A message might be received without counting the number of elements it contains, and the `count` value is often not needed. Also, this allows the same function to be

1 used after a call to `MPI_PROBE` or `MPI_Iprobe`. With a status from `MPI_PROBE`
2 or `MPI_Iprobe`, the same datatypes are allowed as in a call to `MPI_RECV` to receive
3 this message. (*End of rationale.*)
4

5 *Advice to users.* The buffer size required for the receive can be affected by data con-
6 versions and by the stride of the receive datatype. In most cases, the safest approach
7 is to use the same datatype with `MPI_GET_COUNT` and the receive. (*End of advice*
8 *to users.*)
9

10 All send and receive operations use the `buf`, `count`, `datatype`, `source`, `dest`, `tag`, `comm`
11 and `status` arguments in the same way as the blocking `MPI_SEND` and `MPI_RECV` operations
12 described in this section.
13

14 3.3 Data type matching and data conversion

15 3.3.1 Type matching rules

16 One can think of message transfer as consisting of the following three phases.
17

- 18 1. Data is pulled out of the send buffer and a message is assembled.
- 19 2. A message is transferred from sender to receiver.
- 20 3. Data is pulled from the incoming message and disassembled into the receive buffer.

21 Type matching has to be observed at each of these three phases: The type of each
22 variable in the sender buffer has to match the type specified for that entry by the send
23 operation; the type specified by the send operation has to match the type specified by the
24 receive operation; and the type of each variable in the receive buffer has to match the type
25 specified for that entry by the receive operation. A program that fails to observe these three
26 rules is erroneous.
27

28 To define type matching more precisely, we need to deal with two issues: matching of
29 types of the host language with types specified in communication operations; and matching
30 of types at sender and receiver.
31

32 The types of a send and receive match (phase two) if both operations use identical
33 names. That is, `MPI_INTEGER` matches `MPI_INTEGER`, `MPI_REAL` matches `MPI_REAL`,
34 and so on. There is one exception to this rule, discussed in Sec. 3.13, the type `MPI_PACKED`
35 can match any other type.
36

37 The type of a variable in a host program matches the type specified in the commu-
38 nication operation if the datatype name used by that operation corresponds to the basic
39 type of the host program variable. For example, an entry with type name `MPI_INTEGER`
40 matches a Fortran variable of type `INTEGER`. A table giving this correspondence for Fortran
41 and C appears in Sec. 3.2.2. There are two exceptions to this last rule: an entry with
42 type name `MPI_BYTE` or `MPI_PACKED` can be used to match any byte of storage (on a
43 byte-addressable machine), irrespective of the datatype of the variable that contains this
44 byte. The type `MPI_PACKED` is used to send data that has been explicitly packed, or receive
45 data that will be explicitly unpacked, see Section 3.13. The type `MPI_BYTE` allows one to
46 transfer the binary value of a byte in memory unchanged.
47

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

- Communication of typed values (e.g., with datatype different from `MPI_BYTE`), where the datatypes of the corresponding entries in the sender program, in the send call, in the receive call and in the receiver program must all match.
- 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 3.1 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
  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 3.2 Sender and receiver do not 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
  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 3.3 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
  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 bound 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 communications 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 3.4 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
    CALL MPI_RECV(b(6:10), 5, MPI_CHARACTER, 0, tag, comm, status, ierr)
END IF

```

The last five characters of string b at process 1 are replaced by the first five characters of string a at process 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 3.12). 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`.

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

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 or character values, and to convert a floating point value to the nearest value that can be represented on the target system.

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.

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.)

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

Consider the three examples, 3.1–3.3. 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.

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

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.

1 Data representation conversion also applies to the envelope of a message: source, des-
2 tination and tag are all integers that may need to be converted.

3
4 *Advice to implementors.* The current definition does not require messages to carry
5 data type information. Both sender and receiver provide complete data type infor-
6 mation. In a heterogeneous environment, one can either use a machine independent
7 encoding such as XDR, or have the receiver convert from the sender representation
8 to its own, or even have the sender do the conversion.

9 Additional type information might be added to messages in order to allow the sys-
10 tem to detect mismatches between datatype at sender and receiver. This might be
11 particularly useful in a slower but safer debug mode. (*End of advice to implementors.*)
12

13 MPI does not require support for inter-language communication. The behavior of a
14 program is undefined if messages are sent by a C process and received by a Fortran process,
15 or vice-versa.

16
17 *Rationale.* MPI does not handle inter-language communication because there are no
18 agreed standards for the correspondence between C types and Fortran types. There-
19 fore, MPI programs that mix languages would not port. (*End of rationale.*)

20
21 *Advice to implementors.* MPI implementors may want to support inter-language
22 communication by allowing Fortran programs to use “C MPI types,” such as MPI_INT,
23 MPI_CHAR, etc., and allowing C programs to use Fortran types. (*End of advice to*
24 *implementors.*)
25

26 3.4 Communication Modes

27
28 The send call described in Section 3.2.1 is **blocking**: it does not return until the message
29 data and envelope have been safely stored away so that the sender is free to access and
30 overwrite the send buffer. The message might be copied directly into the matching receive
31 buffer, or it might be copied into a temporary system buffer.

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

38 The send call described in Section 3.2.1 used the **standard** communication mode. In
39 this mode, it is up to MPI to decide whether outgoing messages will be buffered. MPI may
40 buffer outgoing messages. In such a case, the send call may complete before a matching
41 receive is invoked. On the other hand, buffer space may be unavailable, or MPI may choose
42 not to buffer outgoing messages, for performance reasons. In this case, the send call will
43 not complete until a matching receive has been posted, and the data has been moved to the
44 receiver.

45 Thus, a send in standard mode can be started whether or not a matching receive has
46 been posted. It may complete before a matching receive is posted. The standard mode send
47 is **non-local**: successful completion of the send operation may depend on the occurrence
48 of a matching receive.

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

There are three additional communication modes.

A **buffered** mode send operation can be started whether or not a matching receive has been posted. It may complete before a matching receive is posted. However, unlike the standard send, this operation is **local**, and its completion does not depend on the occurrence of a matching receive. Thus, if a send is executed and no matching receive is posted, then 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 posted. However, the send will complete successfully only if a matching receive is posted, 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 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 processes rendezvous at the communication. A send executed in this mode is **non-local**.

A send that uses the **ready** communication mode may be started *only* if the matching receive is already posted. Otherwise, the operation is erroneous and its outcome is undefined. On some systems, this allows the removal of a hand-shake operation 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 posted), 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.

```

1 MPI_BSEND (buf, count, datatype, dest, tag, comm)
2     IN      buf                initial address of send buffer (choice)
3
4     IN      count              number of elements in send buffer (integer)
5
6     IN      datatype           datatype of each send buffer element (handle)
7
8     IN      dest               rank of destination (integer)
9
10    IN      tag                 message tag (integer)
11
12    IN      comm                communicator (handle)
13
14 int MPI_Bsend(void* buf, int count, MPI_Datatype datatype, int dest,
15              int tag, MPI_Comm comm)
16
17 MPI_BSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
18     <type> BUF(*)
19     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR
20
21     Send in buffered mode.
22
23 MPI_SSEND (buf, count, datatype, dest, tag, comm)
24     IN      buf                initial address of send buffer (choice)
25
26     IN      count              number of elements in send buffer (integer)
27
28     IN      datatype           datatype of each send buffer element (handle)
29
30     IN      dest               rank of destination (integer)
31
32     IN      tag                 message tag (integer)
33
34     IN      comm                communicator (handle)
35
36 int MPI_Ssend(void* buf, int count, MPI_Datatype datatype, int dest,
37              int tag, MPI_Comm comm)
38
39 MPI_SSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
40     <type> BUF(*)
41     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR
42
43     Send in synchronous mode.
44
45
46
47
48

```

MPI_RSEND (buf, count, datatype, dest, tag, comm)			1
IN	buf	initial address of send buffer (choice)	2
IN	count	number of elements in send buffer (integer)	3
IN	datatype	datatype of each send buffer element (handle)	4
IN	dest	rank of destination (integer)	5
IN	tag	message tag (integer)	6
IN	comm	communicator (handle)	7

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

```

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

```

Send in ready mode.

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

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

Rationale. We prohibit read accesses to a send buffer while it is being used, even though the send operation is not supposed to alter the content of this buffer. This may seem more stringent than necessary, but the additional restriction causes little loss of functionality and allows better performance on some systems — consider the case where data transfer is done by a DMA engine that is not cache-coherent with the main processor. (*End of rationale.*)

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

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

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

ready send: The message is sent as soon as possible.

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

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

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

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

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

7 A standard send can be implemented as a synchronous send. In such a case, no data
8 buffering is needed. However, many (most?) users expect some buffering.

9 In a multi-threaded environment, the execution of a blocking communication should
10 block only the executing thread, allowing the thread scheduler to de-schedule this
11 thread and schedule another thread for execution. (*End of advice to implementors.*)
12
13

14 3.5 Semantics of point-to-point communication

15 A valid MPI implementation guarantees certain general properties of point-to-point com-
16 munication, which are described in this section.
17
18

19 **Order** Messages are *non-overtaking*: If a sender sends two messages in succession to the
20 same destination, and both match the same receive, then this operation cannot receive the
21 second message if the first one is still pending. If a receiver posts two receives in succession,
22 and both match the same message, then the second receive operation cannot be satisfied
23 by this message, if the first one is still pending. This requirement facilitates matching of
24 sends to receives. It guarantees that message-passing code is deterministic, if processes
25 are single-threaded and the wildcard MPI_ANY_SOURCE is not used in receives. (Some of
26 the calls described later, such as MPI_CANCEL or MPI_WAITANY, are additional sources of
27 nondeterminism.)
28

29 If a process has a single thread of execution, then any two communications executed
30 by this process are ordered. On the other hand, if the process is multi-threaded, then the
31 semantics of thread execution may not define a relative order between two send operations
32 executed by two distinct threads. The operations are logically concurrent, even if one
33 physically precedes the other. In such a case, the two messages sent can be received in
34 any order. Similarly, if two receive operations that are logically concurrent receive two
35 successively sent messages, then the two messages can match the two receives in either
36 order.

37 **Example 3.5** An example of non-overtaking messages.

```
38 CALL MPI_COMM_RANK(comm, rank, ierr)
39 IF (rank.EQ.0) THEN
40     CALL MPI_BSEND(buf1, count, MPI_REAL, 1, tag, comm, ierr)
41     CALL MPI_BSEND(buf2, count, MPI_REAL, 1, tag, comm, ierr)
42 ELSE    ! rank.EQ.1
43     CALL MPI_RECV(buf1, count, MPI_REAL, 0, MPI_ANY_TAG, comm, status, ierr)
44     CALL MPI_RECV(buf2, count, MPI_REAL, 0, tag, comm, status, ierr)
45 END IF
46
```

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

Progress If a pair of matching send and receives have been initiated on two processes, 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 posted at the same destination process.

Example 3.6 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 ! rank.EQ.1
  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 processes invoke their first communication call. Since the first send of process zero uses the buffered mode, it must complete, irrespective of the state of process one. Since no matching receive is posted, 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. Process one next invokes its second receive call, which will be satisfied by the buffered message. Note that process one 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 posted. Then it is possible that the destination 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 posted by a multi-threaded 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 posted at this node (by other executing threads). It is the programmer's responsibility to prevent starvation in such situations.

Resource limitations Any pending communication operation consumes system resources that are limited. Errors may occur when lack of resources prevent the execution of an MPI call. A quality implementation 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.

1 A buffered send operation that cannot complete because of a lack of buffer space is
 2 erroneous. When such a situation is detected, an error is signalled that may cause the
 3 program to terminate abnormally. On the other hand, a standard send operation that
 4 cannot complete because of lack of buffer space will merely block, waiting for buffer space
 5 to become available or for a matching receive to be posted. This behavior is preferable in
 6 many situations. Consider a situation where a producer repeatedly produces new values
 7 and sends them to a consumer. Assume that the producer produces new values faster
 8 than the consumer can consume them. If buffered sends are used, then a buffer overflow
 9 will result. Additional synchronization has to be added to the program so as to prevent
 10 this from occurring. If standard sends are used, then the producer will be automatically
 11 throttled, as its send operations will block when buffer space is unavailable.

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

14 **Example 3.7** An exchange of messages.

```
15 CALL MPI_COMM_RANK(comm, rank, ierr)
16
17 IF (rank.EQ.0) THEN
18     CALL MPI_SEND(sendbuf, count, MPI_REAL, 1, tag, comm, ierr)
19     CALL MPI_RECV(recvbuf, count, MPI_REAL, 1, tag, comm, status, ierr)
20 ELSE ! rank.EQ.1
21     CALL MPI_RECV(recvbuf, count, MPI_REAL, 0, tag, comm, status, ierr)
22     CALL MPI_SEND(sendbuf, count, MPI_REAL, 0, tag, comm, ierr)
23 END IF
24
```

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

28 **Example 3.8** An attempt to exchange messages.

```
29
30 CALL MPI_COMM_RANK(comm, rank, ierr)
31 IF (rank.EQ.0) THEN
32     CALL MPI_RECV(recvbuf, count, MPI_REAL, 1, tag, comm, status, ierr)
33     CALL MPI_SEND(sendbuf, count, MPI_REAL, 1, tag, comm, ierr)
34 ELSE ! rank.EQ.1
35     CALL MPI_RECV(recvbuf, count, MPI_REAL, 0, tag, comm, status, ierr)
36     CALL MPI_SEND(sendbuf, count, MPI_REAL, 0, tag, comm, ierr)
37 END IF
38
```

39 The receive operation of the first process must complete before its send, and can complete
 40 only if the matching send of the second processor is executed. The receive operation of the
 41 second process must complete before its send and can complete only if the matching send
 42 of the first process is executed. This program will always deadlock. The same holds for any
 43 other send mode.

44 **Example 3.9** An exchange that relies on buffering.

```
45
46 CALL MPI_COMM_RANK(comm, rank, ierr)
47 IF (rank.EQ.0) THEN
48     CALL MPI_SEND(sendbuf, count, MPI_REAL, 1, tag, comm, ierr)

```

```

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

```

The message sent by each process has to be copied out before the send operation returns 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. When standard send operations are used, then a deadlock situation may occur where both processes are blocked because 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 program is “safe” if no message buffering is required for the program to complete. One can replace all sends in such program with synchronous sends, and the program will still run correctly. This conservative programming style provides the best portability, since program completion does not depend on the amount of buffer space available or in the communication protocol used.

Many programmers prefer to have more leeway and be able to use the “unsafe” programming style shown in example 3.9. In such cases, the use of standard sends is likely to provide the best compromise between performance and robustness: quality implementations will provide sufficient buffering so that “common practice” programs will not deadlock. The buffered send mode can be used for programs that require more buffering, or in situations where the programmer wants more control. This mode might also be used for debugging purposes, as buffer overflow conditions are easier to diagnose than deadlock conditions.

Nonblocking message-passing operations, as described in Section 3.7, can be used to avoid the need for buffering outgoing messages. This prevents deadlocks due to lack of buffer space, and improves performance, by allowing overlap of computation and communication, 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 a buffer to be used for buffering messages sent in buffered mode. Buffering is done by the sender.

```
MPI_BUFFER_ATTACH( buffer, size)
```

```
IN      buffer          initial buffer address (choice)
```

```
IN      size           buffer size, in bytes (integer)
```

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

```

1 MPI_BUFFER_ATTACH( BUFFER, SIZE, IERROR)
2     <type> BUFFER(*)
3     INTEGER SIZE, IERROR

```

Provides to MPI a buffer in the user's memory to be used for buffering outgoing messages. The buffer is used only by messages sent in buffered mode. Only one buffer can be attached to a process at a time.

```

9 MPI_BUFFER_DETACH( buffer_addr, size)
10
11     OUT     buffer_addr           initial buffer address (choice)
12     OUT     size                 buffer size, in bytes (integer)

```

```

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

```

```

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

```

Detach the buffer currently associated with MPI. The call returns the address and the size of the detached buffer. This operation will block until all messages currently in the buffer have been transmitted. Upon return of this function, the user may reuse or deallocate the space taken by the buffer.

Example 3.10 Calls to attach and detach buffers.

```

25 #define BUFFSIZE 10000
26 int size
27 char *buff;
28 MPI_Buffer_attach( malloc(BUFFSIZE), BUFFSIZE);
29 /* a buffer of 10000 bytes can now be used by MPI_Bsend */
30 MPI_Buffer_detach( &buff, &size);
31 /* Buffer size reduced to zero */
32 MPI_Buffer_attach( buff, size);
33 /* Buffer of 10000 bytes available again */

```

Advice to users. Even though the C functions `MPI_Buffer_attach` and `MPI_Buffer_detach` both have a first argument of type `void*`, these arguments are used differently: A pointer to the buffer is passed to `MPI_Buffer_attach`; the address of the pointer is passed to `MPI_Buffer_detach`, so that this call can return the pointer value. (*End of advice to users.*)

Rationale. Both arguments are defined to be of type `void*` (rather than `void*` and `void**`, respectively), so as to avoid complex type casts. E.g., in the last example, `&buff`, which is of type `char**`, can be passed as argument to `MPI_Buffer_detach` without type casting. If the formal parameter had type `void**` then we would need a type cast before and after the call. (*End of rationale.*)

The statements made in this section describe the behavior of MPI for buffered-mode sends. When no buffer is currently associated, MPI behaves as if a zero-sized buffer is associated with the process.

MPI must provide as much buffering for outgoing messages *as if* outgoing message data were buffered by the sending process, in the specified buffer space, using a circular, contiguous-space allocation policy. We outline below a model implementation that defines this policy. MPI may provide more buffering, and may use a better buffer allocation algorithm than described below. On the other hand, MPI may signal an error whenever the simple buffering allocator described below would run out of space. In particular, if no buffer is explicitly associated with the process, then any buffered send may cause an error.

MPI does not provide mechanisms for querying or controlling buffering done by standard mode sends. It is expected that vendors will provide such information for their implementations.

Rationale. There is a wide spectrum of possible implementations of buffered communication: buffering can be done at sender, at receiver, or both; buffers can be dedicated to one sender-receiver pair, or be shared by all communications; buffering can be done in real or in virtual memory; it can use dedicated memory, or memory shared by other processes; buffer space may be allocated statically or be changed dynamically; etc. It does not seem feasible to provide a portable mechanism for querying or controlling buffering that would be compatible with all these choices, yet provide meaningful information. (*End of rationale.*)

3.6.1 Model implementation of buffered mode

The model implementation uses the packing and unpacking functions described in Section 3.13 and the nonblocking communication functions described in Section 3.7.

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

A buffered send call results in the execution of the following code.

- Traverse sequentially the PME queue from head towards the tail, deleting all entries for communications that have completed, up to the first entry with an uncompleted request; update queue head to point to that entry.
- Compute the number, n , of bytes needed to store an entry for the new message. An upper bound on n can be computed as follows: A call to the function `MPI_PACK_SIZE(count, datatype, comm, size)`, with the `count`, `datatype` and `comm` arguments used in the `MPI_BSEND` call, returns an upper bound on the amount of space needed to buffer the message data (see Section 3.13). The MPI constant `MPI_BSEND_OVERHEAD` provides an upper bound on the additional space consumed by the entry (e.g., for pointers or envelope information).
- 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.

- 1 • Post nonblocking send (standard mode) for packed data.
- 2
- 3 • Return
- 4

5 3.7 Nonblocking communication

7 One can improve performance on many systems by overlapping communication and com-
8 putation. This is especially true on systems where communication can be executed au-
9 tonomously by an intelligent communication controller. Light-weight threads are one mech-
10 anism for achieving such overlap. An alternative mechanism that often leads to better
11 performance is to use **nonblocking communication**. A nonblocking **send start** call ini-
12 tiates the send operation, but does not complete it. The send start call will return before
13 the message was copied out of the send buffer. A separate **send complete** call is needed
14 to complete the communication, i.e., to verify that the data has been copied out of the send
15 buffer. With suitable hardware, the transfer of data out of the sender memory may proceed
16 concurrently with computations done at the sender after the send was initiated and before it
17 completed. Similarly, a nonblocking **receive start call** initiates the receive operation, but
18 does not complete it. The call will return before a message is stored into the receive buffer.
19 A separate **receive complete** call is needed to complete the receive operation and verify
20 that the data has been received into the receive buffer. With suitable hardware, the transfer
21 of data into the receiver memory may proceed concurrently with computations done after
22 the receive was initiated and before it completed. The use of nonblocking receives may also
23 avoid system buffering and memory-to-memory copying, as information is provided early
24 on the location of the receive buffer.

25 Nonblocking send start calls can use the same four modes as blocking sends: **standard**,
26 **buffered**, **synchronous** and **ready**. These carry the same meaning. Sends of all modes, **ready**
27 excepted, can be started whether a matching receive has been posted or not; a nonblocking
28 **ready** send can be started only if a matching receive is posted. In all cases, the send start call
29 is local: it returns immediately, irrespective of the status of other processes. If the call causes
30 some system resource to be exhausted, then it will fail and return an error code. Quality
31 implementations of MPI should ensure that this happens only in “pathological” cases. That
32 is, an MPI implementation should be able to support a large number of pending nonblocking
33 operations.

34 The send-complete call returns when data has been copied out of the send buffer. It
35 may carry additional meaning, depending on the send mode.

36 If the send mode is **synchronous**, then the send can complete only if a matching receive
37 has started. That is, a receive has been posted, and has been matched with the send. In
38 this case, the send-complete call is non-local. Note that a synchronous, nonblocking send
39 may complete, if matched by a nonblocking receive, before the receive complete call occurs.
40 (It can complete as soon as the sender “knows” the transfer will complete, but before the
41 receiver “knows” the transfer will complete.)

42 If the send mode is **buffered** then the message must be buffered if there is no pending
43 receive. In this case, the send-complete call is local, and must succeed irrespective of the
44 status of a matching receive.

45 If the send mode is **standard** then the send-complete call may return before a matching
46 receive occurred, if the message is buffered. On the other hand, the send-complete may not
47 complete until a matching receive occurred, and the message was copied into the receive
48 buffer.

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

Advice to users. The completion of a send operation may be delayed, for standard mode, and must be delayed, for synchronous mode, until a matching receive is posted. The use of nonblocking sends in these two cases allows the sender to proceed ahead of the receiver, so that the computation is more tolerant of fluctuations in the speeds of the two processes.

Nonblocking sends in the buffered and ready modes have a more limited impact. A nonblocking send will return as soon as possible, whereas a blocking send will return after the data has been copied out of the sender memory. The use of nonblocking sends is advantageous in these cases only if data copying can be concurrent with computation.

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

3.7.1 Communication Objects

Nonblocking communications use opaque `request` objects to identify communication operations and match the operation that initiates the communication with the operation that terminates it. These are system objects that are accessed via a handle. A request object identifies various properties of a communication operation, such as the send mode, the communication buffer that is associated with it, its context, the tag and 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

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 a prefix of I (for immediate) indicates that the call is nonblocking.

```

1 MPI_ISEND(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 (integer)
5
6     IN      datatype           datatype of each send buffer element (handle)
7
8     IN      dest               rank of destination (integer)
9
10    IN      tag                message tag (integer)
11
12    IN      comm               communicator (handle)
13
14    OUT     request            communication request (handle)
15
16 int MPI_Isend(void* buf, int count, MPI_Datatype datatype, int dest,
17              int tag, MPI_Comm comm, MPI_Request *request)
18
19 MPI_ISEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
20     <type> BUF(*)
21     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
22
23     Start a standard mode, nonblocking send.
24
25 MPI_IBSEND(buf, count, datatype, dest, tag, comm, request)
26     IN      buf                initial address of send buffer (choice)
27
28     IN      count              number of elements in send buffer (integer)
29
30     IN      datatype           datatype of each send buffer element (handle)
31
32     IN      dest               rank of destination (integer)
33
34     IN      tag                message tag (integer)
35
36     IN      comm               communicator (handle)
37
38     OUT     request            communication request (handle)
39
40 int MPI_Ibsend(void* buf, int count, MPI_Datatype datatype, int dest,
41              int tag, MPI_Comm comm, MPI_Request *request)
42
43 MPI_IBSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
44     <type> BUF(*)
45     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
46
47     Start a buffered mode, nonblocking send.
48

```

MPI_ISSEND(buf, count, datatype, dest, tag, comm, request)			1
IN	buf	initial address of send buffer (choice)	2
			3
IN	count	number of elements in send buffer (integer)	4
IN	datatype	datatype of each send buffer element (handle)	5
			6
IN	dest	rank of destination (integer)	7
IN	tag	message tag (integer)	8
IN	comm	communicator (handle)	9
OUT	request	communication request (handle)	10
			11
			12
int MPI_Issend(void* buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm, MPI_Request *request)			13
			14
MPI_ISSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)			15
	<type> BUF(*)		16
	INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR		17
			18
	Start a synchronous mode, nonblocking send.		19
			20
			21
MPI_IRSEND(buf, count, datatype, dest, tag, comm, request)			22
IN	buf	initial address of send buffer (choice)	23
IN	count	number of elements in send buffer (integer)	24
IN	datatype	datatype of each send buffer element (handle)	25
			26
IN	dest	rank of destination (integer)	27
IN	tag	message tag (integer)	28
IN	comm	communicator (handle)	29
OUT	request	communication request (handle)	30
			31
			32
int MPI_Irsend(void* buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm, MPI_Request *request)			33
			34
MPI_IRSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)			35
	<type> BUF(*)		36
	INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR		37
			38
	Start a ready mode nonblocking send.		39
			40
			41
			42
			43
			44
			45
			46
			47
			48


```

1 MPI_RECV (buf, count, datatype, source, tag, comm, request)
2   OUT    buf                initial address of receive buffer (choice)
3
4   IN     count              number of elements in receive buffer (integer)
5
6   IN     datatype           datatype of each receive buffer element (handle)
7
8   IN     source              rank of source (integer)
9
10  IN     tag                  message tag (integer)
11
12  IN     comm                  communicator (handle)
13
14  OUT    request              communication request (handle)

```

```

15 int MPI_Irecv(void* buf, int count, MPI_Datatype datatype, int source,
16              int tag, MPI_Comm comm, MPI_Request *request)

```

```

17 MPI_RECV(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR)
18   <type> BUF(*)
19   INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR

```

Start a nonblocking receive.

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 access 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 subsections “Problems Due to Data Copying and Sequence Association,” and “A Problem with Register Optimization” in Section 10.2.2 of the MPI-2 Standard, pages 286 and 289. (*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 locations in 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 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. A **empty** status is a status which 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 (`MPI_TEST`, `MPI_TESTANY`, `MPI_WAIT`, `MPI_WAITANY`) that return a single `MPI_STATUS` value, the normal MPI error return process should be used (not the `MPI_ERROR` field in the `MPI_STATUS` argument).

`MPI_WAIT(request, status)`

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

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

`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 communication object associated with this request was created by a nonblocking send or receive call, then the object is deallocated by the call to `MPI_WAIT` and the request handle is set to `MPI_REQUEST_NULL`. `MPI_WAIT` is a non-local operation.

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 operation 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 attached with `MPI_BUFFER_ATTACH`. Note that, at this point, we can no longer cancel the send (see Sec. 3.8). If a matching receive is never posted, then the buffer cannot be freed. This runs somewhat counter to the stated goal of `MPI_CANCEL` (always being able to free program space that was committed to the communication subsystem). (*End of advice to users.*)

Advice to implementors. In a multi-threaded environment, a call to `MPI_WAIT` should block only the calling thread, allowing the thread scheduler to schedule another thread for execution. (*End of advice to implementors.*)

```

6
7 MPI_TEST(request, flag, status)

```

```

8     INOUT    request          communication request (handle)
9     OUT     flag              true if operation completed (logical)
10    OUT     status            status object (Status)

```

```

11
12
13 int MPI_Test(MPI_Request *request, int *flag, MPI_Status *status)

```

```

14 MPI_TEST(REQUEST, FLAG, STATUS, IERROR)
15     LOGICAL FLAG
16     INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR

```

A call to `MPI_TEST` returns `flag = true` if the operation identified by `request` is complete. In such a case, the status object is set to contain information on the completed operation; if the communication object was created by a nonblocking send or receive, then it is deallocated and the request handle is set to `MPI_REQUEST_NULL`. The call returns `flag = false`, otherwise. In this case, the value of the status object is undefined. `MPI_TEST` is a local operation.

The return status object for a receive operation carries information that can be accessed as described in section 3.2.5. The status object for a send operation carries information that can be accessed by a call to `MPI_TEST_CANCELLED` (see Section 3.8).

One is allowed to call `MPI_TEST` with a null or inactive `request` argument. In such a case the operation returns with `flag = true` and empty `status`.

The functions `MPI_WAIT` and `MPI_TEST` can be used to complete both sends and receives.

Advice to users. The use of the nonblocking `MPI_TEST` call allows the user to 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.*)

Rationale. The function `MPI_TEST` returns with `flag = true` exactly in those situations where the function `MPI_WAIT` returns; both functions return in such case the same value in `status`. Thus, a blocking `Wait` can be easily replaced by a nonblocking `Test`. (*End of rationale.*)

Example 3.11 Simple usage of nonblocking operations and `MPI_WAIT`.

```

41
42 CALL MPI_COMM_RANK(comm, rank, ierr)
43 IF(rank.EQ.0) THEN
44     CALL MPI_ISEND(a(1), 10, MPI_REAL, 1, tag, comm, request, ierr)
45     **** do some computation to mask latency ****
46     CALL MPI_WAIT(request, status, ierr)
47 ELSE
48

```



```

1 ELSE    ! rank.EQ.1
2     CALL MPI_IRecv(inval, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
3     CALL MPI_WAIT(req, status, ierr)
4     DO I=1, n-1
5         CALL MPI_ISEND(outval, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
6         CALL MPI_REQUEST_FREE(req, ierr)
7         CALL MPI_IRecv(inval, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
8         CALL MPI_WAIT(req, status, ierr)
9     END DO
10    CALL MPI_ISEND(outval, 1, MPI_REAL, 0, 0, MPI_COMM_WORLD, req, ierr)
11    CALL MPI_WAIT(req, status, ierr)
12 END IF

```

13

14 3.7.4 Semantics of Nonblocking Communications

15

16 The semantics of nonblocking communication is defined by suitably extending the definitions
17 in Section 3.5.

18

19 **Order** Nonblocking communication operations are ordered according to the execution order
20 of the calls that initiate the communication. The non-overtaking requirement of Section 3.5
21 is extended to nonblocking communication, with this definition of order being used.

22

23 **Example 3.13** Message ordering for nonblocking operations.

24

```

25 CALL MPI_COMM_RANK(comm, rank, ierr)
26 IF (RANK.EQ.0) THEN
27     CALL MPI_ISEND(a, 1, MPI_REAL, 1, 0, comm, r1, ierr)
28     CALL MPI_ISEND(b, 1, MPI_REAL, 1, 0, comm, r2, ierr)
29 ELSE    ! rank.EQ.1
30     CALL MPI_IRecv(a, 1, MPI_REAL, 0, MPI_ANY_TAG, comm, r1, ierr)
31     CALL MPI_IRecv(b, 1, MPI_REAL, 0, 0, comm, r2, ierr)
32 END IF
33 CALL MPI_WAIT(r1, status, ierr)
34 CALL MPI_WAIT(r2, status, ierr)

```

35 The first send of process zero will match the first receive of process one, even if both messages
36 are sent before process one executes either receive.

37

38 **Progress** A call to MPI_WAIT that completes a receive will eventually terminate and return
39 if a matching send has been started, unless the send is satisfied by another receive. In
40 particular, if the matching send is nonblocking, then the receive should complete even if
41 no call is executed by the sender to complete the send. Similarly, a call to MPI_WAIT that
42 completes a send will eventually return if a matching receive has been started, unless the
43 receive is satisfied by another send, and even if no call is executed to complete the receive.

44

45 **Example 3.14** An illustration of progress semantics.

46

```

47 CALL MPI_COMM_RANK(comm, rank, ierr)
48 IF (RANK.EQ.0) THEN

```

49

```

CALL MPI_SSEND(a, 1, MPI_REAL, 1, 0, comm, ierr)
CALL MPI_SEND(b, 1, MPI_REAL, 1, 1, comm, ierr)
ELSE    ! rank.EQ.1
CALL MPI_IRECV(a, 1, MPI_REAL, 0, 0, comm, r, ierr)
CALL MPI_RECV(b, 1, MPI_REAL, 0, 1, comm, ierr)
CALL MPI_WAIT(r, status, ierr)
END IF

```

This code should not deadlock in a correct MPI implementation. The first synchronous send of process zero must complete after process one posts the matching (nonblocking) receive even if process one has not yet reached the completing wait call. Thus, process zero will continue and execute the second send, allowing process one 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.

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 (integer)
INOUT	array_of_requests	array of requests (array of handles)
OUT	index	index of handle for operation that completed (integer)
OUT	status	status object (Status)

```

int MPI_Waitany(int count, MPI_Request *array_of_requests, int *index,
                MPI_Status *status)

```

```

MPI_WAITANY(COUNT, ARRAY_OF_REQUESTS, INDEX, STATUS, IERROR)
INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE),
IERROR

```

Blocks until one of the operations associated with the active requests in the array has completed. If more than one operation is enabled and can terminate, one is arbitrarily chosen. Returns in `index` the index of that request in the array and returns in `status` the status of the completing communication. (The array is indexed from zero in C, and from one in Fortran.) If the request was allocated by a nonblocking communication operation, then it is deallocated and the request handle is set to `MPI_REQUEST_NULL`.

1 The `array_of_requests` list may contain null or inactive handles. If the list contains no
 2 active handles (list has length zero or all entries are null or inactive), then the call returns
 3 immediately with `index = MPI_UNDEFINED`, and an empty `status`.

4 The execution of `MPI_WAITANY(count, array_of_requests, index, status)` has the same
 5 effect as the execution of `MPI_WAIT(&array_of_requests[i], status)`, where `i` is the value
 6 returned by `index` (unless the value of `index` is `MPI_UNDEFINED`). `MPI_WAITANY` with an
 7 array containing one active entry is equivalent to `MPI_WAIT`.
 8

9
 10 `MPI_TESTANY(count, array_of_requests, index, flag, status)`

11	IN	count	list length (integer)
12	INOUT	array_of_requests	array of requests (array of handles)
13			
14	OUT	index	index of operation that completed, or <code>MPI_UNDEFINED</code> if none completed (integer)
15			
16	OUT	flag	true if one of the operations is complete (logical)
17	OUT	status	status object (<code>Status</code>)
18			

19
 20 `int MPI_Testany(int count, MPI_Request *array_of_requests, int *index,`
 21 `int *flag, MPI_Status *status)`

22 `MPI_TESTANY(COUNT, ARRAY_OF_REQUESTS, INDEX, FLAG, STATUS, IERROR)`
 23 `LOGICAL FLAG`
 24 `INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE),`
 25 `IERROR`
 26

27 Tests for completion of either one or none of the operations associated with active
 28 handles. In the former case, it returns `flag = true`, returns in `index` the index of this request
 29 in the array, and returns in `status` the status of that operation; if the request was allocated
 30 by a nonblocking communication call then the request is deallocated and the handle is set
 31 to `MPI_REQUEST_NULL`. (The array is indexed from zero in C, and from one in Fortran.)
 32 In the latter case (no operation completed), it returns `flag = false`, returns a value of
 33 `MPI_UNDEFINED` in `index` and `status` is undefined.

34 The array may contain null or inactive handles. If the array contains no active handles
 35 then the call returns immediately with `flag = true`, `index = MPI_UNDEFINED`, and an empty
 36 `status`.

37 If the array of requests contains active handles then the execution of `MPI_TESTANY(count,`
 38 `array_of_requests, index, status)` has the same effect as the execution of `MPI_TEST(&ar-`
 39 `ray_of_requests[i], flag, status)`, for `i=0, 1, ..., count-1`, in some arbitrary order, until one call
 40 returns `flag = true`, or all fail. In the former case, `index` is set to the last value of `i`, and in
 41 the latter case, it is set to `MPI_UNDEFINED`. `MPI_TESTANY` with an array containing one
 42 active entry is equivalent to `MPI_TEST`.
 43

44 *Rationale.* The function `MPI_TESTANY` returns with `flag = true` exactly in those
 45 situations where the function `MPI_WAITANY` returns; both functions return in that
 46 case the same values in the remaining parameters. Thus, a blocking `MPI_WAITANY`
 47 can be easily replaced by a nonblocking `MPI_TESTANY`. The same relation holds for
 48 the other pairs of Wait and Test functions defined in this section. (*End of rationale.*)

```

MPI_WAITALL( count, array_of_requests, array_of_statuses) 1
    IN      count                lists length (integer) 2
    INOUT   array_of_requests    array of requests (array of handles) 3
    OUT     array_of_statuses     array of status objects (array of Status) 4
                                                5
                                                6

```

```

int MPI_Waitall(int count, MPI_Request *array_of_requests, 7
                MPI_Status *array_of_statuses) 8
                                                9

```

```

MPI_WAITALL(COUNT, ARRAY_OF_REQUESTS, ARRAY_OF_STATUSES, IERROR) 10
    INTEGER COUNT, ARRAY_OF_REQUESTS(*) 11
    INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), IERROR 12

```

Blocks until all communication operations associated with active handles in the list complete, and return 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. Requests that were created by nonblocking communication operations 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(count, array_of_requests, array_of_statuses)` has the same effect as the execution of `MPI_WAIT(&array_of_request[i], &array_of_statuses[i])`, for $i=0, \dots, \text{count}-1$, in some arbitrary order. `MPI_WAITALL` with an array of length one is equivalent to `MPI_WAIT`.

When one or more of the communications 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.*)

```

MPI_TESTALL(count, array_of_requests, flag, array_of_statuses) 33
    IN      count                lists length (integer) 34
    INOUT   array_of_requests    array of requests (array of handles) 35
    OUT     flag                 (logical) 36
    OUT     array_of_statuses     array of status objects (array of Status) 37
                                                38

```

```

int MPI_Testall(int count, MPI_Request *array_of_requests, int *flag, 39
                MPI_Status *array_of_statuses) 40
                                                41
                                                42
                                                43
                                                44
                                                45

```

```

int MPI_Testall(int count, MPI_Request *array_of_requests, int *flag, 46
                MPI_Status *array_of_statuses) 47
                                                48

```



```

1 MPI_TESTALL(COUNT, ARRAY_OF_REQUESTS, FLAG, ARRAY_OF_STATUSES, IERROR)
2     LOGICAL FLAG
3     INTEGER COUNT, ARRAY_OF_REQUESTS(*),
4     ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), IERROR

```

Returns `flag = true` if all communications 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 handle request is set to the status of the corresponding communication; if the request was allocated by a nonblocking communication call then it is deallocated, and the handle is 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 operation.

Errors that occurred during the execution of `MPI_TESTALL` are handled as errors in `MPI_WAITALL`.

```

17 MPI_WAITSSOME(incount, array_of_requests, outcount, array_of_indices, array_of_statuses)

```

18	IN	incount	length of array_of_requests (integer)
19	INOUT	array_of_requests	array of requests (array of handles)
20	OUT	outcount	number of completed requests (integer)
21	OUT	array_of_indices	array of indices of operations that completed (array of integers)
22	OUT	array_of_statuses	array of status objects for operations that completed (array of Status)
23			
24			
25			
26			
27			

```

28 int MPI_Waitssome(int incount, MPI_Request *array_of_requests, int *outcount,
29                 int *array_of_indices, MPI_Status *array_of_statuses)

```

```

30 MPI_WAITSSOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT, ARRAY_OF_INDICES,
31              ARRAY_OF_STATUSES, IERROR)

```

```

32     INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),
33     ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), IERROR

```

Waits until at least one of the operations associated with active handles in the list have completed. 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 (index 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_status` the status for these completed operations. If a request that completed was allocated by a nonblocking communication call, then it is deallocated, and the associated handle is set to `MPI_REQUEST_NULL`.

If the list contains no active handles, then the call returns immediately with `outcount = MPI_UNDEFINED`.

When one or more of the communications completed by `MPI_WAITSSOME` fails, then it is desirable to return specific information on each communication. The arguments `outcount`, `array_of_indices` and `array_of_statuses` will be adjusted to indicate completion of all communications that have succeeded or failed. The call will return the error code `MPI_ERR_IN_STATUS`

and the error field of each status returned will be set to indicate success or to indicate the specific error that occurred. The call will return `MPI_SUCCESS` if no request resulted in an error, and 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.

```

MPI_TESTSOME(incount, array_of_requests, outcount, array_of_indices, array_of_statuses)
IN      incount          length of array_of_requests (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)

```

```

int MPI_Testsome(int incount, MPI_Request *array_of_requests, int *outcount,
                 int *array_of_indices, MPI_Status *array_of_statuses)

```

```

MPI_TESTSOME(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

```

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 operation, 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 posted, 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 communications, 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 communications as possible. (*End of advice to implementors.*)

Example 3.15 Client-server code (starvation can occur).

```

1  CALL MPI_COMM_SIZE(comm, size, ierr)
2
3  CALL MPI_COMM_RANK(comm, rank, ierr)
4  IF(rank > 0) THEN          ! client code
5      DO WHILE(.TRUE.)
6          CALL MPI_ISEND(a, n, MPI_REAL, 0, tag, comm, request, ierr)
7          CALL MPI_WAIT(request, status, ierr)
8      END DO
9  ELSE          ! rank=0 -- server code
10     DO i=1, size-1
11         CALL MPI_Irecv(a(1,i), n, MPI_REAL, i tag,
12             comm, request_list(i), ierr)
13     END DO
14     DO WHILE(.TRUE.)
15         CALL MPI_WAITANY(size-1, request_list, index, status, ierr)
16         CALL DO_SERVICE(a(1,index)) ! handle one message
17         CALL MPI_Irecv(a(1, index), n, MPI_REAL, index, tag,
18             comm, request_list(index), ierr)
19     END DO
20 END IF
21
22
23

```

Example 3.16 Same code, using MPI_WAITSOME.

```

24  CALL MPI_COMM_SIZE(comm, size, ierr)
25  CALL MPI_COMM_RANK(comm, rank, ierr)
26  IF(rank > 0) THEN          ! client code
27      DO WHILE(.TRUE.)
28          CALL MPI_ISEND(a, n, MPI_REAL, 0, tag, comm, request, ierr)
29          CALL MPI_WAIT(request, status, ierr)
30      END DO
31  ELSE          ! rank=0 -- server code
32      DO i=1, size-1
33         CALL MPI_Irecv(a(1,i), n, MPI_REAL, i, tag,
34             comm, requests(i), ierr)
35     END DO
36     DO WHILE(.TRUE.)
37         CALL MPI_WAITSOME(size, request_list, numdone,
38             indices, statuses, ierr)
39         DO i=1, numdone
40             CALL DO_SERVICE(a(1, indices(i)))
41             CALL MPI_Irecv(a(1, indices(i)), n, MPI_REAL, 0, tag,
42                 comm, requests(indices(i)), ierr)
43         END DO
44     END DO
45 END IF
46
47
48

```

3.8 Probe and Cancel

The `MPI_PROBE` and `MPI_IPROBE` operations allow incoming messages to be checked for, without actually receiving them. The user can then decide how to receive them, based on the information returned by the probe (basically, the information returned by `status`). In particular, the user may allocate memory for the receive buffer, according to the length of the probed message.

The `MPI_CANCEL` operation allows pending communications to be canceled. This is required for cleanup. Posting a send or a receive ties up user resources (send or receive buffers), and a cancel may be needed to free these resources gracefully.

`MPI_IPROBE(source, tag, comm, flag, status)`

IN	source	source rank, or <code>MPI_ANY_SOURCE</code> (integer)
IN	tag	tag value or <code>MPI_ANY_TAG</code> (integer)
IN	comm	communicator (handle)
OUT	flag	(logical)
OUT	status	status object (<code>Status</code>)

```
int MPI_Iprobe(int source, int tag, MPI_Comm comm, int *flag,
               MPI_Status *status)
```

`MPI_IPROBE(SOURCE, TAG, COMM, FLAG, STATUS, IERROR)`

LOGICAL FLAG
 INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR

`MPI_IPROBE(source, tag, comm, flag, status)` 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(..., source, tag, comm, 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.

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 process is multi-threaded, it is the user's responsibility to ensure that the last condition holds.

The `source` argument of `MPI_PROBE` 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.

```

1 MPI_PROBE(source, tag, comm, status)
2     IN      source      source rank, or MPI_ANY_SOURCE (integer)
3     IN      tag         tag value, or MPI_ANY_TAG (integer)
4     IN      comm        communicator (handle)
5     OUT     status      status object (Status)

```

```

8
9 int MPI_Probe(int source, int tag, MPI_Comm comm, MPI_Status *status)

```

```

10 MPI_PROBE(SOURCE, TAG, COMM, STATUS, IERROR)
11     INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR

```

MPI_PROBE behaves like MPI_IPROBE except that it is a blocking 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 a process, and a send that matches the probe has been initiated by some 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 process). Similarly, if a process busy waits with MPI_IPROBE and a matching message has been issued, then the call to MPI_IPROBE will eventually return flag = true unless the message is received by another concurrent receive operation.

Example 3.17 Use blocking probe to wait for an incoming message.

```

24     CALL MPI_COMM_RANK(comm, rank, ierr)
25     IF (rank.EQ.0) THEN
26         CALL MPI_SEND(i, 1, MPI_INTEGER, 2, 0, comm, ierr)
27     ELSE IF(rank.EQ.1) THEN
28         CALL MPI_SEND(x, 1, MPI_REAL, 2, 0, comm, ierr)
29     ELSE ! rank.EQ.2
30         DO i=1, 2
31             CALL MPI_PROBE(MPI_ANY_SOURCE, 0,
32                             comm, status, ierr)
33             IF (status(MPI_SOURCE) .EQ. 0) THEN
34 100         CALL MPI_RECV(i, 1, MPI_INTEGER, 0, 0, comm, status, ierr)
35             ELSE
36 200         CALL MPI_RECV(x, 1, MPI_REAL, 1, 0, comm, status, ierr)
37             END IF
38         END DO
39     END IF
40

```

Each message is received with the right type.

Example 3.18 A similar program to the previous example, but now it has a problem.

```

45     CALL MPI_COMM_RANK(comm, rank, ierr)
46     IF (rank.EQ.0) THEN
47         CALL MPI_SEND(i, 1, MPI_INTEGER, 2, 0, comm, ierr)
48     ELSE IF(rank.EQ.1) THEN

```

```

        CALL MPI_SEND(x, 1, MPI_REAL, 2, 0, comm, ierr)          1
ELSE                                           2
    DO i=1, 2                                           3
        CALL MPI_PROBE(MPI_ANY_SOURCE, 0,                4
                        comm, status, ierr)              5
        IF (status(MPI_SOURCE) .EQ. 0) THEN            6
100         CALL MPI_RECV(i, 1, MPI_INTEGER, MPI_ANY_SOURCE, 7
                        0, comm, status, ierr)          8
        ELSE                                           9
200         CALL MPI_RECV(x, 1, MPI_REAL, MPI_ANY_SOURCE, 10
                        0, comm, status, ierr)          11
        END IF                                         12
    END DO                                             13
END IF                                             14

```

We slightly modified example 3.17, using `MPI_ANY_SOURCE` as the source argument in the two receive calls in statements labeled 100 and 200. 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 implementors. A call to `MPI_PROBE(source, tag, comm, status)` will match the message that would have been received by a call to `MPI_RECV(..., source, tag, comm, 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.*)

`MPI_CANCEL(request)`

IN request communication request (handle)

int MPI_Cancel(MPI_Request *request)

`MPI_CANCEL(REQUEST, IERROR)`

INTEGER REQUEST, IERROR

A call to `MPI_CANCEL` marks for cancellation a pending, nonblocking communication operation (send or receive). The cancel call is local. It returns immediately, possibly before the communication is actually canceled. It is still necessary to complete a communication that has been marked for cancellation, using a call to `MPI_REQUEST_FREE`, `MPI_WAIT` or `MPI_TEST` (or any of the derived operations).

1 If a communication is marked for cancellation, then a `MPI_WAIT` call for that com-
 2 munication is guaranteed to return, irrespective of the activities of other processes (i.e.,
 3 `MPI_WAIT` behaves as a local function); similarly if `MPI_TEST` is repeatedly called in a
 4 busy wait loop for a canceled communication, then `MPI_TEST` will eventually be successful.

5 `MPI_CANCEL` can be used to cancel a communication that uses a persistent request (see
 6 Sec. 3.9), in the same way it is used for nonpersistent requests. A successful cancellation
 7 cancels the active communication, but not the request itself. After the call to `MPI_CANCEL`
 8 and the subsequent call to `MPI_WAIT` or `MPI_TEST`, the request becomes inactive and can
 9 be activated for a new communication.

10 The successful cancellation of a buffered send frees the buffer space occupied by the
 11 pending message.

12 Either the cancellation succeeds, or the communication succeeds, but not both. If a
 13 send is marked for cancellation, then it must be the case that either the send completes
 14 normally, in which case the message sent was received at the destination process, or that
 15 the send is successfully canceled, in which case no part of the message was received at the
 16 destination. Then, any matching receive has to be satisfied by another send. If a receive is
 17 marked for cancellation, then it must be the case that either the receive completes normally,
 18 or that the receive is successfully canceled, in which case no part of the receive buffer is
 19 altered. Then, any matching send has to be satisfied by another receive.

20 If the operation has been canceled, then information to that effect will be returned in
 21 the status argument of the operation that completes the communication.

22
 23
 24 `MPI_TEST_CANCELLED(status, flag)`

25	IN	status	status object (Status)
26	OUT	flag	(logical)

27
 28
 29 `int MPI_Test_cancelled(MPI_Status *status, int *flag)`

30 `MPI_TEST_CANCELLED(STATUS, FLAG, IERROR)`

31 LOGICAL FLAG

32 INTEGER STATUS(MPI_STATUS_SIZE), IERROR

33
 34 Returns `flag = true` if the communication associated with the status object was canceled
 35 successfully. In such a case, all other fields of `status` (such as `count` or `tag`) are undefined.
 36 Returns `flag = false`, otherwise. If a receive operation might be canceled then one should call
 37 `MPI_TEST_CANCELLED` first, to check whether the operation was canceled, before checking
 38 on the other fields of the return status.

39 *Advice to users.* Cancel can be an expensive operation that should be used only
 40 exceptionally. (*End of advice to users.*)

41
 42 *Advice to implementors.* If a send operation uses an “eager” protocol (data is trans-
 43 ferred to the receiver before a matching receive is posted), then the cancellation of this
 44 send may require communication with the intended receiver in order to free allocated
 45 buffers. On some systems this may require an interrupt to the intended receiver. Note
 46 that, while communication may be needed to implement `MPI_CANCEL`, this is still a
 47 local operation, since its completion does not depend on the code executed by other
 48 processes. If processing is required on another process, this should be transparent to

the application (hence the need for an interrupt and an interrupt handler). (*End of advice to implementors.*)

3.9 Persistent communication requests

Often a communication with the same argument list is repeatedly executed within the inner loop of a parallel computation. In such a situation, it may be possible to optimize the communication by binding the list of communication arguments to a **persistent** communication request once and, then, repeatedly using the request to initiate and complete messages. The persistent request thus created can be thought of as a communication port or a “half-channel.” It does not provide the full functionality of a conventional channel, since there is no binding of the send port to the receive port. This construct allows reduction of the overhead for communication between the process and communication controller, but not of the overhead for communication between one communication controller and another. It is not necessary that messages sent with a persistent request be received by a receive operation using a persistent request, or vice versa.

A persistent communication request is created using one of the four following calls. These calls involve no communication.

`MPI_SEND_INIT(buf, count, datatype, dest, tag, comm, request)`

IN	buf	initial address of send buffer (choice)
IN	count	number of elements sent (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)

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

```
MPI_SEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
<type> BUF(*)
INTEGER REQUEST, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
```

Creates a persistent communication request for a standard mode send operation, and binds to it all the arguments of a send operation.


```

1 MPI_BSEND_INIT(buf, count, datatype, dest, tag, comm, request)
2     IN      buf                initial address of send buffer (choice)
3
4     IN      count              number of elements sent (integer)
5
6     IN      datatype           type of each element (handle)
7
8     IN      dest               rank of destination (integer)
9
10    IN      tag                 message tag (integer)
11
12    IN      comm                communicator (handle)
13
14    OUT     request             communication request (handle)
15
16 int MPI_Bsend_init(void* buf, int count, MPI_Datatype datatype, int dest,
17                   int tag, MPI_Comm comm, MPI_Request *request)
18
19 MPI_BSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
20     <type> BUF(*)
21     INTEGER REQUEST, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
22
23     Creates a persistent communication request for a buffered mode send.
24
25 MPI_SSEND_INIT(buf, count, datatype, dest, tag, comm, request)
26     IN      buf                initial address of send buffer (choice)
27
28     IN      count              number of elements sent (integer)
29
30     IN      datatype           type of each element (handle)
31
32     IN      dest               rank of destination (integer)
33
34     IN      tag                 message tag (integer)
35
36     IN      comm                communicator (handle)
37
38     OUT     request             communication request (handle)
39
40 int MPI_Ssend_init(void* buf, int count, MPI_Datatype datatype, int dest,
41                   int tag, MPI_Comm comm, MPI_Request *request)
42
43 MPI_SSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
44     <type> BUF(*)
45     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
46
47     Creates a persistent communication object for a synchronous mode send operation.
48

```



```

1 MPI_START(request)
2     INOUT    request                communication request (handle)
3

```

```

4
5 int MPI_Start(MPI_Request *request)

```

```

6 MPI_START(REQUEST, IERROR)
7     INTEGER REQUEST, IERROR
8

```

The argument, `request`, is a handle returned by one of the previous five calls. The associated request should be inactive. The request becomes active once the call is made.

If the request is for a send with ready mode, then a matching receive should be posted before the call is made. The communication buffer should not be accessed 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.

```

19
20 MPI_STARTALL(count, array_of_requests)
21
22     IN        count                list length (integer)
23     INOUT    array_of_requests    array of requests (array of handle)
24

```

```

25 int MPI_Startall(int count, MPI_Request *array_of_requests)
26

```

```

27 MPI_STARTALL(COUNT, ARRAY_OF_REQUESTS, IERROR)
28     INTEGER COUNT, ARRAY_OF_REQUESTS(*), IERROR
29

```

Start all communications associated with requests in `array_of_requests`. A call to `MPI_STARTALL(count, array_of_requests)` has the same effect as calls to `MPI_START (&array_of_requests[i])`, executed for $i=0, \dots, \text{count}-1$, in some arbitrary order.

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 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. 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 communication object is used in several concurrent threads, it is the user's responsibility to coordinate calls so that the

correct sequence is obeyed.

A send operation initiated with `MPI_START` can be matched with any receive operation and, likewise, a receive operation initiated 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 subsections “Problems Due to Data Copying and Sequence Association,” and “A Problem with Register Optimization” in Section 10.2.2 of the MPI-2 Standard, pages 286 and 289. (*End of advice to users.*)

3.10 Send-receive

The **send-receive** operations combine in one call the sending of a message to one destination and the receiving of another message, from another process. The two (source and destination) are possibly the same. A send-receive operation is very useful for executing a shift operation across a chain of processes. If blocking sends and receives are used for such a shift, then one needs to order the sends and receives correctly (for example, even processes send, then receive, odd processes receive first, then send) so as to prevent cyclic dependencies that may lead to deadlock. When a send-receive operation is used, the communication subsystem takes care of these issues. The send-receive operation can be used in conjunction with the functions described in Chapter 6 in order to perform shifts on various logical topologies. Also, a send-receive operation is useful for implementing remote procedure calls.

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

`MPI_SENDRECV`(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 (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 (integer)
IN	recvtype	type of elements in receive buffer (handle)
IN	source	rank of source (integer)
IN	recvtag	receive tag (integer)
IN	comm	communicator (handle)
OUT	status	status object (Status)

`int MPI_Sendrecv(void *sendbuf, int sendcount, MPI_Datatype sendtype,`

```

1         int dest, int sendtag, void *recvbuf, int recvcount,
2         MPI_Datatype recvtype, int source, int recvtag, MPI_Comm comm,
3         MPI_Status *status)
4
5 MPI_SENDRECV(SENDBUF, SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVBUF,
6             RECVCOUNT, RECVTYPE, SOURCE, RECVTAG, COMM, STATUS, IERROR)
7     <type> SENDBUF(*), RECVBUF(*)
8     INTEGER SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVCOUNT, RECVTYPE,
9     SOURCE, RECVTAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR

```

10 Execute a blocking send and receive operation. Both send and receive use the same
11 communicator, but possibly different tags. The send buffer and receive buffers must be
12 disjoint, and may have different lengths and datatypes.

```

14
15 MPI_SENDRECV_REPLACE(buf, count, datatype, dest, sendtag, source, recvtag, comm, sta-
16 tus)

```

17	INOUT	buf	initial address of send and receive buffer (choice)
18	IN	count	number of elements in send and receive buffer (integer)
19	IN	datatype	type of elements in send and receive buffer (handle)
20	IN	dest	rank of destination (integer)
21	IN	sendtag	send message tag (integer)
22	IN	source	rank of source (integer)
23	IN	recvtag	receive message tag (integer)
24	IN	comm	communicator (handle)
25	OUT	status	status object (Status)

```

26
27
28
29
30 int MPI_Sendrecv_replace(void* buf, int count, MPI_Datatype datatype,
31     int dest, int sendtag, int source, int recvtag, MPI_Comm comm,
32     MPI_Status *status)

```

```

33
34 MPI_SENDRECV_REPLACE(BUF, COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG,
35     COMM, STATUS, IERROR)
36     <type> BUF(*)
37     INTEGER COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG, COMM,
38     STATUS(MPI_STATUS_SIZE), IERROR

```

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

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

44
45 *Advice to implementors.* Additional intermediate buffering is needed for the “replace”
46 variant. (*End of advice to implementors.*)

47
48

3.11 Null 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 non-circular 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 process `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`.

3.12 Derived datatypes

Up to here, all point to point communication have involved only contiguous buffers containing a sequence of elements of the same type. 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 back 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 shape and size. 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.

1 Let

$$2 \quad \text{Typemap} = \{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

3 be such a type map, where $type_i$ are basic types, and $disp_i$ are displacements. Let

$$4 \quad \text{Typesig} = \{type_0, \dots, type_{n-1}\}$$

5 be the associated type signature. This type map, together with a base address buf , specifies
6 a communication buffer: the communication buffer that consists of n entries, where the
7 i -th entry is at address $buf + disp_i$ and has type $type_i$. A message assembled from such a
8 communication buffer will consist of n values, of the types defined by Typesig .

9 We can use a handle to a general datatype as an argument in a send or receive operation,
10 instead of a basic datatype argument. The operation `MPI_SEND(buf, 1, datatype, ...)` will use
11 the send buffer defined by the base address `buf` and the general datatype associated with
12 `datatype`; it will generate a message with the type signature determined by the `datatype`
13 argument. `MPI_RECV(buf, 1, datatype, ...)` will use the receive buffer defined by the base
14 address `buf` and the general datatype associated with `datatype`.

15 General datatypes can be used in all send and receive operations. We discuss, in Sec.
16 [3.12.5](#), the case where the second argument `count` has value > 1 .

17 The basic datatypes presented in section [3.2.2](#) are particular cases of a general datatype,
18 and are predefined. Thus, `MPI_INT` is a predefined handle to a datatype with type map
19 $\{(\text{int}, 0)\}$, with one entry of type `int` and displacement zero. The other basic datatypes are
20 similar.

21 The **extent** of a datatype is defined to be the span from the first byte to the last byte
22 occupied by entries in this datatype, rounded up to satisfy alignment requirements. That
23 is, if

$$24 \quad \text{Typemap} = \{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

25 then

$$26 \quad \begin{aligned} 27 \quad lb(\text{Typemap}) &= \min_j disp_j, \\ 28 \quad ub(\text{Typemap}) &= \max_j (disp_j + sizeof(type_j)) + \epsilon, \text{ and} \\ 29 \quad extent(\text{Typemap}) &= ub(\text{Typemap}) - lb(\text{Typemap}). \end{aligned} \quad (3.1)$$

30 If $type_i$ requires alignment to a byte address that is a multiple of k_i , then ϵ is the least
31 nonnegative increment needed to round $extent(\text{Typemap})$ to the next multiple of $\max_i k_i$.
32 The complete definition of **extent** is given on page [73](#).

33 **Example 3.19** Assume that $\text{Type} = \{(\text{double}, 0), (\text{char}, 8)\}$ (a double at displacement zero,
34 followed by a char at displacement eight). Assume, furthermore, that doubles have to be
35 strictly aligned at addresses that are multiples of eight. Then, the extent of this datatype is
36 16 (9 rounded to the next multiple of 8). A datatype that consists of a character immediately
37 followed by a double will also have an extent of 16.

38 *Rationale.* The definition of extent is motivated by the assumption that the amount
39 of padding added at the end of each structure in an array of structures is the least
40 needed to fulfill alignment constraints. More explicit control of the extent is provided
41 in section [3.12.3](#). Such explicit control is needed in cases where the assumption does
42 not hold, for example, where union types are used. (*End of rationale.*)

3.12.1 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 (nonnegative integer)

IN oldtype old datatype (handle)

OUT newtype new datatype (handle)

```
int MPI_Type_contiguous(int count, MPI_Datatype oldtype,
                        MPI_Datatype *newtype)
```

```
MPI_TYPE_CONTIGUOUS(COUNT, OLDTYPE, NEWTYPE, IERROR)
INTEGER COUNT, OLDTYPE, NEWTYPE, IERROR
```

`newtype` is the datatype obtained by concatenating `count` copies of `oldtype`. Concatenation is defined using *extent* as the size of the concatenated copies.

Example 3.20 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 · 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 function `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.


```

1 MPI_TYPE_VECTOR( count, blocklength, stride, oldtype, newtype)
2     IN          count          number of blocks (nonnegative integer)
3
4     IN          blocklength    number of elements in each block (nonnegative inte-
5                                ger)
6
7     IN          stride         number of elements between start of each block (inte-
8                                ger)
9
10    IN          oldtype        old datatype (handle)
11
12    OUT         newtype        new datatype (handle)

```

```

13 int MPI_Type_vector(int count, int blocklength, int stride,
14                   MPI_Datatype oldtype, MPI_Datatype *newtype)
15 MPI_TYPE_VECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR)
16 INTEGER COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR

```

Example 3.21 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 blocks.

Example 3.22 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:

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1}), \\ (type_0, disp_0 + ex), \dots, (type_{n-1}, disp_{n-1} + ex), \dots, \\ (type_0, disp_0 + (bl - 1) \cdot ex), \dots, (type_{n-1}, disp_{n-1} + (bl - 1) \cdot ex), \\ (type_0, disp_0 + stride \cdot ex), \dots, (type_{n-1}, disp_{n-1} + stride \cdot ex), \dots, \\ (type_0, disp_0 + (stride + bl - 1) \cdot ex), \dots, (type_{n-1}, disp_{n-1} + (stride + bl - 1) \cdot ex), \dots,$$

$(type_0, disp_0 + stride \cdot (count - 1) \cdot ex), \dots,$

$(type_{n-1}, disp_{n-1} + stride \cdot (count - 1) \cdot ex), \dots,$

$(type_0, disp_0 + (stride \cdot (count - 1) + bl - 1) \cdot ex), \dots,$

$(type_{n-1}, disp_{n-1} + (stride \cdot (count - 1) + bl - 1) \cdot ex)\}.$

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)`, n arbitrary.

Hvector The function `MPI_TYPE_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 Sec. 3.12.7. (H stands for “heterogeneous”).

`MPI_TYPE_HVECTOR(count, blocklength, stride, oldtype, newtype)`

IN	count	number of blocks (nonnegative integer)
IN	blocklength	number of elements in each block (nonnegative integer)
IN	stride	number of bytes between start of each block (integer)
IN	oldtype	old datatype (handle)
OUT	newtype	new datatype (handle)

```
int MPI_Type_hvector(int count, int blocklength, MPI_Aint stride,
                    MPI_Datatype oldtype, MPI_Datatype *newtype)
```

```
MPI_TYPE_HVECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDDTYPE, NEWTYPE, IERROR)
INTEGER COUNT, BLOCKLENGTH, STRIDE, OLDDTYPE, NEWTYPE, IERROR
```

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:

$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1}),$

$(type_0, disp_0 + ex), \dots, (type_{n-1}, disp_{n-1} + ex), \dots,$

$(type_0, disp_0 + (bl - 1) \cdot ex), \dots, (type_{n-1}, disp_{n-1} + (bl - 1) \cdot ex),$

$(type_0, disp_0 + stride), \dots, (type_{n-1}, disp_{n-1} + stride), \dots,$

$(type_0, disp_0 + stride + (bl - 1) \cdot ex), \dots,$

```

1      ( $type_{n-1}, disp_{n-1} + stride + (bl - 1) \cdot ex$ ), ...,
2
3      ( $type_0, disp_0 + stride \cdot (count - 1)$ ), ..., ( $type_{n-1}, disp_{n-1} + stride \cdot (count - 1)$ ), ...,
4
5      ( $type_0, disp_0 + stride \cdot (count - 1) + (bl - 1) \cdot ex$ ), ...,
6
7      ( $type_{n-1}, disp_{n-1} + stride \cdot (count - 1) + (bl - 1) \cdot ex$ )}.

```

Indexed The function `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.

```

17 MPI_TYPE_INDEXED( count, array_of_blocklengths, array_of_displacements, oldtype, newtype)

```

19	IN	count	number of blocks – also number of entries in array_of_displacements and array_of_blocklengths (non-negative integer)
23	IN	array_of_blocklengths	number of elements per block (array of nonnegative integers)
25	IN	array_of_displacements	displacement for each block, in multiples of oldtype extent (array of integer)
27	IN	oldtype	old datatype (handle)
28	OUT	newtype	new datatype (handle)

```

31 int MPI_Type_indexed(int count, int *array_of_blocklengths,
32                     int *array_of_displacements, MPI_Datatype oldtype,
33                     MPI_Datatype *newtype)

```

```

34 MPI_TYPE_INDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
35                 OLDTYPE, NEWTYPE, IERROR)
36 INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*),
37 OLDTYPE, NEWTYPE, IERROR

```

Example 3.23 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,

```

43      {(double, 64), (char, 72), (double, 80), (char, 88), (double, 96), (char, 104),
44
45      (double, 0), (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_blocklength` 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:

$$\begin{aligned} &\{(type_0, disp_0 + D[0] \cdot ex), \dots, (type_{n-1}, disp_{n-1} + D[0] \cdot ex), \dots, \\ &(type_0, disp_0 + (D[0] + B[0] - 1) \cdot ex), \dots, (type_{n-1}, disp_{n-1} + (D[0] + B[0] - 1) \cdot ex), \dots, \\ &(type_0, disp_0 + D[\text{count} - 1] \cdot ex), \dots, (type_{n-1}, disp_{n-1} + D[\text{count} - 1] \cdot ex), \dots, \\ &(type_0, disp_0 + (D[\text{count} - 1] + B[\text{count} - 1] - 1) \cdot ex), \dots, \\ &(type_{n-1}, disp_{n-1} + (D[\text{count} - 1] + B[\text{count} - 1] - 1) \cdot ex)\}. \end{aligned}$$

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 function `MPI_TYPE_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.

`MPI_TYPE_HINDEXED(count, array_of_blocklengths, array_of_displacements, oldtype, newtype)`

IN	count	number of blocks – also number of entries in <code>array_of_displacements</code> and <code>array_of_blocklengths</code> (integer)
IN	array_of_blocklengths	number of elements in each block (array of nonnegative integers)
IN	array_of_displacements	byte displacement of each block (array of integer)
IN	oldtype	old datatype (handle)
OUT	newtype	new datatype (handle)

```
int MPI_Type_hindexed(int count, int *array_of_blocklengths,
                     MPI_Aint *array_of_displacements, MPI_Datatype oldtype,
                     MPI_Datatype *newtype)
```

```
MPI_TYPE_HINDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
                  OLDTYPE, NEWTYPE, IERROR)
```

```

1     INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*),
2     OLDTYPE, NEWTYPE, IERROR
3
4

```

Assume that oldtype has type map,

```

6     {(type0, disp0), ..., (typen-1, dispn-1)},
7

```

with extent *ex*. Let **B** be the array_of_blocklength 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:

```

11     {(type0, disp0 + D[0]), ..., (typen-1, dispn-1 + D[0]), ...,
12     (type0, disp0 + D[0] + (B[0] - 1) · ex), ...,
13     (typen-1, dispn-1 + D[0] + (B[0] - 1) · ex), ...,
14     (type0, disp0 + D[count - 1]), ..., (typen-1, dispn-1 + D[count - 1]), ...,
15     (type0, disp0 + D[count - 1] + (B[count - 1] - 1) · ex), ...,
16     (typen-1, dispn-1 + D[count - 1] + (B[count - 1] - 1) · ex)}.
17
18
19
20
21
22
23

```

Struct MPI_Type_Struct is the most general type constructor. It further generalizes the previous one in that it allows each block to consist of replications of different datatypes.

```

24
25
26
27
28 MPI_Type_Struct(count, array_of_blocklengths, array_of_displacements, array_of_types, new-
29 type)

```

30	IN	count	number of blocks (integer) – also number of entries in arrays array_of_types, array_of_displacements and array_of_blocklengths
31			
32			
33			
34	IN	array_of_blocklength	number of elements in each block (array of integer)
35	IN	array_of_displacements	byte displacement of each block (array of integer)
36	IN	array_of_types	type of elements in each block (array of handles to datatype objects)
37			
38			
39	OUT	newtype	new datatype (handle)
40			

```

41 int MPI_Type_struct(int count, int *array_of_blocklengths,
42                    MPI_Aint *array_of_displacements, MPI_Datatype *array_of_types,
43                    MPI_Datatype *newtype)
44

```

```

45 MPI_Type_Struct(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
46                ARRAY_OF_TYPES, NEWTYPE, IERROR)
47
48 INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*),
49 ARRAY_OF_TYPES(*), NEWTYPE, IERROR

```

Example 3.24 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_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. (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_HINDEXED(count, B, D, oldtype, newtype)` is equivalent to a call to `MPI_TYPE_STRUCT(count, B, D, T, newtype)`, where each entry of T is equal to `oldtype`.

3.12.2 Address and extent functions

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`.

The address of a location in memory can be found by invoking the function `MPI_ADDRESS`.

`MPI_ADDRESS(location, address)`

IN	location	location in caller memory (choice)
OUT	address	address of location (integer)

`int MPI_Address(void* location, MPI_Aint *address)`

```

1 MPI_ADDRESS(LOCATION, ADDRESS, IERROR)
2   <type> LOCATION(*)
3   INTEGER ADDRESS, IERROR

```

Returns the (byte) address of location.

Example 3.25 Using MPI_ADDRESS for an array.

```

7   REAL A(100,100)
8   INTEGER I1, I2, DIFF
9   CALL MPI_ADDRESS(A(1,1), I1, IERROR)
10  CALL MPI_ADDRESS(A(10,10), I2, IERROR)
11  DIFF = I2 - I1
12  ! The value of DIFF is 909*sizeofreal; the values of I1 and I2 are
13  ! implementation dependent.

```

Advice to users. C users may be tempted to avoid the usage of MPI_ADDRESS and rely on the availability of the address operator &. Note, however, that & *cast-expression* is a pointer, not an address. ANSI 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_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 subsections “Problems Due to Data Copying and Sequence Association,” and “A Problem with Register Optimization” in Section 10.2.2 of the MPI-2 Standard, pages 286 and 289. (*End of advice to users.*)

The following auxiliary functions provide useful information on derived datatypes.

```

31 MPI_TYPE_EXTENT(datatype, extent)

```

```

33   IN      datatype          datatype (handle)
34   OUT     extent            datatype extent (integer)

```

```

36 int MPI_Type_extent(MPI_Datatype datatype, MPI_Aint *extent)

```

```

38 MPI_TYPE_EXTENT(DATATYPE, EXTENT, IERROR)
39   INTEGER DATATYPE, EXTENT, IERROR

```

Returns the extent of a datatype, where extent is as defined on page 73.

```

43 MPI_TYPE_SIZE(datatype, size)

```

```

45   IN      datatype          datatype (handle)
46   OUT     size              datatype size (integer)

```

```

48 int MPI_Type_size(MPI_Datatype datatype, int *size)

```

MPI_TYPE_SIZE(DATATYPE, SIZE, IERROR)
 INTEGER DATATYPE, SIZE, IERROR

MPI_TYPE_SIZE returns 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.

Advice to users. The MPI-1 Standard specifies that the output argument of MPI_TYPE_SIZE in C is of type `int`. The MPI Forum considered proposals to change this and decided to reiterate the original decision. (*End of advice to users.*)

3.12.3 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 73. 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 Sec. 3.12.7. 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 “pseudo-datatypes,” MPI_LB and MPI_UB, that can be used, respectively, to mark the lower bound or the upper bound of a datatype. These pseudo-datatypes occupy no space ($extent(\text{MPI_LB}) = extent(\text{MPI_UB}) = 0$). 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 3.26 Let $D = (-3, 0, 6)$; $T = (\text{MPI_LB}, \text{MPI_INT}, \text{MPI_UB})$, and $B = (1, 1, 1)$. Then a call to `MPI_TYPE_STRUCT(3, B, D, T, 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 sequence $\{(\text{lb}, -3), (\text{int}, 0), (\text{ub}, 6)\}$. 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 sequence $\{(\text{lb}, -3), (\text{int}, 0), (\text{int}, 9), (\text{ub}, 15)\}$. (An entry of type `ub` can be deleted if there is another entry of type `ub` with a higher displacement; an entry of type `lb` can be deleted if there is another entry of type `lb` with a lower displacement.)

In general, if

$$Typemap = \{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\},$$

then the **lower bound** of *Typemap* is defined to be

$$lb(Typemap) = \begin{cases} \min_j disp_j & \text{if no entry has basic type lb} \\ \min_j \{disp_j \text{ such that } type_j = \text{lb}\} & \text{otherwise} \end{cases}$$

Similarly, the **upper bound** of *Typemap* is defined to be

$$ub(Typemap) = \begin{cases} \max_j disp_j + sizeof(type_j) + \epsilon & \text{if no entry has basic type ub} \\ \max_j \{disp_j \text{ such that } type_j = \text{ub}\} & \text{otherwise} \end{cases}$$

1 Then

$$2 \quad extent(Typemap) = ub(Typemap) - lb(Typemap)$$

4 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 $extent(Typemap)$ to the next multiple of $\max_i k_i$.

7 The formal definitions given for the various datatype constructors apply now, with the amended definition of **extent**.

9 The two functions below can be used for finding the lower bound and the upper bound of a datatype.

12 MPI_TYPE_LB(datatype, displacement)

13	IN	datatype	datatype (handle)
14			
15	OUT	displacement	displacement of lower bound from origin, in bytes (integer)
16			

17
18 int MPI_Type_lb(MPI_Datatype datatype, MPI_Aint* displacement)

19 MPI_TYPE_LB(DATATYPE, DISPLACEMENT, IERROR)
20 INTEGER DATATYPE, DISPLACEMENT, IERROR
21

23 MPI_TYPE_UB(datatype, displacement)

24			
25	IN	datatype	datatype (handle)
26			
27	OUT	displacement	displacement of upper bound from origin, in bytes (integer)
28			

29
30 int MPI_Type_ub(MPI_Datatype datatype, MPI_Aint* displacement)

31 MPI_TYPE_UB(DATATYPE, DISPLACEMENT, IERROR)
32 INTEGER DATATYPE, DISPLACEMENT, IERROR
33

35 3.12.4 Commit and free

36 A datatype object has to be **committed** before it can be used in a communication. A committed datatype can still be used as a argument in datatype constructors. There is no need to commit basic datatypes. They are “pre-committed.”

40 MPI_TYPE_COMMIT(datatype)

41			
42	INOUT	datatype	datatype that is committed (handle)
43			

44
45 int MPI_Type_commit(MPI_Datatype *datatype)

46 MPI_TYPE_COMMIT(DATATYPE, IERROR)
47 INTEGER DATATYPE, IERROR
48

The commit operation commits the datatype, that is, the formal description of a communication buffer, not the content of that buffer. Thus, after a datatype has been committed, it can be repeatedly reused to communicate the changing content of a buffer or, indeed, 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. (*End of advice to implementors.*)

`MPI_TYPE_FREE(datatype)`

INOUT datatype datatype that is freed (handle)

`int MPI_Type_free(MPI_Datatype *datatype)`

`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. Derived datatypes that were defined from the freed datatype are not affected.

Example 3.27 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
```

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 may implement constructors of derived datatypes so that they keep pointers to their datatype arguments, rather than copying them. In this case, one needs to keep track of active datatype definition references in order to know when a datatype object can be freed. (*End of advice to implementors.*)

3.12.5 Use of general datatypes in communication

Handles to derived datatypes can be passed to a communication call wherever a datatype argument is required. A call of the form `MPI_SEND(buf, count, datatype, ...)`, where `count > 1`, is interpreted as if the call was passed a new datatype which is the concatenation of `count` copies of `datatype`. Thus, `MPI_SEND(buf, count, datatype, dest, tag, comm)` is equivalent to,

```
MPI_TYPE_CONTIGUOUS(count, datatype, newtype)
MPI_TYPE_COMMIT(newtype)
MPI_SEND(buf, 1, newtype, dest, tag, comm).
```

Similar statements apply to all other communication functions that have a `count` and `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*. (Empty entries of “pseudo-type” `MPI_UB` and `MPI_LB` 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, empty entries of “pseudo-type” `MPI_UB` and `MPI_LB` 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 3.28 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, ...) 1
CALL MPI_SEND( a, 1, type4, ...) 2
... 3
CALL MPI_RECV( a, 4, MPI_REAL, ...) 4
CALL MPI_RECV( a, 2, type2, ...) 5
CALL MPI_RECV( a, 1, type22, ...) 6
CALL MPI_RECV( a, 1, type4, ...) 7

```

Each of the sends matches any of the receives. 8

A datatype may specify overlapping entries. The use of such a datatype in a receive operation is erroneous. (This is erroneous even if the actual message received is short enough not to write any entry more than once.) 9

Suppose that `MPI_RECV(buf, count, datatype, dest, tag, comm, status)` is executed, where `datatype` has type map, 10

$$\{(type_0, disp_0), \dots, (type_{n-1}, disp_{n-1})\}.$$

The received message need not fill all the receive buffer, nor does it need to fill a number of locations which is a multiple of n . Any number, k , of basic elements can be received, where $0 \leq k \leq count \cdot n$. The number of basic elements received can be retrieved from `status` using the query function `MPI_GET_ELEMENTS`. 11

`MPI_GET_ELEMENTS(status, datatype, count)` 12

IN	status	return status of receive operation (Status)	13
IN	datatype	datatype used by receive operation (handle)	14
OUT	count	number of received basic elements (integer)	15

```
int MPI_Get_elements(MPI_Status *status, MPI_Datatype datatype, int *count) 16
```

```
MPI_GET_ELEMENTS(STATUS, DATATYPE, COUNT, IERROR) 17
```

```
INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR 18
```

The previously defined function, `MPI_GET_COUNT` (Sec. 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 $0 \leq k \leq count$. 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` returns the value `MPI_UNDEFINED`. The `datatype` argument should match the argument provided by the receive call that set the `status` variable. 19

Example 3.29 Usage of `MPI_GET_COUNT` and `MPI_GET_ELEMENT`. 20

```

... 21
CALL MPI_TYPE_CONTIGUOUS(2, MPI_REAL, Type2, ierr) 22
CALL MPI_TYPE_COMMIT(Type2, ierr) 23
... 24

```

```

1 CALL MPI_COMM_RANK(comm, rank, ierr)
2 IF(rank.EQ.0) THEN
3     CALL MPI_SEND(a, 2, MPI_REAL, 1, 0, comm, ierr)
4     CALL MPI_SEND(a, 3, MPI_REAL, 1, 0, comm, ierr)
5 ELSE
6     CALL MPI_RECV(a, 2, Type2, 0, 0, comm, stat, ierr)
7     CALL MPI_GET_COUNT(stat, Type2, i, ierr)      ! returns i=1
8     CALL MPI_GET_ELEMENTS(stat, Type2, i, ierr)  ! returns i=2
9     CALL MPI_RECV(a, 2, Type2, 0, 0, comm, stat, ierr)
10    CALL MPI_GET_COUNT(stat, Type2, i, ierr)     ! returns i=MPI_UNDEFINED
11    CALL MPI_GET_ELEMENTS(stat, Type2, i, ierr)  ! returns i=3
12 END IF

```

The function `MPI_GET_ELEMENTS` can also be used after a probe to find the number of elements in the probed message. Note that the two functions `MPI_GET_COUNT` and `MPI_GET_ELEMENTS` return the same values when they are used with basic datatypes.

Rationale. The extension given to the definition of `MPI_GET_COUNT` seems natural: one would expect this function 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 function `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.*)

3.12.6 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.

Variables belong to the same **sequential storage** if they belong to the same array, to the same `COMMON` block in Fortran, or to the same structure in C. Valid addresses are defined recursively as follows:

1. The function `MPI_ADDRESS` returns a valid address, when passed as argument a variable of the calling program.

2. The `buf` argument of a communication function evaluates to a valid address, when passed as argument a variable of the calling program.
3. If `v` is a valid address, and `i` is an integer, then `v+i` is a valid address, provided `v` and `v+i` are in the same sequential storage.
4. If `v` is a valid address then `MPI_BOTTOM + v` is a valid address.

A correct program uses only valid addresses to identify the locations of entries in communication buffers. Furthermore, if `u` and `v` are two valid addresses, then the (integer) difference `u - v` can be computed only if both `u` and `v` are in the same sequential storage. No other arithmetic operations can be meaningfully executed on addresses.

The rules above impose no constraints on the use of derived datatypes, as long as they are used to define a communication buffer that is wholly contained within the same sequential storage. However, the construction of a communication buffer that contains variables that are not within the same sequential storage must obey certain restrictions. Basically, a communication buffer with variables that are not within the same sequential storage can be used only by specifying in the communication call `buf = MPI_BOTTOM`, `count = 1`, and using a `datatype` argument where all displacements are valid (absolute) addresses.

Advice to users. It is not expected that MPI implementations will be able to detect erroneous, “out of bound” displacements — unless those overflow the user address space — since the MPI call may not know the extent of the arrays and records in the host program. (*End of advice to users.*)

Advice to implementors. There is no need to distinguish (absolute) addresses and (relative) displacements on a machine with contiguous address space: `MPI_BOTTOM` is zero, and both addresses and displacements are integers. On machines where the distinction is required, addresses are recognized as expressions that involve `MPI_BOTTOM`. (*End of advice to implementors.*)

Note that in Fortran, Fortran INTEGERS may be too small to contain an address (e.g., 32 bit INTEGERS on a machine with 64bit pointers). Because of this, in Fortran, implementations may restrict the use of absolute addresses to only part of the process memory, and restrict the use of relative displacements to subranges of the process memory where they are constrained by the size of Fortran INTEGERS.

3.12.7 Examples

The following examples illustrate the use of derived datatypes.

Example 3.30 Send and receive a section of a 3D array.

```

REAL a(100,100,100), e(9,9,9)
INTEGER oneslice, twoslice, threeslice, sizeofreal, myrank, ierr
INTEGER status(MPI_STATUS_SIZE)

C      extract the section a(1:17:2, 3:11, 2:10)
C      and store it in e(:, :, :).

```

```

1
2     CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank)
3
4     CALL MPI_TYPE_EXTENT( MPI_REAL, sizeofreal, ierr)
5
6     C     create datatype for a 1D section
7     CALL MPI_TYPE_VECTOR( 9, 1, 2, MPI_REAL, oneslice, ierr)
8
9     C     create datatype for a 2D section
10    CALL MPI_TYPE_HVECTOR(9, 1, 100*sizeofreal, oneslice, twoslice, ierr)
11
12    C     create datatype for the entire section
13    CALL MPI_TYPE_HVECTOR( 9, 1, 100*100*sizeofreal, twoslice,
14                          threeslice, ierr)
15
16    CALL MPI_TYPE_COMMIT( threeslice, ierr)
17    CALL MPI_SENDRECV(a(1,3,2), 1, threeslice, myrank, 0, e, 9*9*9,
18                      MPI_REAL, myrank, 0, MPI_COMM_WORLD, status, ierr)
19

```

Example 3.31 Copy the (strictly) lower triangular part of a matrix.

```

21    REAL a(100,100), b(100,100)
22    INTEGER disp(100), blocklen(100), ltype, myrank, ierr
23    INTEGER status(MPI_STATUS_SIZE)
24
25    C     copy lower triangular part of array a
26    C     onto lower triangular part of array b
27
28
29    CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank)
30
31    C     compute start and size of each column
32    DO i=1, 100
33        disp(i) = 100*(i-1) + i
34        block(i) = 100-i
35    END DO
36
37    C     create datatype for lower triangular part
38    CALL MPI_TYPE_INDEXED( 100, block, disp, MPI_REAL, ltype, ierr)
39
40    CALL MPI_TYPE_COMMIT(ltype, ierr)
41    CALL MPI_SENDRECV( a, 1, ltype, myrank, 0, b, 1,
42                      ltype, myrank, 0, MPI_COMM_WORLD, status, ierr)
43

```

Example 3.32 Transpose a matrix.

```

45    REAL a(100,100), b(100,100)
46    INTEGER row, xpose, sizeofreal, myrank, ierr
47    INTEGER status(MPI_STATUS_SIZE)
48

```

```

C transpose matrix a onto b 1
2
CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank) 3
4
CALL MPI_TYPE_EXTENT( MPI_REAL, sizeofreal, ierr) 5
6
C create datatype for one row 7
CALL MPI_TYPE_VECTOR( 100, 1, 100, MPI_REAL, row, ierr) 8
9
C create datatype for matrix in row-major order 10
CALL MPI_TYPE_HVECTOR( 100, 1, sizeofreal, row, xpose, ierr) 11
12
CALL MPI_TYPE_COMMIT( xpose, ierr) 13
14
C send matrix in row-major order and receive in column major order 15
CALL MPI_SENDRECV( a, 1, xpose, myrank, 0, b, 100*100, 16
MPI_REAL, myrank, 0, MPI_COMM_WORLD, status, ierr) 17
18

```

Example 3.33 Another approach to the transpose problem: 19

```

REAL a(100,100), b(100,100) 20
INTEGER disp(2), blocklen(2), type(2), row, row1, sizeofreal 21
INTEGER myrank, ierr 22
INTEGER status(MPI_STATUS_SIZE) 23
24
CALL MPI_COMM_RANK(MPI_COMM_WORLD, myrank) 25
26
C transpose matrix a onto b 27
28
CALL MPI_TYPE_EXTENT( MPI_REAL, sizeofreal, ierr) 29
30
C create datatype for one row 31
CALL MPI_TYPE_VECTOR( 100, 1, 100, MPI_REAL, row, ierr) 32
33
C create datatype for one row, with the extent of one real number 34
disp(1) = 0 35
disp(2) = sizeofreal 36
type(1) = row 37
type(2) = MPI_UB 38
blocklen(1) = 1 39
blocklen(2) = 1 40
CALL MPI_TYPE_STRUCT( 2, blocklen, disp, type, row1, ierr) 41
42
CALL MPI_TYPE_COMMIT( row1, ierr) 43
44
C send 100 rows and receive in column major order 45
CALL MPI_SENDRECV( a, 100, row1, myrank, 0, b, 100*100, 46
MPI_REAL, myrank, 0, MPI_COMM_WORLD, status, ierr) 47
48

```



```

1  Example 3.34 We manipulate an array of structures.
2
3  struct Partstruct
4  {
5      int    class; /* particle class */
6      double d[6]; /* particle coordinates */
7      char   b[7]; /* some additional information */
8  };
9
10 struct Partstruct    particle[1000];
11
12 int                  i, dest, rank;
13 MPI_Comm             comm;
14
15
16 /* build datatype describing structure */
17
18 MPI_Datatype Particletype;
19 MPI_Datatype type[3] = {MPI_INT, MPI_DOUBLE, MPI_CHAR};
20 int           blocklen[3] = {1, 6, 7};
21 MPI_Aint     disp[3];
22 MPI_Aint     base;
23
24
25 /* compute displacements of structure components */
26
27 MPI_Address( particle, disp);
28 MPI_Address( particle[0].d, disp+1);
29 MPI_Address( particle[0].b, disp+2);
30 base = disp[0];
31 for (i=0; i <3; i++) disp[i] -= base;
32
33 MPI_Type_struct( 3, blocklen, disp, type, &Particletype);
34
35     /* If compiler does padding in mysterious ways,
36     the following may be safer */
37
38 MPI_Datatype type1[4] = {MPI_INT, MPI_DOUBLE, MPI_CHAR, MPI_UB};
39 int           blocklen1[4] = {1, 6, 7, 1};
40 MPI_Aint     disp1[4];
41
42 /* compute displacements of structure components */
43
44 MPI_Address( particle, disp1);
45 MPI_Address( particle[0].d, disp1+1);
46 MPI_Address( particle[0].b, disp1+2);
47 MPI_Address( particle+1, disp1+3);
48 base = disp1[0];

```

```

for (i=0; i <4; i++) disp1[i] -= base;
1
2
/* build datatype describing structure */
3
4
MPI_Type_struct( 4, blocklen1, disp1, type1, &Particletype);
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
/* 4.1:
send the entire array */

MPI_Type_commit( &Particletype);
MPI_Send( particle, 1000, Particletype, dest, tag, comm);

/* 4.2:
send only the entries of class zero particles,
preceded by the number of such entries */

MPI_Datatype Zparticles; /* datatype describing all particles
with class zero (needs to be recomputed
if classes change) */

MPI_Datatype Ztype;

MPI_Aint zdisp[1000];
int zblock[1000], j, k;
int zzblock[2] = {1,1};
MPI_Aint zzdisp[2];
MPI_Datatype zztype[2];

/* compute displacements of class zero particles */
j = 0;
for(i=0; i < 1000; i++)
    if (particle[i].class==0)
        {
            zdisp[j] = i;
            zblock[j] = 1;
            j++;
        }

/* create datatype for class zero particles */
MPI_Type_indexed( j, zblock, zdisp, Particletype, &Zparticles);

/* prepend particle count */
MPI_Address(&j, zzdisp);
MPI_Address(particle, zzdisp+1);
zztype[0] = MPI_INT;
zztype[1] = Zparticles;
MPI_Type_struct(2, zzblock, zzdisp, zztype, &Ztype);

```

```

1
2 MPI_Type_commit( &Ztype);
3 MPI_Send( MPI_BOTTOM, 1, Ztype, dest, tag, comm);
4
5
6     /* A probably more efficient way of defining Zparticles */
7
8 /* consecutive particles with index zero are handled as one block */
9 j=0;
10 for (i=0; i < 1000; i++)
11     if (particle[i].index==0)
12         {
13             for (k=i+1; (k < 1000)&&(particle[k].index == 0) ; k++);
14             zdisp[j] = i;
15             zblock[j] = k-i;
16             j++;
17             i = k;
18         }
19 MPI_Type_indexed( j, zblock, zdisp, Particletype, &Zparticles);
20
21
22     /* 4.3:
23     send the first two coordinates of all entries */
24
25 MPI_Datatype Allpairs;     /* datatype for all pairs of coordinates */
26
27 MPI_Aint sizeofentry;
28
29 MPI_Type_extent( Particletype, &sizeofentry);
30
31     /* sizeofentry can also be computed by subtracting the address
32     of particle[0] from the address of particle[1] */
33
34 MPI_Type_hvector( 1000, 2, sizeofentry, MPI_DOUBLE, &Allpairs);
35 MPI_Type_commit( &Allpairs);
36 MPI_Send( particle[0].d, 1, Allpairs, dest, tag, comm);
37
38     /* an alternative solution to 4.3 */
39
40 MPI_Datatype Onepair;     /* datatype for one pair of coordinates, with
41                             the extent of one particle entry */
42 MPI_Aint disp2[3];
43 MPI_Datatype type2[3] = {MPI_LB, MPI_DOUBLE, MPI_UB};
44 int blocklen2[3] = {1, 2, 1};
45
46 MPI_Address( particle, disp2);
47 MPI_Address( particle[0].d, disp2+1);
48 MPI_Address( particle+1, disp2+2);

```

```

base = disp2[0];
for (i=0; i<2; i++) disp2[i] -= base;

MPI_Type_struct( 3, blocklen2, disp2, type2, &Onepair);
MPI_Type_commit( &Onepair);
MPI_Send( particle[0].d, 1000, Onepair, dest, tag, comm);

```

Example 3.35 The same manipulations as in the previous example, but use absolute addresses in datatypes.

```

struct Partstruct
{
    int class;
    double d[6];
    char b[7];
};

struct Partstruct particle[1000];

    /* build datatype describing first array entry */

MPI_Datatype Particletype;
MPI_Datatype type[3] = {MPI_INT, MPI_DOUBLE, MPI_CHAR};
int          block[3] = {1, 6, 7};
MPI_Aint     disp[3];

MPI_Address( particle, disp);
MPI_Address( particle[0].d, disp+1);
MPI_Address( particle[0].b, disp+2);
MPI_Type_struct( 3, block, disp, type, &Particletype);

/* Particletype describes first array entry -- using absolute
   addresses */

    /* 5.1:
       send the entire array */

MPI_Type_commit( &Particletype);
MPI_Send( MPI_BOTTOM, 1000, Particletype, dest, tag, comm);

    /* 5.2:
       send the entries of class zero,
       preceded by the number of such entries */

MPI_Datatype Zparticles, Ztype;

```

```

1  MPI_Aint zdisp[1000]
2  int zblock[1000], i, j, k;
3  int zzblock[2] = {1,1};
4  MPI_Datatype zztype[2];
5  MPI_Aint    zzdisp[2];
6
7  j=0;
8  for (i=0; i < 1000; i++)
9      if (particle[i].index==0)
10         {
11             for (k=i+1; (k < 1000)&&(particle[k].index = 0) ; k++);
12             zdisp[j] = i;
13             zblock[j] = k-i;
14             j++;
15             i = k;
16         }
17 MPI_Type_indexed( j, zblock, zdisp, Particletype, &Zparticles);
18 /* Zparticles describe particles with class zero, using
19    their absolute addresses*/
20
21 /* prepend particle count */
22 MPI_Address(&j, zzdisp);
23 zzdisp[1] = MPI_BOTTOM;
24 zztype[0] = MPI_INT;
25 zztype[1] = Zparticles;
26 MPI_Type_struct(2, zzblock, zzdisp, zztype, &Ztype);
27
28 MPI_Type_commit( &Ztype);
29 MPI_Send( MPI_BOTTOM, 1, Ztype, dest, tag, comm);
30
31
32 Example 3.36 Handling of unions.
33
34 union {
35     int    ival;
36     float  fval;
37     } u[1000]
38
39 int    utype;
40
41 /* All entries of u have identical type; variable
42    utype keeps track of their current type */
43
44 MPI_Datatype  type[2];
45 int          blocklen[2] = {1,1};
46 MPI_Aint     disp[2];
47 MPI_Datatype mpi_utype[2];
48 MPI_Aint     i,j;

```

```

1
2 /* compute an MPI datatype for each possible union type;
3    assume values are left-aligned in union storage. */
4
5 MPI_Address( u, &i);
6 MPI_Address( u+1, &j);
7 disp[0] = 0; disp[1] = j-i;
8 type[1] = MPI_UB;
9
10 type[0] = MPI_INT;
11 MPI_Type_struct(2, blocklen, disp, type, &mpi_utype[0]);
12
13 type[0] = MPI_FLOAT;
14 MPI_Type_struct(2, blocklen, disp, type, &mpi_utype[1]);
15
16 for(i=0; i<2; i++) MPI_Type_commit(&mpi_utype[i]);
17
18 /* actual communication */
19
20 MPI_Send(u, 1000, mpi_utype[utype], dest, tag, comm);
21

```

3.13 Pack and unpack

22
23
24 Some existing communication libraries provide pack/unpack functions for sending noncon-
25 tiguous data. In these, the user explicitly packs data into a contiguous buffer before sending
26 it, and unpacks it from a contiguous buffer after receiving it. Derived datatypes, which are
27 described in Section 3.12, allow one, in most cases, to avoid explicit packing and unpacking.
28 The user specifies the layout of the data to be sent or received, and the communication
29 library directly accesses a noncontiguous buffer. The pack/unpack routines are provided
30 for compatibility with previous libraries. Also, they provide some functionality that is not
31 otherwise available in MPI. For instance, a message can be received in several parts, where
32 the receive operation done on a later part may depend on the content of a former part.
33 Another use is that outgoing messages may be explicitly buffered in user supplied space,
34 thus overriding the system buffering policy. Finally, the availability of pack and unpack
35 operations facilitates the development of additional communication libraries layered on top
36 of MPI.

```
1 MPI_PACK(inbuf, incount, datatype, outbuf, outsize, position, comm)
```

```
2     IN      inbuf          input buffer start (choice)
3
4     IN      incount       number of input data items (integer)
5
6     IN      datatype      datatype of each input data item (handle)
7
8     OUT     outbuf        output buffer start (choice)
9
10    IN      outsize       output buffer size, in bytes (integer)
11
12    INOUT   position      current position in buffer, in bytes (integer)
13
14    IN      comm          communicator for packed message (handle)
```

```
15 int MPI_Pack(void* inbuf, int incount, MPI_Datatype datatype, void *outbuf,
16             int outsize, int *position, MPI_Comm comm)
```

```
17 MPI_PACK(INBUF, INCOUNT, DATATYPE, OUTBUF, OUTSIZE, POSITION, COMM, IERROR)
18     <type> INBUF(*), OUTBUF(*)
19     INTEGER INCOUNT, DATATYPE, OUTSIZE, POSITION, COMM, IERROR
```

19 Packs the message in the send buffer specified by `inbuf`, `incount`, `datatype` into the buffer space specified by `outbuf` and `outsize`. The input buffer can be any communication buffer allowed in `MPI_SEND`. The output buffer is a contiguous storage area containing `outsize` 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`).

24 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 packed message. The `comm` argument is the communicator that will be subsequently used for sending the packed message.

```
30 MPI_UNPACK(inbuf, insize, position, outbuf, outcount, datatype, comm)
```

```
31     IN      inbuf          input buffer start (choice)
32
33     IN      insize        size of input buffer, in bytes (integer)
34
35     INOUT   position      current position in bytes (integer)
36
37     OUT     outbuf        output buffer start (choice)
38
39     IN      outcount      number of items to be unpacked (integer)
40
41     IN      datatype      datatype of each output data item (handle)
42
43     IN      comm          communicator for packed message (handle)
```

```
44 int MPI_Unpack(void* inbuf, int insize, int *position, void *outbuf,
45              int outcount, MPI_Datatype datatype, MPI_Comm comm)
```

```
46 MPI_UNPACK(INBUF, INSIZE, POSITION, OUTBUF, OUTCOUNT, DATATYPE, COMM,
47            IERROR)
48     <type> INBUF(*), OUTBUF(*)
49     INTEGER INSIZE, POSITION, OUTCOUNT, DATATYPE, COMM, IERROR
```

Unpacks a message into the receive buffer specified by `outbuf`, `outcount`, `datatype` from the buffer space specified by `inbuf` and `insize`. The output buffer can be any communication 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 `scanf` 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 function 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 (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 (integer)

```
int MPI_Pack_size(int incount, MPI_Datatype datatype, MPI_Comm comm,
                 int *size)
```

```
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).

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 3.37 An example using MPI_PACK.

```
int position, i, j, a[2];
char buff[1000];
....

MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
if (myrank == 0)
{
    /* SENDER CODE */

    position = 0;
    MPI_Pack(&i, 1, MPI_INT, buff, 1000, &position, MPI_COMM_WORLD);
    MPI_Pack(&j, 1, MPI_INT, buff, 1000, &position, MPI_COMM_WORLD);
    MPI_Send( buff, position, MPI_PACKED, 1, 0, MPI_COMM_WORLD);
```

```

}
else /* RECEIVER CODE */
    MPI_Recv( a, 2, MPI_INT, 0, 0, MPI_COMM_WORLD)
}

```

Example 3.38 A elaborate example.

```

int position, i;
float a[1000];
char buff[1000]

....

MPI_Comm_rank(MPI_Comm_world, &myrank);
if (myrank == 0)
{
    /* SENDER CODE */

    int len[2];
    MPI_Aint disp[2];
    MPI_Datatype type[2], newtype;

    /* build datatype for i followed by a[0]...a[i-1] */

    len[0] = 1;
    len[1] = i;
    MPI_Address( &i, disp);
    MPI_Address( a, disp+1);
    type[0] = MPI_INT;
    type[1] = MPI_FLOAT;
    MPI_Type_struct( 2, len, disp, type, &newtype);
    MPI_Type_commit( &newtype);

    /* Pack i followed by a[0]...a[i-1]*/

    position = 0;
    MPI_Pack( MPI_BOTTOM, 1, newtype, buff, 1000, &position, MPI_COMM_WORLD);

    /* Send */

    MPI_Send( buff, position, MPI_PACKED, 1, 0,
              MPI_COMM_WORLD)

/* *****
   One can replace the last three lines with
   MPI_Send( MPI_BOTTOM, 1, newtype, 1, 0, MPI_COMM_WORLD);
   ***** */

```

```

1  }
2  else /* myrank == 1 */
3  {
4      /* RECEIVER CODE */
5
6      MPI_Status status;
7
8      /* Receive */
9
10     MPI_Recv( buff, 1000, MPI_PACKED, 0, 0, &status);
11
12     /* Unpack i */
13
14     position = 0;
15     MPI_Unpack(buff, 1000, &position, &i, 1, MPI_INT, MPI_COMM_WORLD);
16
17     /* Unpack a[0]...a[i-1] */
18     MPI_Unpack(buff, 1000, &position, a, i, MPI_FLOAT, MPI_COMM_WORLD);
19 }
20

```

Example 3.39 Each process sends a count, followed by count characters to the root; the root concatenate all characters into one string.

```

23 int count, gsize, counts[64], totalcount, k1, k2, k,
24     displs[64], position, concat_pos;
25 char chr[100], *lbuf, *rbuf, *cbuf;
26
27 ...
28 MPI_Comm_size(comm, &gsize);
29 MPI_Comm_rank(comm, &myrank);
30
31     /* allocate local pack buffer */
32     MPI_Pack_size(1, MPI_INT, comm, &k1);
33     MPI_Pack_size(count, MPI_CHAR, comm, &k2);
34     k = k1+k2;
35     lbuf = (char *)malloc(k);
36
37     /* pack count, followed by count characters */
38     position = 0;
39     MPI_Pack(&count, 1, MPI_INT, lbuf, k, &position, comm);
40     MPI_Pack(chr, count, MPI_CHAR, lbuf, k, &position, comm);
41
42     if (myrank != root)
43         /* gather at root sizes of all packed messages */
44         MPI_Gather( &position, 1, MPI_INT, NULL, NULL,
45                   NULL, root, comm);
46
47         /* gather at root packed messages */
48     MPI_Gatherv( &lbuf, position, MPI_PACKED, NULL,

```

```
        NULL, NULL, NULL, root, comm); 1
2
else { /* root code */ 3
    /* gather sizes of all packed messages */ 4
    MPI_Gather( &position, 1, MPI_INT, counts, 1, 5
               MPI_INT, root, comm); 6
7
    /* gather all packed messages */ 8
    displs[0] = 0; 9
    for (i=1; i < gsize; i++) 10
        displs[i] = displs[i-1] + counts[i-1]; 11
    totalcount = displs[gsizes-1] + counts[gsizes-1]; 12
    rbuf = (char *)malloc(totalcount); 13
    cbuf = (char *)malloc(totalcount); 14
    MPI_Gatherv( lbuf, position, MPI_PACKED, rbuf, 15
                counts, displs, MPI_PACKED, root, comm); 16
17
    /* unpack all messages and concatenate strings */ 18
    concat_pos = 0; 19
    for (i=0; i < gsize; i++) { 20
        position = 0; 21
        MPI_Unpack( rbuf+displs[i], totalcount-displs[i], 22
                   &position, &count, 1, MPI_INT, comm); 23
        MPI_Unpack( rbuf+displs[i], totalcount-displs[i], 24
                   &position, cbuf+concat_pos, count, MPI_CHAR, comm); 25
        concat_pos += count; 26
    } 27
    cbuf[concat_pos] = '\0'; 28
} 29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
```

Chapter 4

Collective Communication

4.1 Introduction and Overview

Collective communication is defined as communication that involves a group of processes. The functions of this type provided by MPI are the following:

- Barrier synchronization across all group members (Sec. 4.3).
- Broadcast from one member to all members of a group (Sec. 4.4). This is shown in figure 4.1.
- Gather data from all group members to one member (Sec. 4.5). This is shown in figure 4.1.
- Scatter data from one member to all members of a group (Sec. 4.6). This is shown in figure 4.1.
- A variation on Gather where all members of the group receive the result (Sec. 4.7). This is shown as “allgather” in figure 4.1.
- Scatter/Gather data from all members to all members of a group (also called complete exchange or all-to-all) (Sec. 4.8). This is shown as “alltoall” in figure 4.1.
- Global reduction operations such as sum, max, min, or user-defined functions, where the result is returned to all group members and a variation where the result is returned to only one member (Sec. 4.9).
- A combined reduction and scatter operation (Sec. 4.10).
- Scan across all members of a group (also called prefix) (Sec. 4.11).

A collective operation is executed by having all processes in the group call the communication routine, with matching arguments. The syntax and semantics of the collective operations are defined to be consistent with the syntax and semantics of the point-to-point operations. Thus, general datatypes are allowed and must match between sending and receiving processes as specified in Chapter 3. One of the key arguments is a communicator that defines the group of participating processes and provides a context for the operation. Several collective routines such as broadcast and gather have a single originating or receiving process. Such processes are called the *root*. Some arguments in the collective functions




Figure 4.1: Collective move functions illustrated for a group of six processes. In each case, each row of boxes represents data locations in one process. Thus, in the broadcast, initially just the first process contains the data A_0 , but after the broadcast all processes contain it.

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 are specified as “significant only at root,” and are ignored for all participants except the
2 root. The reader is referred to Chapter 3 for information concerning communication buffers,
3 general datatypes and type matching rules, and to Chapter 5 for information on how to
4 define groups and create communicators.

5 The type-matching conditions for the collective operations are more strict than the cor-
6 responding conditions between sender and receiver in point-to-point. Namely, for collective
7 operations, the amount of data sent must exactly match the amount of data specified by
8 the receiver. Distinct type maps (the layout in memory, see Sec. 3.12) between sender and
9 receiver are still allowed.

10 Collective routine calls can (but are not required to) return as soon as their participa-
11 tion in the collective communication is complete. The completion of a call indicates that the
12 caller is now free to access locations in the communication buffer. It does not indicate that
13 other processes in the group have completed or even started the operation (unless otherwise
14 indicated in the description of the operation). Thus, a collective communication call may,
15 or may not, have the effect of synchronizing all calling processes. This statement excludes,
16 of course, the barrier function.

17 Collective communication calls may use the same communicators as point-to-point
18 communication; MPI guarantees that messages generated on behalf of collective communi-
19 cation calls will not be confused with messages generated by point-to-point communication.
20 A more detailed discussion of correct use of collective routines is found in Sec. 4.12.

21
22 *Rationale.* The equal-data restriction (on type matching) was made so as to avoid
23 the complexity of providing a facility analogous to the status argument of MPI_RECV
24 for discovering the amount of data sent. Some of the collective routines would require
25 an array of status values.

26 The statements about synchronization are made so as to allow a variety of implemen-
27 tations of the collective functions.

28 The collective operations do not accept a message tag argument. If future revisions of
29 MPI define non-blocking collective functions, then tags (or a similar mechanism) will
30 need to be added so as to allow the dis-ambiguation of multiple, pending, collective
31 operations. (*End of rationale.*)

32
33 *Advice to users.* It is dangerous to rely on synchronization side-effects of the col-
34 lective operations for program correctness. For example, even though a particular
35 implementation may provide a broadcast routine with a side-effect of synchroniza-
36 tion, the standard does not require this, and a program that relies on this will not be
37 portable.

38 On the other hand, a correct, portable program must allow for the fact that a collective
39 call *may* be synchronizing. Though one cannot rely on any synchronization side-effect,
40 one must program so as to allow it. These issues are discussed further in Sec. 4.12.
41 (*End of advice to users.*)

42
43 *Advice to implementors.* While vendors may write optimized collective routines
44 matched to their architectures, a complete library of the collective communication
45 routines can be written entirely using the MPI point-to-point communication func-
46 tions and a few auxiliary functions. If implementing on top of point-to-point, a hidden,
47

special communicator must be created for the collective operation so as to avoid interference with any on-going point-to-point communication at the time of the collective call. This is discussed further in Sec. 4.12. (*End of advice to implementors.*)

4.2 Communicator argument

The key concept of the collective functions is to have a “group” of participating processes. The routines do not have a group identifier as an explicit argument. Instead, there is a communicator argument. For the purposes of this chapter, a communicator can be thought of as a group identifier linked with a context. An inter-communicator, that is, a communicator that spans two groups, is *not* allowed as an argument to a collective function.

4.3 Barrier synchronization

`MPI_BARRIER(comm)`

IN `comm` communicator (handle)

`int MPI_Barrier(MPI_Comm comm)`

`MPI_BARRIER(COMM, IERROR)`
 INTEGER `COMM, IERROR`

`MPI_BARRIER` blocks the caller until all group members have called it. The call returns at any process only after all group members have entered the call.

4.4 Broadcast

`MPI_BCAST(buffer, count, datatype, root, comm)`

INOUT `buffer` starting address of buffer (choice)

IN `count` number of entries in buffer (integer)

IN `datatype` data type of buffer (handle)

IN `root` rank of broadcast root (integer)

IN `comm` communicator (handle)

`int MPI_Bcast(void* buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm)`

`MPI_BCAST(BUFFER, COUNT, DATATYPE, ROOT, COMM, IERROR)`
 <type> `BUFFER(*)`
 INTEGER `COUNT, DATATYPE, ROOT, COMM, IERROR`

`MPI_BCAST` broadcasts a message from the process with rank `root` to all processes of the group, itself included. It is called by all members of group using the same arguments

1 for `comm`, `root`. On return, the contents of `root`'s communication buffer has been copied to
 2 all processes.

3 General, derived datatypes are allowed for `datatype`. The type signature of `count`,
 4 `datatype` on any process must be equal to the type signature of `count`, `datatype` at the root.
 5 This implies that the amount of data sent must be equal to the amount received, pairwise
 6 between each process and the root. `MPI_BCAST` and all other data-movement collective
 7 routines make this restriction. Distinct type maps between sender and receiver are still
 8 allowed.

10 4.4.1 Example using `MPI_BCAST`

11 **Example 4.1** Broadcast 100 ints from process 0 to every process in the group.

```
13     MPI_Comm comm;
14     int array[100];
15     int root=0;
16     ...
17     MPI_Bcast( array, 100, MPI_INT, root, comm);
```

19 As in many of our example code fragments, we assume that some of the variables (such as
 20 `comm` in the above) have been assigned appropriate values.

23 4.5 Gather

26 `MPI_GATHER(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root, comm)`

28	IN	<code>sendbuf</code>	starting address of send buffer (choice)
29	IN	<code>sendcount</code>	number of elements in send buffer (integer)
30	IN	<code>sendtype</code>	data type of send buffer elements (handle)
31	IN	<code>sendtype</code>	data type of send buffer elements (handle)
32	OUT	<code>recvbuf</code>	address of receive buffer (choice, significant only at root)
33			root)
34	IN	<code>recvcount</code>	number of elements for any single receive (integer, significant only at root)
35			nificant only at root)
36	IN	<code>recvtype</code>	data type of recv buffer elements (significant only at root) (handle)
37			root) (handle)
38			
39	IN	<code>root</code>	rank of receiving process (integer)
40	IN	<code>comm</code>	communicator (handle)
41			

```
42 int MPI_Gather(void* sendbuf, int sendcount, MPI_Datatype sendtype,
43              void* recvbuf, int recvcount, MPI_Datatype recvtype, int root,
44              MPI_Comm comm)
```

```
46 MPI_GATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
47           ROOT, COMM, IERROR)
48 <type> SENDBUF(*), RECVBUF(*)
```

INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVMODE, ROOT, COMM, IERROR

Each process (root process included) sends the contents of its send buffer to the root process. The root process receives the messages and stores them in rank order. The outcome is *as if* each of the n processes in the group (including the root process) 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_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 non-root processes.

General, derived datatypes are allowed for both `sendtype` and `recvtype`. The type signature of `sendcount`, `sendtype` on process `i` 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 process and the root. Distinct type maps between sender and receiver are still allowed.

All arguments to the function are significant on process `root`, while on other processes, only arguments `sendbuf`, `sendcount`, `sendtype`, `root`, `comm` are significant. The arguments `root` and `comm` must have identical values on all 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* process, not the total number of items it receives.

```

1 MPI_GATHERV( sendbuf, sendcount, sendtype, recvbuf, recvcunts, displs, recvtype, root,
2 comm)
3     IN      sendbuf      starting address of send buffer (choice)
4     IN      sendcount    number of elements in send buffer (integer)
5     IN      sendtype     data type of send buffer elements (handle)
6     OUT     recvbuf      address of receive buffer (choice, significant only at
7                          root)
8     IN      recvcunts    integer array (of length group size) containing the num-
9                          ber of elements that are received from each process
10                        (significant only at root)
11
12     IN      displs      integer array (of length group size). Entry i specifies
13                        the displacement relative to recvbuf at which to place
14                        the incoming data from process i (significant only at
15                        root)
16
17     IN      recvtype     data type of recv buffer elements (significant only at
18                          root) (handle)
19
20     IN      root        rank of receiving process (integer)
21
22     IN      comm        communicator (handle)

```

```

23 int MPI_Gatherv(void* sendbuf, int sendcount, MPI_Datatype sendtype,
24                void* recvbuf, int *recvcunts, int *displs,
25                MPI_Datatype recvtype, int root, MPI_Comm comm)
26
27 MPI_GATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNTS, DISPLS,
28             RECVTYPE, ROOT, COMM, IERROR)
29 <type> SENDBUF(*), RECVBUF(*)
30 INTEGER SENDCOUNT, SENDTYPE, REVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT,
31 COMM, IERROR

```

32 `MPI_GATHERV` extends the functionality of `MPI_GATHER` by allowing a varying count
33 of data from each process, since `recvcunts` is now an array. It also allows more flexibility
34 as to where the data is placed on the root, by providing the new argument, `displs`.

35 The outcome is *as if* each process, including the root process, sends a message to the
36 root,

```

37
38     MPI_Send(sendbuf, sendcount, sendtype, root, ...),

```

39 and the root executes `n` receives,

```

40
41     MPI_Recv(recvbuf + displs[i] · extent(recvtype), recvcunts[i], recvtype, i, ...).

```

42
43 Messages are placed in the receive buffer of the root process in rank order, that is, the
44 data sent from process `j` is placed in the `j`th portion of the receive buffer `recvbuf` on process
45 `root`. The `j`th portion of `recvbuf` begins at offset `displs[j]` elements (in terms of `recvtype`) into
46 `recvbuf`.

47 The receive buffer is ignored for all non-root processes.

48




Figure 4.2: The root process gathers 100 ints from each process in the group.

The type signature implied by `sendcount`, `sendtype` on process `i` must be equal to the type signature implied by `recvcounts[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 process and the root. Distinct type maps between sender and receiver are still allowed, as illustrated in Example 4.6.

All arguments to the function are significant on process `root`, while on other processes, only arguments `sendbuf`, `sendcount`, `sendtype`, `root`, `comm` are significant. The arguments `root` and `comm` must have identical values on all processes.

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.

4.5.1 Examples using `MPI_GATHER`, `MPI_GATHERV`

Example 4.2 Gather 100 ints from every process in group to root. See figure 4.2.

```

MPI_Comm comm;
int gsize, sendarray[100];
int root, *rbuf;
...
MPI_Comm_size( comm, &gsize);
rbuf = (int *)malloc(gsize*100*sizeof(int));
MPI_Gather( sendarray, 100, MPI_INT, rbuf, 100, MPI_INT, root, comm);

```

Example 4.3 Previous example modified – only the root allocates memory for the receive buffer.

```

MPI_Comm comm;
int gsize, sendarray[100];
int root, myrank, *rbuf;
...
MPI_Comm_rank( comm, myrank);
if ( myrank == root) {
    MPI_Comm_size( comm, &gsize);
    rbuf = (int *)malloc(gsize*100*sizeof(int));
}
MPI_Gather( sendarray, 100, MPI_INT, rbuf, 100, MPI_INT, root, comm);

```

Example 4.4 Do the same as the previous example, but use a derived datatype. Note that the type cannot be the entire set of `gsize*100` ints since type matching is defined pairwise between the root and each process in the gather.

```

5     MPI_Comm comm;
6     int gsize,sendarray[100];
7     int root, *rbuf;
8     MPI_Datatype rtype;
9     ...
10    MPI_Comm_size( comm, &gsize);
11    MPI_Type_contiguous( 100, MPI_INT, &rtype );
12    MPI_Type_commit( &rtype );
13    rbuf = (int *)malloc(gsize*100*sizeof(int));
14    MPI_Gather( sendarray, 100, MPI_INT, rbuf, 1, rtype, root, comm);

```

Example 4.5 Now have each process send 100 ints to root, but place each set (of 100) *stride* ints apart at receiving end. Use `MPI_GATHERV` and the `displs` argument to achieve this effect. Assume *stride* ≥ 100 . See figure 4.3.

```

19    MPI_Comm comm;
20    int gsize,sendarray[100];
21    int root, *rbuf, stride;
22    int *displs,i,*rcounts;
23    ...
24
25    ...
26
27    MPI_Comm_size( comm, &gsize);
28    rbuf = (int *)malloc(gsize*stride*sizeof(int));
29    displs = (int *)malloc(gsize*sizeof(int));
30    rcounts = (int *)malloc(gsize*sizeof(int));
31    for (i=0; i<gsize; ++i) {
32        displs[i] = i*stride;
33        rcounts[i] = 100;
34    }
35    MPI_Gatherv( sendarray, 100, MPI_INT, rbuf, rcounts, displs, MPI_INT,
36                root, comm);

```

Note that the program is erroneous if *stride* < 100 .

Example 4.6 Same as Example 4.5 on the receiving side, but send the 100 ints from the 0th column of a 100×150 int array, in C. See figure 4.4.

```

43    MPI_Comm comm;
44    int gsize,sendarray[100][150];
45    int root, *rbuf, stride;
46    MPI_Datatype stype;
47    int *displs,i,*rcounts;
48

```




Figure 4.3: The root process gathers 100 ints from each process in the group, each set is placed `stride` ints apart.




Figure 4.4: The root process gathers column 0 of a 100×150 C array, and each set is placed `stride` ints apart.

...

```

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;
}
/* Create datatype for 1 column of array
*/
MPI_Type_vector( 100, 1, 150, MPI_INT, &styp);
MPI_Type_commit( &styp );
MPI_Gatherv( sendarray, 1, styp, rbuf, rcounts, displs, MPI_INT,
             root, comm);

```

Example 4.7 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 4.5.

```

MPI_Comm comm;
int gsize, sendarray[100][150], *sptr;
int root, *rbuf, stride, myrank;

```




Figure 4.5: The root process gathers $100-i$ ints from column i of a 100×150 C array, and each set is placed stride ints apart.

```

14 MPI_Datatype stype;
15 int *displs,i,*rcounts;
16
17 ...
18
19 MPI_Comm_size( comm, &gsize);
20 MPI_Comm_rank( comm, &myrank );
21 rbuf = (int *)malloc(gsize*stride*sizeof(int));
22 displs = (int *)malloc(gsize*sizeof(int));
23 rcounts = (int *)malloc(gsize*sizeof(int));
24 for (i=0; i<gsize; ++i) {
25     displs[i] = i*stride;
26     rcounts[i] = 100-i;    /* note change from previous example */
27 }
28 /* Create datatype for the column we are sending
29 */
30 MPI_Type_vector( 100-myrank, 1, 150, MPI_INT, &stype);
31 MPI_Type_commit( &stype );
32 /* sptr is the address of start of "myrank" column
33 */
34 sptr = &sendarray[0][myrank];
35 MPI_Gatherv( sptr, 1, stype, rbuf, rcounts, displs, MPI_INT,
36             root, comm);

```

Note that a different amount of data is received from each process.

Example 4.8 Same as Example 4.7, 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 that we read a column of a C array. A similar thing was done in Example 3.33, Section 3.12.7.

```

44 MPI_Comm comm;
45 int gsize,sendarray[100][150],*sptr;
46 int root,*rbuf, stride, myrank, disp[2], blocklen[2];
47 MPI_Datatype stype,type[2];
48 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;
}
/* Create datatype for one int, with extent of entire row
*/
disp[0] = 0;      disp[1] = 150*sizeof(int);
type[0] = MPI_INT; type[1] = MPI_UB;
blocklen[0] = 1;  blocklen[1] = 1;
MPI_Type_struct( 2, blocklen, disp, type, &stype );
MPI_Type_commit( &stype );
sptr = &sendarray[0][myrank];
MPI_Gatherv( sptr, 100-myrank, stype, rbuf, rcounts, displs, MPI_INT,
            root, comm);

```

Example 4.9 Same as Example 4.7 at sending side, but at receiving side we make the stride between received blocks vary from block to block. See figure 4.6.

```

MPI_Comm comm;
int gsize,sendarray[100][150],*sptr;
int root, *rbuf, *stride, myrank, bufsize;
MPI_Datatype stype;
int *displs,i,*rcounts,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
*/

/* set up displs and rcounts vectors first
*/
displs = (int *)malloc(gsize*sizeof(int));
rcounts = (int *)malloc(gsize*sizeof(int));
offset = 0;
for (i=0; i<gsize; ++i) {

```





Figure 4.6: The root process 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).

```

14     displs[i] = offset;
15     offset += stride[i];
16     rcounts[i] = 100-i;
17 }
18 /* the required buffer size for rbuf is now easily obtained
19 */
20 bufsize = displs[gsize-1]+rcounts[gsize-1];
21 rbuf = (int *)malloc(bufsize*sizeof(int));
22 /* Create datatype for the column we are sending
23 */
24 MPI_Type_vector( 100-myrank, 1, 150, MPI_INT, &stype);
25 MPI_Type_commit( &stype );
26 sptr = &sendarray[0][myrank];
27 MPI_Gatherv( sptr, 1, stype, rbuf, rcounts, displs, MPI_INT,
28             root, comm);

```

Example 4.10 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.

```

35 MPI_Comm comm;
36 int gsize, sendarray[100][150], *sptr;
37 int root, *rbuf, stride, myrank, disp[2], blocklen[2];
38 MPI_Datatype stype, types[2];
39 int *displs, i, *rcounts, num;
40
41 ...
42
43 MPI_Comm_size( comm, &gsize);
44 MPI_Comm_rank( comm, &myrank );
45
46 /* First, gather nums to root
47 */
48 rcounts = (int *)malloc(gsize*sizeof(int));

```

```

MPI_Gather( &num, 1, MPI_INT, rcounts, 1, MPI_INT, root, comm);      1
/* root now has correct rcounts, using these we set displs[] so    2
 * that data is placed contiguously (or concatenated) at receive end  3
 */                                                                    4
displs = (int *)malloc(gsize*sizeof(int));                          5
displs[0] = 0;                                                       6
for (i=1; i<gsize; ++i) {                                           7
    displs[i] = displs[i-1]+rcounts[i-1];                             8
}                                                                      9
/* And, create receive buffer                                       10
 */                                                                    11
rbuf = (int *)malloc(gsize*(displs[gsize-1]+rcounts[gsize-1])    12
                    *sizeof(int));                                  13
/* Create datatype for one int, with extent of entire row         14
 */                                                                    15
disp[0] = 0;      disp[1] = 150*sizeof(int);                         16
type[0] = MPI_INT; type[1] = MPI_UB;                                 17
blocklen[0] = 1;  blocklen[1] = 1;                                  18
MPI_Type_struct( 2, blocklen, disp, type, &stype );                 19
MPI_Type_commit( &stype );                                         20
sptr = &sendarray[0][myrank];                                       21
MPI_Gatherv( sptr, num, stype, rbuf, rcounts, displs, MPI_INT,     22
            root, comm);                                             23

```

4.6 Scatter

```

MPI_SCATTER( sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, root, comm) 29
IN      sendbuf      address of send buffer (choice, significant only at root) 30
IN      sendcount    number of elements sent to each process (integer, sig- 31
                    nificant only at root)                               32
IN      sendtype     data type of send buffer elements (significant only at 33
                    root) (handle)                                     34
OUT     recvbuf      address of receive buffer (choice)                 35
IN      recvcount    number of elements in receive buffer (integer)     36
IN      recvtype     data type of receive buffer elements (handle)       37
IN      root         rank of sending process (integer)                   38
IN      comm         communicator (handle)                                39

```

```

int MPI_Scatter(void* sendbuf, int sendcount, MPI_Datatype sendtype, 43
               void* recvbuf, int recvcount, MPI_Datatype recvtype, int root, 44
               MPI_Comm comm)                                          45
MPI_SCATTER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE, 46
            ROOT, COMM, IERROR)                                        47

```

```

1 <type> SENDBUF(*), RECVBUF(*)
2 INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVMYPE, ROOT, COMM, IERROR

```

3 MPI_SCATTER is the inverse operation to MPI_GATHER.

4 The outcome is *as if* the root executed *n* send operations,

```

5 MPI_Send(sendbuf + i · sendcount · extent(sendtype), sendcount, sendtype, i, ...),

```

6 and each process executed a receive,

```

7 MPI_Recv(recvbuf, recvcount, recvmype, i, ...).
8
9

```

10 An alternative description is that the root sends a message with MPI_Send(sendbuf,
11 sendcount·*n*, sendtype, ...). This message is split into *n* equal segments, the *i*th segment is
12 sent to the *i*th process in the group, and each process receives this message as above.

13 The send buffer is ignored for all non-root processes.

14 The type signature associated with sendcount, sendtype at the root must be equal to
15 the type signature associated with recvcount, recvmype at all processes (however, the type
16 maps may be different). This implies that the amount of data sent must be equal to the
17 amount of data received, pairwise between each process and the root. Distinct type maps
18 between sender and receiver are still allowed.

19 All arguments to the function are significant on process root, while on other processes,
20 only arguments recvbuf, recvcount, recvmype, root, comm are significant. The arguments root
21 and comm must have identical values on all processes.

22 The specification of counts and types should not cause any location on the root to be
23 read more than once.

24 *Rationale.* Though not needed, the last restriction is imposed so as to achieve
25 symmetry with MPI_GATHER, where the corresponding restriction (a multiple-write
26 restriction) is necessary. (*End of rationale.*)
27
28

```

29 MPI_SCATTERV( sendbuf, sendcounts, displs, sendtype, recvbuf, recvcount, recvmype, root,
30 comm)

```

31			
32	IN	sendbuf	address of send buffer (choice, significant only at root)
33	IN	sendcounts	integer array (of length group size) specifying the number of elements to send to each processor
34			
35	IN	displs	integer array (of length group size). Entry <i>i</i> specifies the displacement (relative to sendbuf from which to take the outgoing data to process <i>i</i>
36			
37			
38	IN	sendtype	data type of send buffer elements (handle)
39	OUT	recvbuf	address of receive buffer (choice)
40	IN	recvcount	number of elements in receive buffer (integer)
41	IN	recvmype	data type of receive buffer elements (handle)
42	IN	root	rank of sending process (integer)
43	IN	comm	communicator (handle)
44			
45			
46			

```

47 int MPI_Scatterv(void* sendbuf, int *sendcounts, int *displs,
48 MPI_Datatype sendtype, void* recvbuf, int recvcount,

```

```

        MPI_Datatype recvtype, int root, MPI_Comm comm)
MPI_SCATTERV(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, REVCOUNT,
            RECVTYPE, ROOT, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, REVCOUNT, 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 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 the new argument, `displs`.

The outcome is as if the root executed `n` send operations,

```
MPI_Send(sendbuf + displs[i] · extent(sendtype), sendcounts[i], sendtype, i, ...),
```

and each process executed a receive,

```
MPI_Recv(recvbuf, recvcnt, recvtype, i, ...).
```

The send buffer is ignored for all non-root processes.

The type signature implied by `sendcount[i]`, `sendtype` at the root must be equal to the type signature implied by `recvcnt`, `recvtype` at 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 process and the root. Distinct type maps between sender and receiver are still allowed.

All arguments to the function are significant on process `root`, while on other processes, only arguments `recvbuf`, `recvcnt`, `recvtype`, `root`, `comm` are significant. The arguments `root` and `comm` must have identical values on all processes.

The specification of counts, types, and displacements should not cause any location on the root to be read more than once.

4.6.1 Examples using MPI_SCATTER, MPI_SCATTERV


Example 4.11 The reverse of Example 4.2. Scatter sets of 100 ints from the root to each process in the group. See figure 4.7.

```

MPI_Comm comm;
int gsize,*sendbuf;
int root, rbuf[100];
...
MPI_Comm_size( comm, &gsize);
sendbuf = (int *)malloc(gsize*100*sizeof(int));
...
MPI_Scatter( sendbuf, 100, MPI_INT, rbuf, 100, MPI_INT, root, comm);


```

Example 4.12 The reverse of Example 4.5. The root process scatters sets of 100 ints to the other processes, but the sets of 100 are *stride* ints apart in the sending buffer. Requires use of MPI_SCATTERV. Assume *stride* ≥ 100. See figure 4.8.



9
10
11
12
13
14
15
16
17
18
19

Figure 4.7: The root process scatters sets of 100 ints to each process in the group.



42
43
44
45

Figure 4.8: The root process scatters sets of 100 ints, moving by `stride` ints from send to send in the scatter.

```

23 MPI_Comm comm;
24 int gsize,*sendbuf;
25 int root, rbuf[100], i, *displs, *scounts;
26
27 ...
28
29 MPI_Comm_size( comm, &gsize);
30 sendbuf = (int *)malloc(gsize*stride*sizeof(int));
31 ...
32 displs = (int *)malloc(gsize*sizeof(int));
33 scounts = (int *)malloc(gsize*sizeof(int));
34 for (i=0; i<gsize; ++i) {
35     displs[i] = i*stride;
36     scounts[i] = 100;
37 }
38 MPI_Scatterv( sendbuf, scounts, displs, MPI_INT, rbuf, 100, MPI_INT,
39             root, comm);
40
41

```

42
43
44
45

Example 4.13 The reverse of Example 4.9. We have a varying stride between blocks at sending (root) side, at the receiving side we receive into the *i*th column of a 100×150 C array. See figure 4.9.

```

46 MPI_Comm comm;
47 int gsize,recvarray[100][150],*rptr;
48 int root, *sendbuf, myrank, bufsize, *stride;

```




Figure 4.9: The root scatters blocks of $100-i$ ints into column i of a 100×150 C array. At the sending side, the blocks are `stride[i]` ints apart.

```

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);

```

4.7 Gather-to-all

```

1  MPI_ALLGATHER( sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm)
2
3
4
5  IN      sendbuf      starting address of send buffer (choice)
6
7  IN      sendcount    number of elements in send buffer (integer)
8
9  IN      sendtype     data type of send buffer elements (handle)
10
11 OUT     recvbuf      address of receive buffer (choice)
12
13 IN      recvcount    number of elements received from any process (inte-
14                      ger)
15
16 IN      recvtype     data type of receive buffer elements (handle)
17
18 IN      comm         communicator (handle)

```

```

19 int MPI_Allgather(void* sendbuf, int sendcount, MPI_Datatype sendtype,
20                 void* recvbuf, int recvcount, MPI_Datatype recvtype,
21                 MPI_Comm comm)
22
23 MPI_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, REVCOUNT, RECVTYPE,
24              COMM, IERROR)
25
26 <type> SENDBUF(*), RECVBUF(*)
27
28 INTEGER SENDCOUNT, SENDTYPE, REVCOUNT, RECVTYPE, COMM, IERROR

```

MPI_ALLGATHER can be thought of as MPI_GATHER, but where all processes receive the result, instead of just the root. The block of data sent from the *j*th process is received by every process and placed in the *j*th block of the buffer *recvbuf*.

The type signature associated with *sendcount*, *sendtype*, at a process must be equal to the type signature associated with *recvcount*, *recvtype* at any other process.

The outcome of a call to MPI_ALLGATHER(...) is as if all processes executed *n* calls to

```

31 MPI_GATHER(sendbuf, sendcount, sendtype, recvbuf, recvcount,
32           recvtype, root, comm),

```

for *root* = 0 , . . . , *n*-1. The rules for correct usage of MPI_ALLGATHER are easily found from the corresponding rules for MPI_GATHER.

MPI_ALLGATHERV(sendbuf, sendcount, sendtype, recvbuf, recvcunts, displs, recvtype, comm)			1
			2
IN	sendbuf	starting address of send buffer (choice)	3
			4
IN	sendcount	number of elements in send buffer (integer)	5
IN	sendtype	data type of send buffer elements (handle)	6
OUT	recvbuf	address of receive buffer (choice)	7
			8
IN	recvcunts	integer array (of length group size) containing the number of elements that are received from each process	9
			10
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 process <i>i</i>	11
			12
			13
IN	recvtype	data type of receive buffer elements (handle)	14
			15
IN	comm	communicator (handle)	16
			17
int MPI_Allgatherv(void* sendbuf, int sendcount, MPI_Datatype sendtype, void* recvbuf, int *recvcunts, int *displs, MPI_Datatype recvtype, MPI_Comm comm)			18
			19
			20
MPI_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS, RECVTYPE, COMM, IERROR)			21
			22
			23
			24
			25
			26

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 process is received by every 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 process *j* must be equal to the type signature associated with *recvcunts*[*j*], *recvtype* at any other process.

The outcome is as if all processes executed calls to

```
MPI_GATHERV(sendbuf, sendcount, sendtype, recvbuf, recvcunts, displs,
            recvtype, root, comm),
```

for *root* = 0 , . . . , *n*-1. The rules for correct usage of MPI_ALLGATHERV are easily found from the corresponding rules for MPI_GATHERV.

4.7.1 Examples using MPI_ALLGATHER, MPI_ALLGATHERV

Example 4.14 The all-gather version of Example 4.2. Using MPI_ALLGATHER, we will gather 100 ints from every process in the group to every process.

```
MPI_Comm comm;
int gsize, sendarray[100];
int *rbuf;
...
```



```

1   MPI_Comm_size( comm, &gsize);
2   rbuf = (int *)malloc(gsize*100*sizeof(int));
3   MPI_Allgather( sendarray, 100, MPI_INT, rbuf, 100, MPI_INT, comm);

```

After the call, every process has the group-wide concatenation of the sets of data.

4.8 All-to-All Scatter/Gather

```

11  MPI_ALLTOALL(sendbuf, sendcount, sendtype, recvbuf, recvcount, recvtype, comm)
12
13  IN      sendbuf      starting address of send buffer (choice)
14  IN      sendcount    number of elements sent to each process (integer)
15  IN      sendtype     data type of send buffer elements (handle)
16  OUT     recvbuf      address of receive buffer (choice)
17  IN      recvcount    number of elements received from any process (inte-
18  ger)
19
20  IN      recvtype     data type of receive buffer elements (handle)
21  IN      comm         communicator (handle)

```

```

23  int MPI_Alltoall(void* sendbuf, int sendcount, MPI_Datatype sendtype,
24                  void* recvbuf, int recvcount, MPI_Datatype recvtype,
25                  MPI_Comm comm)
26
27  MPI_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
28              COMM, IERROR)
29  <type> SENDBUF(*), RECVBUF(*)
30  INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, IERROR

```

MPI_ALLTOALL is an extension of MPI_ALLGATHER to the case where each process sends distinct data to each of the receivers. The j th block sent from process i is received by process j and is placed in the i th block of `recvbuf`.

The type signature associated with `sendcount`, `sendtype`, at a process must be equal to the type signature associated with `recvcount`, `recvtype` at any other process. This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of processes. As usual, however, the type maps may be different.

The outcome is as if each process executed a send to each process (itself included) with a call to,

```
MPI_Send(sendbuf + i · sendcount · extent(sendtype), sendcount, sendtype, i, ...),
```

and a receive from every other process with a call to,

```
MPI_Recv(recvbuf + i · recvcount · extent(recvtype), recvcount, i, ...).
```

All arguments on all processes are significant. The argument `comm` must have identical values on all processes.

MPI_ALLTOALLV(sendbuf, sendcounts, sdispls, sendtype, recvbuf, recvcounts, rdispls, recvtype, comm)			1
			2
IN	sendbuf	starting address of send buffer (choice)	3
			4
IN	sendcounts	integer array equal to the group size specifying the number of elements to send to each processor	5
			6
IN	sdispls	integer array (of length group size). Entry j specifies the displacement (relative to sendbuf from which to take the outgoing data destined for process j	7
			8
			9
IN	sendtype	data type of send buffer elements (handle)	10
OUT	recvbuf	address of receive buffer (choice)	11
			12
IN	recvcounts	integer array equal to the group size specifying the number of elements that can be received from each processor	13
			14
			15
IN	rdispls	integer array (of length group size). Entry i specifies the displacement (relative to recvbuf at which to place the incoming data from process i	16
			17
			18
			19
IN	recvtype	data type of receive buffer elements (handle)	20
			21
IN	comm	communicator (handle)	22

```

int MPI_Alltoallv(void* sendbuf, int *sendcounts, int *sdispls,
                 MPI_Datatype sendtype, void* recvbuf, int *recvcounts,
                 int *rdispls, MPI_Datatype recvtype, MPI_Comm comm)
MPI_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, RECVCOUNTS,
              RDISPLS, RECVTYPE, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), 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`.

The j th block sent from process i is received by 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 `sendcount[j]`, `sendtype` at process i must be equal to the type signature associated with `recvcount[i]`, `recvtype` at process j . This implies that the amount of data sent must be equal to the amount of data received, pairwise between every pair of processes. Distinct type maps between sender and receiver are still allowed.

The outcome is as if each process sent a message to every other process with,

```
MPI_Send(sendbuf + displs[i] · extent(sendtype), sendcounts[i], sendtype, i, ...),
```

and received a message from every other process with a call to

```
MPI_Recv(recvbuf + displs[i] · extent(recvtype), recvcounts[i], recvtype, i, ...).
```

All arguments on all processes are significant. The argument `comm` must have identical values on all processes.

Rationale. The definitions of `MPI_ALLTOALL` and `MPI_ALLTOALLV` give as much flexibility as one would achieve by specifying `n` independent, point-to-point communications, with two exceptions: all messages use the same datatype, and messages are scattered from (or gathered to) sequential storage. (*End of rationale.*)

Advice to implementors. Although the discussion of collective communication in terms of point-to-point operation implies that each message is transferred directly from sender to receiver, implementations may use a tree communication pattern. Messages can be forwarded by intermediate nodes where they are split (for scatter) or concatenated (for gather), if this is more efficient. (*End of advice to implementors.*)

4.9 Global Reduction Operations

The functions in this section perform a global reduce operation (such as sum, max, logical AND, etc.) across all the members of a group. The reduction operation can be either one of a predefined list of operations, or a user-defined operation. The global reduction functions come in several flavors: a reduce that returns the result of the reduction at one node, an all-reduce that returns this result at all nodes, and a scan (parallel prefix) operation. In addition, a reduce-scatter operation combines the functionality of a reduce and of a scatter operation.

4.9.1 Reduce

`MPI_REDUCE(sendbuf, recvbuf, count, datatype, op, root, comm)`

IN	<code>sendbuf</code>	address of send buffer (choice)
OUT	<code>recvbuf</code>	address of receive buffer (choice, significant only at root)
IN	<code>count</code>	number of elements in send buffer (integer)
IN	<code>datatype</code>	data type of elements of send buffer (handle)
IN	<code>op</code>	reduce operation (handle)
IN	<code>root</code>	rank of root process (integer)
IN	<code>comm</code>	communicator (handle)

```
int MPI_Reduce(void* sendbuf, void* recvbuf, int count,
              MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)
```

```
MPI_REDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER COUNT, DATATYPE, OP, ROOT, COMM, IERROR
```

`MPI_REDUCE` combines the elements provided in the input buffer of each process in the group, using the operation `op`, and returns the combined value in the output buffer of the 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 processes provide input buffers and output buffers of the same length, with elements of the same type. Each 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))`.

Sec. 4.9.2, lists the set of predefined operations provided by MPI. That section also enumerates the datatypes each operation can be applied to. 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 Sec. 4.9.4.

The operation `op` is always assumed to be associative. All predefined operations are also assumed to be commutative. Users may define operations that are assumed to be associative, but not commutative. The “canonical” evaluation order of a reduction is determined by the ranks of the processes in the group. However, the implementation can take advantage of associativity, or associativity and commutativity in order to change the order of evaluation. This may change the result of the reduction for operations that are not strictly associative and commutative, such as floating point addition.

Advice to implementors. It is strongly recommended that `MPI_REDUCE` 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 processors. (*End of advice to implementors.*)

The `datatype` argument of `MPI_REDUCE` must be compatible with `op`. Predefined operators work only with the MPI types listed in Sec. 4.9.2 and Sec. 4.9.3. Furthermore, the `datatype` and `op` given for predefined operators must be the same on all processes.

Note that it is possible for users to supply different user-defined operations to `MPI_REDUCE` in each process. MPI does not define which operations are used on which operands in this case. User-defined operators may operate on general, derived datatypes. In this case, each argument that the reduce operation is applied to is one element described by such a datatype, which may contain several basic values. This is further explained in Section 4.9.4.

Advice to users. Users should make no assumptions about how `MPI_REDUCE` is implemented. Safest is to ensure that the same function is passed to `MPI_REDUCE` by each process. (*End of advice to users.*)

Overlapping datatypes are permitted in “send” buffers. Overlapping datatypes in “receive” buffers are erroneous and may give unpredictable results.

4.9.2 Predefined reduce operations

The following predefined operations are supplied for `MPI_REDUCE` and related functions `MPI_ALLREDUCE`, `MPI_REDUCE_SCATTER`, and `MPI_SCAN`. These operations are invoked by placing the following in `op`.

Name	Meaning
------	---------

1		
2	MPI_MAX	maximum
3	MPI_MIN	minimum
4	MPI_SUM	sum
5	MPI_PROD	product
6	MPI_LAND	logical and
7	MPI_BAND	bit-wise and
8	MPI_LOR	logical or
9	MPI_BOR	bit-wise or
10	MPI_LXOR	logical xor
11	MPI_BXOR	bit-wise xor
12	MPI_MAXLOC	max value and location
13	MPI_MINLOC	min value and location

14 The two operations MPI_MINLOC and MPI_MAXLOC are discussed separately in Sec.
 15 4.9.3. For the other predefined operations, we enumerate below the allowed combinations
 16 of op and datatype arguments. First, define groups of MPI basic datatypes in the following
 17 way.

18		
19		
20	C integer:	MPI_INT, MPI_LONG, MPI_SHORT,
21		MPI_UNSIGNED_SHORT, MPI_UNSIGNED,
22		MPI_UNSIGNED_LONG
23	Fortran integer:	MPI_INTEGER
24	Floating point:	MPI_FLOAT, MPI_DOUBLE, MPI_REAL,
25		MPI_DOUBLE_PRECISION, MPI_LONG_DOUBLE
26	Logical:	MPI_LOGICAL
27	Complex:	MPI_COMPLEX
28	Byte:	MPI_BYTE

29 Now, the valid datatypes for each option is specified below.

30		
31		
32	Op	Allowed Types
33		
34	MPI_MAX, MPI_MIN	C integer, Fortran integer, Floating point
35	MPI_SUM, MPI_PROD	C integer, Fortran integer, Floating point, Complex
36	MPI_LAND, MPI_LOR, MPI_LXOR	C integer, Logical
37	MPI_BAND, MPI_BOR, MPI_BXOR	C integer, Fortran integer, Byte
38		

39 **Example 4.15** A routine that computes the dot product of two vectors that are distributed
 40 across a group of processes and returns the answer at node zero.

```

41
42 SUBROUTINE PAR_BLAS1(m, a, b, c, comm)
43 REAL a(m), b(m)      ! local slice of array
44 REAL c               ! result (at node zero)
45 REAL sum
46 INTEGER m, comm, i, ierr
47
48 ! 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

```

Example 4.16 A routine that computes the product of a vector and an array that are distributed across a group of processes and returns the answer at node zero.

```

SUBROUTINE PAR_BLAS2(m, n, a, b, c, comm)
REAL a(m), b(m,n)    ! local slice of array
REAL c(n)           ! result
REAL sum(n)
INTEGER 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

```

4.9.3 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 application of these is to compute a global minimum (maximum) and the rank of the process containing this value.

The operation that defines `MPI.MAXLOC` is:

$$\begin{pmatrix} u \\ i \end{pmatrix} \circ \begin{pmatrix} v \\ j \end{pmatrix} = \begin{pmatrix} w \\ k \end{pmatrix}$$

where

$$w = \max(u, v)$$

and

$$k = \begin{cases} i & \text{if } u > v \\ \min(i, j) & \text{if } u = v \\ j & \text{if } u < v \end{cases}$$

1 MPI_MINLOC is defined similarly:

$$2 \quad 3 \quad 4 \quad 5 \quad \begin{pmatrix} u \\ i \end{pmatrix} \circ \begin{pmatrix} v \\ j \end{pmatrix} = \begin{pmatrix} w \\ k \end{pmatrix}$$

6 where

$$7 \quad 8 \quad w = \min(u, v)$$

9 and

$$10 \quad 11 \quad 12 \quad 13 \quad k = \begin{cases} i & \text{if } u < v \\ \min(i, j) & \text{if } u = v \\ j & \text{if } u > v \end{cases}$$

14 Both operations are associative and commutative. Note that if MPI_MAXLOC is applied
15 to reduce a sequence of pairs $(u_0, 0), (u_1, 1), \dots, (u_{n-1}, n-1)$, then the value returned is
16 (u, r) , where $u = \max_i u_i$ and r is the index of the first global maximum in the sequence.
17 Thus, if each process supplies a value and its rank within the group, then a reduce operation
18 with `op = MPI_MAXLOC` will return the maximum value and the rank of the first process
19 with that value. Similarly, MPI_MINLOC can be used to return a minimum and its index.
20 More generally, MPI_MINLOC computes a *lexicographic minimum*, where elements are ordered
21 according to the first component of each pair, and ties are resolved according to the second
22 component.

23 The reduce operation is defined to operate on arguments that consist of a pair: value
24 and index. For both Fortran and C, types are provided to describe the pair. The potentially
25 mixed-type nature of such arguments is a problem in Fortran. The problem is circumvented,
26 for Fortran, by having the MPI-provided type consist of a pair of the same type as value,
27 and coercing the index to this type also. In C, the MPI-provided pair type has distinct
28 types and the index is an `int`.

29 In order to use MPI_MINLOC and MPI_MAXLOC in a reduce operation, one must provide
30 a `datatype` argument that represents a pair (value and index). MPI provides nine such
31 predefined datatypes. The operations MPI_MAXLOC and MPI_MINLOC can be used with each
32 of the following datatypes.

33
34 Fortran:

35 Name	Description
36 MPI_2REAL	pair of REALs
37 MPI_2DOUBLE_PRECISION	pair of DOUBLE PRECISION variables
38 MPI_2INTEGER	pair of INTEGERS

39
40
41 C:

42 Name	Description
43 MPI_FLOAT_INT	float and int
44 MPI_DOUBLE_INT	double and int
45 MPI_LONG_INT	long and int
46 MPI_2INT	pair of int
47 MPI_SHORT_INT	short and int
48 MPI_LONG_DOUBLE_INT	long double and int

The datatype MPI_2REAL is *as if* defined by the following (see Section 3.12).

```
MPI_TYPE_CONTIGUOUS(2, MPI_REAL, MPI_2REAL)
```

Similar statements apply for MPI_2INTEGER, MPI_2DOUBLE_PRECISION, and MPI_2INT.

The datatype MPI_FLOAT_INT is *as if* defined by the following sequence of instructions.

```
type[0] = MPI_FLOAT
type[1] = MPI_INT
disp[0] = 0
disp[1] = sizeof(float)
block[0] = 1
block[1] = 1
MPI_TYPE_STRUCT(2, block, disp, type, MPI_FLOAT_INT)
```

Similar statements apply for MPI_LONG_INT and MPI_DOUBLE_INT.

Example 4.17 Each process has an array of 30 doubles, in C. For each of the 30 locations, compute the value and rank of the process containing the largest value.

```
...
/* each 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(MPI_COMM_WORLD, &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 process root
*/
if (myrank == root) {
    /* read ranks out
    */
    for (i=0; i<30; ++i) {
        aout[i] = out[i].val;
        ind[i] = out[i].rank;
    }
}
```

Example 4.18 Same example, in Fortran.


```

1      ...
2      ! each process has an array of 30 double: ain(30)
3
4      DOUBLE PRECISION ain(30), aout(30)
5      INTEGER ind(30);
6      DOUBLE PRECISION in(2,30), out(2,30)
7      INTEGER i, myrank, root, ierr;
8
9      MPI_COMM_RANK(MPI_COMM_WORLD, myrank);
10     DO I=1, 30
11         in(1,i) = ain(i)
12         in(2,i) = myrank    ! myrank is coerced to a double
13     END DO
14
15     MPI_REDUCE( in, out, 30, MPI_2DOUBLE_PRECISION, MPI_MAXLOC, root,
16                                     comm, ierr );
17     ! At this point, the answer resides on process root
18
19     IF (myrank .EQ. root) THEN
20         ! read ranks out
21         DO I= 1, 30
22             aout(i) = out(1,i)
23             ind(i) = out(2,i) ! rank is coerced back to an integer
24         END DO
25     END IF

```

Example 4.19 Each process has a non-empty array of values. Find the minimum global value, the rank of the process that holds it and its index on this process.

```

29     #define LEN 1000
30
31     float val[LEN];          /* local array of values */
32     int count;              /* local number of values */
33     int myrank, minrank, minindex;
34     float minval;
35
36
37     struct {
38         float value;
39         int index;
40     } in, out;
41
42     /* local minloc */
43     in.value = val[0];
44     in.index = 0;
45     for (i=1; i < count; i++)
46         if (in.value > val[i]) {
47             in.value = val[i];
48             in.index = i;

```

```

    }
    /* global minloc */
MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
in.index = myrank*LEN + in.index;
MPI_Reduce( in, out, 1, MPI_FLOAT_INT, MPI_MINLOC, root, comm );
    /* At this point, the answer resides on process root
    */
if (myrank == root) {
    /* read answer out
    */
    minval = out.value;
    minrank = out.index / LEN;
    minindex = out.index % LEN;
}

```

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. A programmer can provide his or her own 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.*)

4.9.4 User-Defined Operations

MPI_OP_CREATE(function, commute, op)

IN	function	user defined function (function)
IN	commute	true if commutative; false otherwise.
OUT	op	operation (handle)

int MPI_Op_create(MPI_User_function *function, int commute, MPI_Op *op)

MPI_OP_CREATE(FUNCTION, COMMUTE, OP, IERROR)

EXTERNAL FUNCTION
 LOGICAL COMMUTE
 INTEGER OP, IERROR

MPI_OP_CREATE binds a user-defined global operation to an op handle that can subsequently be used in MPI_REDUCE, MPI_ALLREDUCE, MPI_REDUCE_SCATTER, and MPI_SCAN. 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 process zero. 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.

`function` is the user-defined function, which must have the following four arguments: `invec`, `inoutvec`, `len` and `datatype`.

1 The ANSI-C prototype for the function is the following.

```
2
3 typedef void MPI_User_function( void *invec, void *inoutvec, int *len,
4                                 MPI_Datatype *datatype);
```

5 The Fortran declaration of the user-defined function appears below.

```
6
7 SUBROUTINE USER_FUNCTION( INVEC, INOUTVEC, LEN, TYPE)
8   <type> INVEC(LEN), INOUTVEC(LEN)
9   INTEGER LEN, TYPE
```

10 The `datatype` argument is a handle to the data type that was passed into the call to
 11 `MPI_REDUCE`. The user reduce function should be written such that the following holds:
 12 Let $u[0], \dots, u[\text{len}-1]$ be the `len` elements in the communication buffer described by the
 13 arguments `invec`, `len` and `datatype` when the function is invoked; let $v[0], \dots, v[\text{len}-1]$ be `len`
 14 elements in the communication buffer described by the arguments `inoutvec`, `len` and `datatype`
 15 when the function is invoked; let $w[0], \dots, w[\text{len}-1]$ be `len` elements in the communication
 16 buffer described by the arguments `inoutvec`, `len` and `datatype` when the function returns;
 17 then $w[i] = u[i] \circ v[i]$, for $i=0, \dots, \text{len}-1$, where \circ is the reduce operation that the function
 18 computes.

19 Informally, we can think of `invec` and `inoutvec` as arrays of `len` elements that `function`
 20 is combining. The result of the reduction over-writes values in `inoutvec`, hence the name.
 21 Each invocation of the function results in the pointwise evaluation of the reduce operator
 22 on `len` elements: I.e, the function returns in `inoutvec[i]` the value `invec[i] \circ inoutvec[i]`, for
 23 $i = 0, \dots, \text{count} - 1$, where \circ is the combining operation computed by the function.

24
 25 *Rationale.* The `len` argument allows `MPI_REDUCE` to avoid calling the function for
 26 each element in the input buffer. Rather, the system can choose to apply the function
 27 to chunks of input. In C, it is passed in as a reference for reasons of compatibility
 28 with Fortran.

29 By internally comparing the value of the `datatype` argument to known, global handles,
 30 it is possible to overload the use of a single user-defined function for several, different
 31 data types. (*End of rationale.*)

32
 33 General datatypes may be passed to the user function. However, use of datatypes that
 34 are not contiguous is likely to lead to inefficiencies.

35 No MPI communication function may be called inside the user function. `MPI_ABORT`
 36 may be called inside the function in case of an error.

37
 38 *Advice to users.* Suppose one defines a library of user-defined reduce functions that
 39 are overloaded: the `datatype` argument is used to select the right execution path at each
 40 invocation, according to the types of the operands. The user-defined reduce function
 41 cannot “decode” the `datatype` argument that it is passed, and cannot identify, by itself,
 42 the correspondence between the `datatype` handles and the `datatype` they represent.
 43 This correspondence was established when the datatypes were created. Before the
 44 library is used, a library initialization preamble must be executed. This preamble
 45 code will define the datatypes that are used by the library, and store handles to these
 46 datatypes in global, static variables that are shared by the user code and the library
 47 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`.

```

if (rank > 0) {
    RECV(tempbuf, count, datatype, rank-1,...)
    User_reduce( tempbuf, sendbuf, count, datatype)
}
if (rank < groupsize-1) {
    SEND( sendbuf, count, datatype, rank+1, ...)
}
/* answer now resides in process groupsize-1 ... now send to root
*/
if (rank == groupsize-1) {
    SEND( sendbuf, count, datatype, root, ...)
}
if (rank == root) {
    RECV(recvbuf, count, datatype, groupsize-1,...)
}

```

The reduction computation proceeds, sequentially, from process 0 to process `group-size-1`. This order is chosen so as to respect the order of a possibly non-commutative 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 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)

`int MPI_op_free(MPI_Op *op)`

`MPI_OP_FREE(OP, IERROR)`

INTEGER OP, IERROR

Marks a user-defined reduction operation for deallocation and sets `op` to `MPI_OP_NULL`.

1 Example of User-defined Reduce

2 It is time for an example of user-defined reduction.

3
4 **Example 4.20** Compute the product of an array of complex numbers, in C.

```

5
6 typedef struct {
7     double real,imag;
8 } Complex;
9
10 /* the user-defined function
11 */
12 void myProd( Complex *in, Complex *inout, int *len, MPI_Datatype *dptr )
13 {
14     int i;
15     Complex c;
16
17     for (i=0; i< *len; ++i) {
18         c.real = inout->real*in->real -
19             inout->imag*in->imag;
20         c.imag = inout->real*in->imag +
21             inout->imag*in->real;
22         *inout = c;
23         in++; inout++;
24     }
25 }
26
27 /* and, to call it...
28 */
29 ...
30
31 /* each process has an array of 100 Complexes
32 */
33 Complex a[100], answer[100];
34 MPI_Op myOp;
35 MPI_Datatype ctype;
36
37 /* explain to MPI how type Complex is defined
38 */
39 MPI_Type_contiguous( 2, MPI_DOUBLE, &ctype );
40 MPI_Type_commit( &ctype );
41 /* create the complex-product user-op
42 */
43 MPI_Op_create( myProd, True, &myOp );
44
45 MPI_Reduce( a, answer, 100, ctype, myOp, root, comm );
46
47 /* At this point, the answer, which consists of 100 Complexes,
48 * resides on process root

```

*/

4.9.5 All-Reduce

MPI includes variants of each of the reduce operations where the result is returned to all processes in the group. MPI requires that all processes 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 (integer)
IN	datatype	data type of elements of send buffer (handle)
IN	op	operation (handle)
IN	comm	communicator (handle)

```
int MPI_Allreduce(void* sendbuf, void* recvbuf, int count,
                 MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
```

```
MPI_ALLREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER COUNT, DATATYPE, OP, COMM, IERROR
```

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.*)

Example 4.21 A routine that computes the product of a vector and an array that are distributed across a group of processes and returns the answer at all nodes (see also Example 4.16).

```
SUBROUTINE PAR_BLAS2(m, n, a, b, c, comm)
REAL a(m), b(m,n)    ! local slice of array
REAL c(n)           ! result
REAL sum(n)
INTEGER 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
```

```

1
2 ! global sum
3 CALL MPI_ALLREDUCE(sum, c, n, MPI_REAL, MPI_SUM, comm, ierr)
4
5 ! return result at all nodes
6 RETURN
7

```

4.10 Reduce-Scatter

MPI includes variants of each of the reduce operations where the result is scattered to all processes in the group on return.

`MPI_REDUCE_SCATTER(sendbuf, recvbuf, recvcnts, datatype, op, comm)`

IN	sendbuf	starting address of send buffer (choice)
OUT	recvbuf	starting address of receive buffer (choice)
IN	recvcnts	integer array specifying the number of elements in result distributed to each process. Array must be identical on all calling processes.
IN	datatype	data type of elements of input buffer (handle)
IN	op	operation (handle)
IN	comm	communicator (handle)

```

26 int MPI_Reduce_scatter(void* sendbuf, void* recvbuf, int *recvcnts,
27                       MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)
28

```

```

29 MPI_REDUCE_SCATTER(SENDBUF, RECVBUF, RECVCOUNTS, DATATYPE, OP, COMM,
30                   IERROR)
31 <type> SENDBUF(*), RECVBUF(*)
32 INTEGER RECVCOUNTS(*), DATATYPE, OP, COMM, IERROR
33

```

`MPI_REDUCE_SCATTER` first does an element-wise reduction on vector of `count = $\sum_i \text{recvcnts}[i]$` elements in the send buffer defined by `sendbuf`, `count` and `datatype`. Next, the resulting vector of results is split into `n` disjoint segments, where `n` is the number of members in the group. Segment `i` contains `recvcnts[i]` elements. The `i`th segment is sent to process `i` and stored in the receive buffer defined by `recvbuf`, `recvcnts[i]` and `datatype`.

Advice to implementors. The `MPI_REDUCE_SCATTER` routine is functionally equivalent to: A `MPI_REDUCE` operation function with `count` equal to the sum of `recvcnts[i]` followed by `MPI_SCATTERV` with `sendcounts` equal to `recvcnts`. However, a direct implementation may run faster. (*End of advice to implementors.*)

4.11 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 (integer)
IN	datatype	data type of elements of input buffer (handle)
IN	op	operation (handle)
IN	comm	communicator (handle)

```
int MPI_Scan(void* sendbuf, void* recvbuf, int count,
             MPI_Datatype datatype, MPI_Op op, MPI_Comm comm )
```

```
MPI_SCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
<type> SENDBUF(*), RECVBUF(*)
INTEGER COUNT, DATATYPE, OP, COMM, IERROR
```

MPI_SCAN is used to perform a prefix reduction on data distributed across the group. The operation returns, in the receive buffer of the process with rank i , the reduction of the values in the send buffers of processes with ranks $0, \dots, i$ (inclusive). The type of operations supported, their semantics, and the constraints on send and receive buffers are as for MPI_REDUCE.

Rationale. We have defined an inclusive scan, that is, the prefix reduction on process i includes the data from process i . An alternative is to define scan in an exclusive manner, where the result on i only includes data up to $i-1$. Both definitions are useful. The latter has some advantages: the inclusive scan can always be computed from the exclusive scan with no additional communication; for non-invertible operations such as max and min, communication is required to compute the exclusive scan from the inclusive scan. There is, however, a complication with exclusive scan since one must define the “unit” element for the reduction in this case. That is, one must explicitly say what occurs for process 0. This was thought to be complex for user-defined operations and hence, the exclusive scan was dropped. (*End of rationale.*)

4.11.1 Example using MPI_SCAN

Example 4.22 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 non-commutative operator. C code that implements it is given below.

```

11 typedef struct {
12     double val;
13     int log;
14 } SegScanPair;
15
16 /* the user-defined function
17 */
18 void segScan( SegScanPair *in, SegScanPair *inout, int *len,
19              MPI_Datatype *dptr )
20 {
21     int i;
22     SegScanPair c;
23
24     for (i=0; i< *len; ++i) {
25         if ( in->log == inout->log )
26             c.val = in->val + inout->val;
27         else
28             c.val = inout->val;
29         c.log = inout->log;
30         *inout = c;
31         in++; inout++;
32     }
33 }
34

```

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 non-commutative, as in the following.

```

39     int i,base;
40     SeqScanPair a, answer;
41     MPI_Op      myOp;
42     MPI_Datatype type[2] = {MPI_DOUBLE, MPI_INT};
43     MPI_Aint     disp[2];
44     int          blocklen[2] = { 1, 1};
45     MPI_Datatype sspair;
46
47     /* explain to MPI how type SegScanPair is defined
48 */

```

```

MPI_Address( a, disp);           1
MPI_Address( a.log, disp+1);     2
base = disp[0];                  3
for (i=0; i<2; ++i) disp[i] -= base; 4
MPI_Type_struct( 2, blocklen, disp, type, &sspair ); 5
MPI_Type_commit( &sspair );     6
/* create the segmented-scan user-op 7
  */                               8
MPI_Op_create( segScan, False, &myOp ); 9
...                               10
MPI_Scan( a, answer, 1, sspair, myOp, comm ); 11

```

4.12 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.

Example 4.23 The following is erroneous.

```

switch(rank) {                  21
  case 0:                        22
    MPI_Bcast(buf1, count, type, 0, comm); 23
    MPI_Bcast(buf2, count, type, 1, comm); 24
    break;                       25
  case 1:                        26
    MPI_Bcast(buf2, count, type, 1, comm); 27
    MPI_Bcast(buf1, count, type, 0, comm); 28
    break;                       29
}                                 30

```

We assume that the group of `comm` is $\{0,1\}$. Two 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 4.24 The following is erroneous.

```

switch(rank) {                  39
  case 0:                        40
    MPI_Bcast(buf1, count, type, 0, comm0); 41
    MPI_Bcast(buf2, count, type, 2, comm2); 42
    break;                       43
  case 1:                        44
    MPI_Bcast(buf1, count, type, 1, comm1); 45
    MPI_Bcast(buf2, count, type, 0, comm0); 46
    break;                       47
  case 2:                        48

```

```

1     MPI_Bcast(buf1, count, type, 2, comm2);
2     MPI_Bcast(buf2, count, type, 1, comm1);
3     break;
4 }

```

5

6 Assume that the group of comm0 is {0,1}, of comm1 is {1, 2} and of comm2 is {2,0}. If
7 the broadcast is a synchronizing operation, then there is a cyclic dependency: the broadcast
8 in comm2 completes only after the broadcast in comm0; the broadcast in comm0 completes
9 only after the broadcast in comm1; and the broadcast in comm1 completes only after the
10 broadcast in comm2. Thus, the code will deadlock.

11 Collective operations must be executed in an order so that no cyclic dependences occur.

12

13 **Example 4.25** The following is erroneous.

14

```

15 switch(rank) {
16     case 0:
17         MPI_Bcast(buf1, count, type, 0, comm);
18         MPI_Send(buf2, count, type, 1, tag, comm);
19         break;
20     case 1:
21         MPI_Recv(buf2, count, type, 0, tag, comm, status);
22         MPI_Bcast(buf1, count, type, 0, comm);
23         break;
24 }

```

25

26 Process zero executes a broadcast, followed by a blocking send operation. Process one
27 first executes a blocking receive that matches the send, followed by broadcast call that
28 matches the broadcast of process zero. This program may deadlock. The broadcast call on
29 process zero *may* block until process one executes the matching broadcast call, so that the
30 send is not executed. Process one will definitely block on the receive and so, in this case,
31 never executes the broadcast.

32 The relative order of execution of collective operations and point-to-point operations
33 should be such, so that even if the collective operations and the point-to-point operations
34 are synchronizing, no deadlock will occur.

35 **Example 4.26** A correct, but non-deterministic program.

36

```

37 switch(rank) {
38     case 0:
39         MPI_Bcast(buf1, count, type, 0, comm);
40         MPI_Send(buf2, count, type, 1, tag, comm);
41         break;
42     case 1:
43         MPI_Recv(buf2, count, type, MPI_ANY_SOURCE, tag, comm, status);
44         MPI_Bcast(buf1, count, type, 0, comm);
45         MPI_Recv(buf2, count, type, MPI_ANY_SOURCE, tag, comm, status);
46         break;
47     case 2:
48         MPI_Send(buf2, count, type, 1, tag, comm);

```

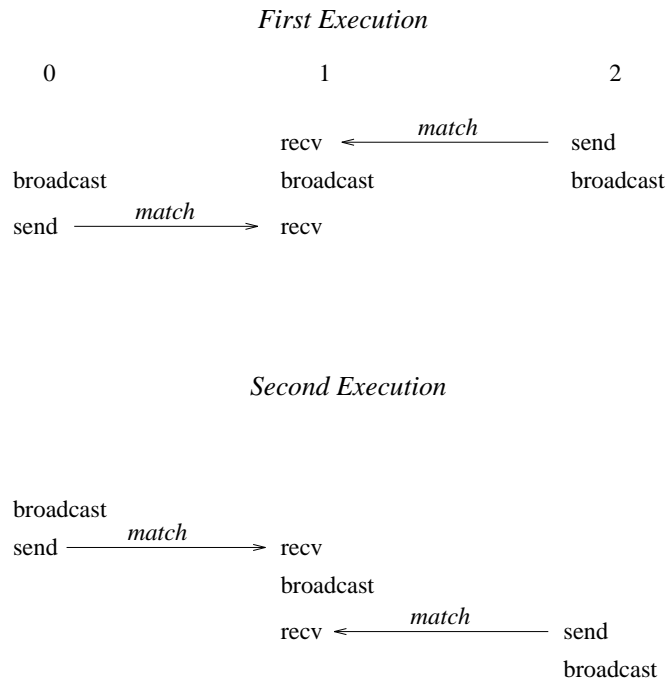


Figure 4.10: A race condition causes non-deterministic matching of sends and receives. One cannot rely on synchronization from a broadcast to make the program deterministic.

```

MPI_Bcast(buf1, count, type, 0, comm);
break;
}

```

All three processes participate in a broadcast. Process 0 sends a message to process 1 after the broadcast, and process 2 sends a message to process 1 before the broadcast. Process 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 4.10. 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 call at a 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 calls at the same process.

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 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
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
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

Chapter 5

Groups, Contexts, and Communicators

5.1 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 [26] and [3] for further information on writing libraries in MPI, using the features described in this chapter.

5.1.1 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 processes (potentially running unrelated code),
- Abstract process naming 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 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 processes, to house abstract process naming, and to store adornments.

5.1.2 MPI's Support for Libraries

The corresponding concepts that MPI provides, specifically to support robust libraries, are as follows:

- **Contexts** of communication,
- **Groups** of processes,
- **Virtual topologies**,
- **Attribute caching**,
- **Communicators**.

Communicators (see [16, 24, 27]) encapsulate all of these ideas in order to provide the appropriate scope for all communication operations in MPI. Communicators are divided into two kinds: intra-communicators for operations within a single group of processes, and inter-communicators, for point-to-point communication between two groups of processes.

Caching. Communicators (see below) provide a “caching” mechanism that allows one to associate new attributes with communicators, on a par with MPI built-in features. This can be used by advanced users to adorn communicators further, and by MPI to implement some communicator functions. For example, the virtual-topology functions described in Chapter 6 are likely to be supported this way.

Groups. Groups define an ordered collection of processes, each with a rank, and it is this group that defines the low-level names for inter-process communication (ranks are used for sending and receiving). Thus, groups define a scope for process names in point-to-point communication. In addition, groups define the scope of collective operations. Groups may be manipulated separately from communicators in MPI, but only communicators can be used in communication operations.

Intra-communicators. The most commonly used means for message passing in MPI is via intra-communicators. Intra-communicators contain an instance of a group, contexts of communication for both point-to-point and collective communication, and the ability to include virtual topology and other attributes. These features work as follows:

- **Contexts** provide the ability to have separate safe “universes” of message passing in MPI. A context is akin to an additional tag that differentiates messages. 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 communications on “other” communicators, and avoids the need to synchronize entry or exit into library code. Pending point-to-point communications are also guaranteed not to interfere with collective communications within a single communicator.
- **Groups** define the participants in the communication (see above) of a communicator.

- A **virtual topology** defines a special mapping of the ranks in a group to and from a topology. Special constructors for communicators are defined in chapter 6 to provide this feature. Intra-communicators as described in this chapter do not have topologies.
- **Attributes** define the local information that the user or library has added to a communicator for later reference.

Advice to users. The current practice in many communication libraries is that there is a unique, predefined communication universe that includes all processes available when the parallel program is initiated; the 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 processes. This practice can be followed in MPI by using the predefined communicator `MPI_COMM_WORLD`. *Users who are satisfied with this practice can plug in `MPI_COMM_WORLD` wherever a communicator argument is required, and can consequently disregard the rest of this chapter. (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 non-overlapping 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 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 in the local group is always a receive 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 communications on “other” communicators, and avoids the need to synchronize entry or exit into library code. There is no general-purpose collective communication on inter-communicators, so contexts are used just to isolate point-to-point communication.
- 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.

1 MPI provides mechanisms for creating and manipulating inter-communicators. They
2 are used for point-to-point communication in an related manner to intra-communicators.
3 Users who do not need inter-communication in their applications can safely ignore this
4 extension. Users who need collective operations via inter-communicators must layer it on
5 top of MPI. Users who require inter-communication between overlapping groups must also
6 layer this capability on top of MPI.

8 5.2 Basic Concepts

10 In this section, we turn to a more formal definition of the concepts introduced above.

12 5.2.1 Groups

14 A **group** is an ordered set of process identifiers (henceforth processes); processes are
15 implementation-dependent objects. Each process in a group is associated with an inte-
16 ger **rank**. Ranks are contiguous and start from zero. Groups are represented by opaque
17 **group objects**, and hence cannot be directly transferred from one process to another. A
18 group is used within a communicator to describe the participants in a communication “uni-
19 verse” and to rank such participants (thus giving them unique names within that “universe”
20 of communication).

21 There is a special pre-defined group: `MPI_GROUP_EMPTY`, which is a group with no
22 members. The predefined constant `MPI_GROUP_NULL` is the value used for invalid group
23 handles.

24
25 *Advice to users.* `MPI_GROUP_EMPTY`, which is a valid handle to an empty group,
26 should not be confused with `MPI_GROUP_NULL`, which in turn is an invalid handle.
27 The former may be used as an argument to group operations; the latter, which is
28 returned when a group is freed, is not a valid argument. (*End of advice to users.*)

29
30 *Advice to implementors.* A group may be represented by a virtual-to-real process-
31 address-translation table. Each communicator object (see below) would have a pointer
32 to such a table.

33 Simple implementations of MPI will enumerate groups, such as in a table. However,
34 more advanced data structures make sense in order to improve scalability and memory
35 usage with large numbers of processes. Such implementations are possible with MPI.
36 (*End of advice to implementors.*)

38 5.2.2 Contexts

39
40 A **context** is a property of communicators (defined next) that allows partitioning of the
41 communication space. A message sent in one context cannot be received in another context.
42 Furthermore, where permitted, collective operations are independent of pending point-to-
43 point operations. Contexts are not explicit MPI objects; they appear only as part of the
44 realization of communicators (below).

45
46 *Advice to implementors.* Distinct communicators in the same process have distinct
47 contexts. A context is essentially a system-managed tag (or tags) needed to make
48 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 don't 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 (which is strictly point-to-point 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.*)

5.2.3 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 6), communicators may also “cache” additional information (see section 5.7). 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 process. The source and destination of a message is identified by process rank within that group.

For collective communication, the intra-communicator specifies the set of 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 process addressing through ranks.

Intra-communicators are represented by opaque **intra-communicator objects**, and hence cannot be directly transferred from one process to another.

5.2.4 Predefined Intra-Communicators

An initial intra-communicator `MPI_COMM_WORLD` of all processes the local process can communicate with after initialization (itself included) is defined once `MPI_INIT` has been called. In addition, the communicator `MPI_COMM_SELF` is provided, which includes only the process itself.

The predefined constant `MPI_COMM_NULL` is the value used for invalid communicator handles.

In a static-process-model implementation of MPI, all processes that participate in the computation are available after MPI is initialized. For this case, `MPI_COMM_WORLD` is a communicator of all processes available for the computation; this communicator has the same value in all processes. In an implementation of MPI where processes can dynamically join an MPI execution, it may be the case that a process starts an MPI computation without having access to all other processes. In such situations, `MPI_COMM_WORLD` is a communicator incorporating all processes with which the joining process can immediately

communicate. Therefore, `MPI_COMM_WORLD` may simultaneously have different values in different processes.

All MPI implementations are required to provide the `MPI_COMM_WORLD` communicator. It cannot be deallocated during the life of a 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 process rank in `MPI_COMM_WORLD` and its (machine-dependent) absolute address. Neither does MPI specify the function of the host process, if any. Other implementation-dependent, predefined communicators may also be provided.

5.3 Group Management

This section describes the manipulation of process groups in MPI. These operations are local and their execution do not require interprocess communication.

5.3.1 Group Accessors

`MPI_GROUP_SIZE(group, size)`

IN	group	group (handle)
OUT	size	number of processes in the group (integer)

`int MPI_Group_size(MPI_Group group, int *size)`

`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 process in group, or <code>MPI_UNDEFINED</code> if the process is not a member (integer)

`int MPI_Group_rank(MPI_Group group, int *rank)`

`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 ranks 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_UNDE- FINED when no correspondence exists.	6

```

int MPI_Group_translate_ranks (MPI_Group group1, int n, int *ranks1,
                             MPI_Group group2, int *ranks2)
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 processes in two different groups. For instance, if one knows the ranks of certain processes in the group of MPI_COMM_WORLD, one might want to know their ranks in a subset of that group.

MPI_GROUP_COMPARE(group1, group2, result)			19
IN	group1	first group (handle)	20
IN	group2	second group (handle)	21
OUT	result	result (integer)	22

```

int MPI_Group_compare(MPI_Group group1, MPI_Group group2, int *result)
MPI_GROUP_COMPARE(GROUP1, GROUP2, RESULT, IERROR)
INTEGER GROUP1, GROUP2, RESULT, IERROR

```

MPI_IDENT results if the group members and group order is exactly the same in both groups. This happens for instance if group1 and group2 are the same handle. MPI_SIMILAR results if the group members are the same but the order is different. MPI_UNEQUAL results otherwise.

5.3.2 Group Constructors

Group constructors are used to subset and superset existing groups. These constructors construct new groups from existing groups. These are local operations, and distinct groups may be defined on different processes; a process may also define a group that does not include itself. Consistent definitions are required when groups are used as arguments in communicator-building functions. MPI does not provide a mechanism to build a group from scratch, but only from other, previously defined groups. The base group, upon which all other groups are defined, is the group associated with the initial communicator MPI_COMM_WORLD (accessible through the function MPI_COMM_GROUP).

Rationale. In what follows, there is no group duplication function analogous to MPI_COMM_DUP, defined later in this chapter. There is no need for a group duplicator. A group, once created, can have several references to it by making copies of

1 the handle. The following constructors address the need for subsets and supersets of
 2 existing groups. (*End of rationale.*)

3
 4 *Advice to implementors.* Each group constructor behaves as if it returned a new
 5 group object. When this new group is a copy of an existing group, then one can
 6 avoid creating such new objects, using a reference-count mechanism. (*End of advice
 7 to implementors.*)

8
 9
 10 MPI_COMM_GROUP(comm, group)

11
 12 IN comm communicator (handle)
 13 OUT group group corresponding to comm (handle)

14
 15 int MPI_Comm_group(MPI_Comm comm, MPI_Group *group)

16
 17 MPI_COMM_GROUP(COMM, GROUP, IERROR)

18 INTEGER COMM, GROUP, IERROR

19 MPI_COMM_GROUP returns in group a handle to the group of comm.

20
 21
 22 MPI_GROUP_UNION(group1, group2, newgroup)

23 IN group1 first group (handle)
 24 IN group2 second group (handle)
 25 OUT newgroup union group (handle)

26
 27
 28 int MPI_Group_union(MPI_Group group1, MPI_Group group2, MPI_Group *newgroup)

29
 30 MPI_GROUP_UNION(GROUP1, GROUP2, NEWGROUP, IERROR)

31 INTEGER GROUP1, GROUP2, NEWGROUP, IERROR

32
 33
 34 MPI_GROUP_INTERSECTION(group1, group2, newgroup)

35 IN group1 first group (handle)
 36 IN group2 second group (handle)
 37 OUT newgroup intersection group (handle)

38
 39
 40 int MPI_Group_intersection(MPI_Group group1, MPI_Group group2,
 41 MPI_Group *newgroup)

42
 43 MPI_GROUP_INTERSECTION(GROUP1, GROUP2, NEWGROUP, IERROR)

44 INTEGER GROUP1, GROUP2, NEWGROUP, IERROR

45
 46
 47
 48

```

MPI_GROUP_DIFFERENCE(group1, group2, newgroup) 1
    IN      group1          first group (handle) 2
    IN      group2          second group (handle) 3
    OUT     newgroup        difference group (handle) 4
                                                    5
                                                    6
int MPI_Group_difference(MPI_Group group1, MPI_Group group2, 7
                        MPI_Group *newgroup) 8
                                                    9

```

```

MPI_GROUP_DIFFERENCE(GROUP1, GROUP2, NEWGROUP, IERROR) 10
    INTEGER GROUP1, GROUP2, NEWGROUP, IERROR 11

```

The set-like operations are defined as follows: 12

union All elements of the first group (*group1*), followed by all elements of second group (*group2*) not in first. 14

intersect all elements of the first group that are also in the second group, ordered as in first group. 17

difference all elements of the first group that are not in the second group, ordered as in the first group. 19

Note that for these operations the order of 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. 22

The new group can be empty, that is, equal to MPI_GROUP_EMPTY. 25

```

MPI_GROUP_INCL(group, n, ranks, newgroup) 28
    IN      group          group (handle) 29
    IN      n              number of elements in array ranks (and size of new- 30
                        group) (integer) 31
    IN      ranks          ranks of processes in group to appear in newgroup (ar- 32
                        ray of integers) 33
    OUT     newgroup        new group derived from above, in the order defined by 34
                        ranks (handle) 35
                                                    36
                                                    37

```

```

int MPI_Group_incl(MPI_Group group, int n, int *ranks, MPI_Group *newgroup) 38
MPI_GROUP_INCL(GROUP, N, RANKS, NEWGROUP, IERROR) 39
    INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR 40

```

The function MPI_GROUP_INCL creates a group *newgroup* that consists of the *n* processes in *group* with ranks *rank[0]*, ..., *rank[n-1]*; the process with rank *i* in *newgroup* is the 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. 42

```

1 MPI_GROUP_EXCL(group, n, ranks, newgroup)
2     IN      group          group (handle)
3
4     IN      n              number of elements in array ranks (integer)
5
6     IN      ranks          array of integer ranks in group not to appear in new-
7                          group
8
9     OUT     newgroup       new group derived from above, preserving the order
10                          defined by group (handle)

```

```

11 int MPI_Group_excl(MPI_Group group, int n, int *ranks, MPI_Group *newgroup)

```

```

12 MPI_GROUP_EXCL(GROUP, N, RANKS, NEWGROUP, IERROR)
13     INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR

```

14 The function `MPI_GROUP_EXCL` creates a group of processes `newgroup` that is obtained
15 by deleting from `group` those processes with ranks `ranks[0]` ,... `ranks[n-1]`. The ordering of
16 processes in `newgroup` is identical to the ordering in `group`. Each of the `n` elements of `ranks`
17 must be a valid rank in `group` and all elements must be distinct; otherwise, the program is
18 erroneous. If `n = 0`, then `newgroup` is identical to `group`.

```

21 MPI_GROUP_RANGE_INCL(group, n, ranges, newgroup)

```

```

22     IN      group          group (handle)
23
24     IN      n              number of triplets in array ranges (integer)
25
26     IN      ranges          a one-dimensional array of integer triplets, of the form
27                          (first rank, last rank, stride) indicating ranks in group
28                          of processes to be included in newgroup
29
30     OUT     newgroup       new group derived from above, in the order defined by
31                          ranges (handle)

```

```

32 int MPI_Group_range_incl(MPI_Group group, int n, int ranges[] [3],
33                         MPI_Group *newgroup)

```

```

34 MPI_GROUP_RANGE_INCL(GROUP, N, RANGES, NEWGROUP, IERROR)
35     INTEGER GROUP, N, RANGES(3,*), NEWGROUP, IERROR

```

36 If `ranges` consist of the triplets

```

37     (first1, last1, stride1), ..., (firstn, lastn, striden)

```

38 then `newgroup` consists of the sequence of processes in `group` with ranks

```

39     first1, first1 + stride1, ..., first1 +  $\left\lfloor \frac{\textit{last}_1 - \textit{first}_1}{\textit{stride}_1} \right\rfloor \textit{stride}_1, \dots$ 

```

```

40     firstn, firstn + striden, ..., firstn +  $\left\lfloor \frac{\textit{last}_n - \textit{first}_n}{\textit{stride}_n} \right\rfloor \textit{stride}_n$ .

```

41 Each computed rank must be a valid rank in `group` and all computed ranks must be
42 distinct, or else the program is erroneous. Note that we may have $\textit{first}_i > \textit{last}_i$, and \textit{stride}_i
43 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)`

IN	group	group (handle)
IN	n	number of elements in array ranges (integer)
IN	ranges	a one-dimensional array of integer triplets of the form (first rank, last rank, stride), indicating the ranks in group of processes to be excluded from the output group <code>newgroup</code> .
OUT	newgroup	new group derived from above, preserving the order in group (handle)

```
int MPI_Group_range_excl(MPI_Group group, int n, int ranges[][3],
                        MPI_Group *newgroup)
```

```
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.*)

5.3.3 Group Destructors

`MPI_GROUP_FREE(group)`

INOUT	group	group (handle)
-------	-------	----------------

```
int MPI_Group_free(MPI_Group *group)
```



```

1 MPI_GROUP_FREE(GROUP, IERROR)
2     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_CREATE` and `MPI_COMM_DUP`, 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.*)

5.4 Communicator Management

This section describes the manipulation of communicators in MPI. Operations that access communicators are local and their execution does not require interprocess communication. Operations that create communicators are collective and may require interprocess communication.

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.*)

5.4.1 Communicator Accessors

The following are all local operations.

```

29 MPI_COMM_SIZE(comm, size)

```

IN	comm	communicator (handle)
OUT	size	number of processes in the group of comm (integer)

```

34 int MPI_Comm_size(MPI_Comm comm, int *size)

```

```

35 MPI_COMM_SIZE(COMM, SIZE, IERROR)
36     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 function is so commonly used, that this shortcut was introduced. (*End of rationale.*)

Advice to users. This function indicates the number of processes involved in a communicator. For `MPI_COMM_WORLD`, it indicates the total number of processes available (for this version of MPI, there is no standard way to change the number of processes once initialization has taken place).

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 process that calls it in the range from $0 \dots \text{size}-1$, 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 process in group of <code>comm</code> (integer)

`int MPI_Comm_rank(MPI_Comm comm, int *rank)`

`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 function is so commonly used, that this shortcut was introduced. (*End of rationale.*)

Advice to users. This function gives the rank of the process in the particular communicator's group. It is useful, as noted above, in conjunction with `MPI_COMM_SIZE`.

Many programs will be written with the master-slave model, where one process (such as the rank-zero process) will play a supervisory role, and the other processes will serve as compute nodes. In this framework, the two preceding calls are useful for determining the roles of the various processes of a communicator. (*End of advice to users.*)

`MPI_COMM_COMPARE(comm1, comm2, result)`

IN	<code>comm1</code>	first communicator (handle)
IN	<code>comm2</code>	second communicator (handle)
OUT	<code>result</code>	result (integer)

`int MPI_Comm_compare(MPI_Comm comm1, MPI_Comm comm2, int *result)`

`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.

5.4.2 Communicator Constructors

The following are collective functions that are invoked by all processes in the group associated with `comm`.

Rationale. Note that there is a chicken-and-egg aspect to MPI in that a communicator is needed to create a new communicator. The base communicator for all MPI communicators is predefined outside of MPI, and is `MPI_COMM_WORLD`. This model was arrived at after considerable debate, and was chosen to increase “safety” of programs written in MPI. (*End of rationale.*)

```
MPI_COMM_DUP(comm, newcomm)
```

IN	<code>comm</code>	communicator (handle)
OUT	<code>newcomm</code>	copy of <code>comm</code> (handle)

```
int MPI_Comm_dup(MPI_Comm comm, MPI_Comm *newcomm)
```

```
MPI_COMM_DUP(COMM, NEWCOMM, IERROR)
INTEGER COMM, NEWCOMM, IERROR
```

`MPI_COMM_DUP` Duplicates the existing communicator `comm` with associated key values. For each key value, the respective copy callback function determines the attribute value associated with this key in the new communicator; one particular action that a copy callback may take is to delete the attribute from the new communicator. Returns in `newcomm` a new communicator with the same group, any copied cached information, but a new context (see section 5.7.1).

Advice to users. This operation is used to provide a parallel library call with a duplicate communication space that has the same properties as the original communicator. This includes any attributes (see below), and topologies (see chapter 6). This call is valid even if there are pending point-to-point communications involving the communicator `comm`. A typical call might involve a `MPI_COMM_DUP` at the beginning of the parallel call, and an `MPI_COMM_FREE` of that duplicated communicator at the end of the call. Other models of communicator management are also possible.

This call applies to both intra- and inter-communicators. (*End of advice to users.*)

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_CREATE(comm, group, newcomm)

IN	comm	communicator (handle)
IN	group	Group, which is a subset of the group of comm (handle)
OUT	newcomm	new communicator (handle)

```
int MPI_Comm_create(MPI_Comm comm, MPI_Group group, MPI_Comm *newcomm)
```

```
MPI_COMM_CREATE(COMM, GROUP, NEWCOMM, IERROR)
```

```
INTEGER COMM, GROUP, NEWCOMM, IERROR
```

This function creates a new communicator `newcomm` with communication group defined by `group` and a new context. No cached information propagates from `comm` to `newcomm`. The function returns `MPI_COMM_NULL` to processes that are not in `group`. The call is erroneous if not all `group` arguments have the same value, or if `group` is not a subset of the group associated with `comm`. Note that the call is to be executed by all processes in `comm`, even if they do not belong to the new group. This call applies only to intra-communicators.

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.
- It permits implementations sometimes to avoid communication related to context creation.

(End of rationale.)

Advice to users. `MPI_COMM_CREATE` provides a means to subset a group of 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) further to subdivide a computation into parallel sub-computations. A more general service is provided by `MPI_COMM_SPLIT`, below. *(End of advice to users.)*

Advice to implementors. Since all processes calling `MPI_COMM_DUP` or `MPI_COMM_CREATE` provide the same `group` argument, 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.

Important: If new communicators are created without synchronizing the processes involved then the communication system should be able to cope with messages arriving in a context that has not yet been allocated at the receiving process. *(End of advice to implementors.)*

```
1 MPI_COMM_SPLIT(comm, color, key, newcomm)
```

```
2     IN      comm      communicator (handle)
```

```
3     IN      color     control of subset assignment (integer)
```

```
4     IN      key       control of rank assignment (integer)
```

```
5     OUT     newcomm   new communicator (handle)
```

```
6
7
8
9 int MPI_Comm_split(MPI_Comm comm, int color, int key, MPI_Comm *newcomm)
```

```
10 MPI_COMM_SPLIT(COMM, COLOR, KEY, NEWCOMM, IERROR)
```

```
11     INTEGER COMM, COLOR, KEY, NEWCOMM, IERROR
```

12
13 This function partitions the group associated with `comm` into disjoint subgroups, one for
14 each value of `color`. Each subgroup contains all processes of the same color. Within each
15 subgroup, the processes are ranked in the order defined by the value of the argument
16 `key`, with ties broken according to their rank in the old group. A new communicator is
17 created for each subgroup and returned in `newcomm`. A process may supply the color value
18 `MPI_UNDEFINED`, in which case `newcomm` returns `MPI_COMM_NULL`. This is a collective
19 call, but each process is permitted to provide different values for `color` and `key`.

20 A call to `MPI_COMM_CREATE(comm, group, newcomm)` is equivalent to
21 a call to `MPI_COMM_SPLIT(comm, color, key, newcomm)`, where all members of `group` pro-
22 vide `color = 0` and `key = rank` in `group`, and all processes that are not members of
23 `group` provide `color = MPI_UNDEFINED`. The function `MPI_COMM_SPLIT` allows more
24 general partitioning of a group into one or more subgroups with optional reordering. This
25 call applies only intra-communicators.

26 The value of `color` must be nonnegative.

27
28 *Advice to users.* This is an extremely powerful mechanism for dividing a single com-
29 municating group of processes into k subgroups, with k chosen implicitly by the user
30 (by the number of colors asserted over all the processes). Each resulting communica-
31 tor will be non-overlapping. Such a division could be useful for defining a hierarchy
32 of computations, such as for multigrid, or linear algebra.

33 Multiple calls to `MPI_COMM_SPLIT` can be used to overcome the requirement that
34 any call have no overlap of the resulting communicators (each process is of only one
35 color per call). In this way, multiple overlapping communication structures can be
36 created. Creative use of the `color` and `key` in such splitting operations is encouraged.

37 Note that, for a fixed `color`, the keys need not be unique. It is `MPI_COMM_SPLIT`'s
38 responsibility to sort processes in ascending order according to this key, and to break
39 ties in a consistent way. If all the keys are specified in the same way, then all the
40 processes in a given color will have the relative rank order as they did in their parent
41 group. (In general, they will have different ranks.)

42 Essentially, making the key value zero for all processes of a given color means that one
43 doesn't really care about the rank-order of the processes in the new communicator.
44 (*End of advice to users.*)

45
46 *Rationale.* `color` is restricted to be nonnegative, so as not to conflict with the value
47 assigned to `MPI_UNDEFINED`. (*End of rationale.*)
48

5.4.3 Communicator Destructors

MPI_COMM_FREE(comm)

INOUT comm communicator to be destroyed (handle)

```
int MPI_Comm_free(MPI_Comm *comm)
```

```
MPI_COMM_FREE(COMM, IERROR)
```

```
INTEGER COMM, IERROR
```

This collective operation marks the communication object for deallocation. The handle is set to MPI_COMM_NULL. Any pending operations that use this communicator will complete normally; the object is actually deallocated only if there are no other active references to it. This call applies to intra- and inter-communicators. The delete callback functions for all cached attributes (see section 5.7) are called in arbitrary order.

Advice to implementors. A reference-count mechanism may be used: the reference count is incremented by each call to MPI_COMM_DUP, and decremented by each call to MPI_COMM_FREE. The object is ultimately deallocated when the count reaches zero.

Though collective, it is anticipated that this operation will normally be implemented to be local, though the debugging version of an MPI library might choose to synchronize. (*End of advice to implementors.*)

5.5 Motivating Examples

5.5.1 Current Practice #1

Example #1a:

```
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 ("Process %d size %d\n", me, size);
    ...
    MPI_Finalize();
}
```

Example #1a is a do-nothing program that initializes itself legally, and refers to the “all” communicator, and prints a message. It terminates itself legally too. This example does not imply that MPI supports printf-like communication itself.

Example #1b (supposing that size is even):

```

1   main(int argc, char **argv)
2   {
3       int me, size;
4       int SOME_TAG = 0;
5       ...
6       MPI_Init(&argc, &argv);
7
8       MPI_Comm_rank(MPI_COMM_WORLD, &me); /* local */
9       MPI_Comm_size(MPI_COMM_WORLD, &size); /* local */
10
11      if((me % 2) == 0)
12      {
13          /* send unless highest-numbered process */
14          if((me + 1) < size)
15              MPI_Send(..., me + 1, SOME_TAG, MPI_COMM_WORLD);
16      }
17      else
18          MPI_Recv(..., me - 1, SOME_TAG, MPI_COMM_WORLD);
19
20      ...
21      MPI_Finalize();
22  }

```

Example #1b schematically illustrates message exchanges between “even” and “odd” processes in the “all” communicator.

5.5.2 Current Practice #2

```

28  main(int argc, char **argv)
29  {
30      int me, count;
31      void *data;
32      ...
33
34      MPI_Init(&argc, &argv);
35      MPI_Comm_rank(MPI_COMM_WORLD, &me);
36
37      if(me == 0)
38      {
39          /* get input, create buffer “data” */
40          ...
41      }
42
43      MPI_Bcast(data, count, MPI_BYTE, 0, MPI_COMM_WORLD);
44
45      ...
46
47      MPI_Finalize();
48  }

```

This example illustrates the use of a collective communication.

5.5.3 (Approximate) Current Practice #3

```

main(int argc, char **argv)
{
    int me, count, count2;
    void *send_buf, *recv_buf, *send_buf2, *recv_buf2;
    MPI_Group MPI_GROUP_WORLD, grprem;
    MPI_Comm commslave;
    static int ranks[] = {0};
    ...
    MPI_Init(&argc, &argv);
    MPI_Comm_group(MPI_COMM_WORLD, &MPI_GROUP_WORLD);
    MPI_Comm_rank(MPI_COMM_WORLD, &me); /* local */

    MPI_Group_excl(MPI_GROUP_WORLD, 1, ranks, &grprem); /* local */
    MPI_Comm_create(MPI_COMM_WORLD, grprem, &commslave);

    if(me != 0)
    {
        /* compute on slave */
        ...
        MPI_Reduce(send_buf,recv_buff,count, MPI_INT, MPI_SUM, 1, commslave);
        ...
    }
    /* zero falls through immediately to this reduce, others do later... */
    MPI_Reduce(send_buf2, recv_buff2, count2,
               MPI_INT, MPI_SUM, 0, MPI_COMM_WORLD);

    MPI_Comm_free(&commslave);
    MPI_Group_free(&MPI_GROUP_WORLD);
    MPI_Group_free(&grprem);
    MPI_Finalize();
}

```

This example illustrates how a group consisting of all but the zeroth process of the “all” group is created, and then how a communicator is formed (`commslave`) for that new group. The new communicator is used in a collective call, and all 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 `commslave`, and vice versa.

In summary, “group safety” is achieved via communicators because distinct contexts within communicators are enforced to be unique on any process.

5.5.4 Example #4

The following example 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

1 collective communication.

```

2
3     #define TAG_ARBITRARY 12345
4     #define SOME_COUNT    50
5
6     main(int argc, char **argv)
7     {
8         int me;
9         MPI_Request request[2];
10        MPI_Status status[2];
11        MPI_Group MPI_GROUP_WORLD, subgroup;
12        int ranks[] = {2, 4, 6, 8};
13        MPI_Comm the_comm;
14        ...
15        MPI_Init(&argc, &argv);
16        MPI_Comm_group(MPI_COMM_WORLD, &MPI_GROUP_WORLD);
17
18        MPI_Group_incl(MPI_GROUP_WORLD, 4, ranks, &subgroup); /* local */
19        MPI_Group_rank(subgroup, &me); /* local */
20
21        MPI_Comm_create(MPI_COMM_WORLD, subgroup, &the_comm);
22
23        if(me != MPI_UNDEFINED)
24        {
25            MPI_Irecv(buff1, count, MPI_DOUBLE, MPI_ANY_SOURCE, TAG_ARBITRARY,
26                    the_comm, request);
27            MPI_Isend(buff2, count, MPI_DOUBLE, (me+1)%4, TAG_ARBITRARY,
28                    the_comm, request+1);
29        }
30
31        for(i = 0; i < SOME_COUNT, i++)
32            MPI_Reduce(..., the_comm);
33        MPI_Waitall(2, request, status);
34
35        MPI_Comm_free(&the_comm);
36        MPI_Group_free(&MPI_GROUP_WORLD);
37        MPI_Group_free(&subgroup);
38        MPI_Finalize();
39    }

```

41 5.5.5 Library Example #1

42 The main program:

```

43
44     main(int argc, char **argv)
45     {
46         int done = 0;
47         user_lib_t *libh_a, *libh_b;
48         void *dataset1, *dataset2;

```

```

...
MPI_Init(&argc, &argv);
...
init_user_lib(MPI_COMM_WORLD, &libh_a);
init_user_lib(MPI_COMM_WORLD, &libh_b);
...
user_start_op(libh_a, dataset1);
user_start_op(libh_b, dataset2);
...
while(!done)
{
    /* work */
    ...
    MPI_Reduce(..., MPI_COMM_WORLD);
    ...
    /* see if done */
    ...
}
user_end_op(libh_a);
user_end_op(libh_b);

uninit_user_lib(libh_a);
uninit_user_lib(libh_b);
MPI_Finalize();
}

```

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:

```

1 void user_end_op(user_lib_t *handle)
2 {
3     MPI_Status *status;
4     MPI_Wait(handle -> isend_handle, status);
5     MPI_Wait(handle -> irecv_handle, status);
6 }

```

7
8 User object clean-up code:

```

9 void uninit_user_lib(user_lib_t *handle)
10 {
11     MPI_Comm_free(&(handle -> comm));
12     free(handle);
13 }
14

```

15 5.5.6 Library Example #2

16
17 The main program:

```

18 main(int argc, char **argv)
19 {
20     int ma, mb;
21     MPI_Group MPI_GROUP_WORLD, group_a, group_b;
22     MPI_Comm comm_a, comm_b;
23
24     static int list_a[] = {0, 1};
25 #if defined(EXAMPLE_2B) | defined(EXAMPLE_2C)
26     static int list_b[] = {0, 2 ,3};
27 #else/* EXAMPLE_2A */
28     static int list_b[] = {0, 2};
29 #endif
30
31     int size_list_a = sizeof(list_a)/sizeof(int);
32     int size_list_b = sizeof(list_b)/sizeof(int);
33
34     ...
35     MPI_Init(&argc, &argv);
36     MPI_Comm_group(MPI_COMM_WORLD, &MPI_GROUP_WORLD);
37
38     MPI_Group_incl(MPI_GROUP_WORLD, size_list_a, list_a, &group_a);
39     MPI_Group_incl(MPI_GROUP_WORLD, size_list_b, list_b, &group_b);
40
41     MPI_Comm_create(MPI_COMM_WORLD, group_a, &comm_a);
42     MPI_Comm_create(MPI_COMM_WORLD, group_b, &comm_b);
43
44     if(comm_a != MPI_COMM_NULL)
45         MPI_Comm_rank(comm_a, &ma);
46     if(comm_b != MPI_COMM_NULL)
47         MPI_Comm_rank(comm_b, &mb);
48

```

```

if(comm_a != MPI_COMM_NULL) 1
    lib_call(comm_a); 2
3
if(comm_b != MPI_COMM_NULL) 4
{ 5
    lib_call(comm_b); 6
    lib_call(comm_b); 7
} 8
9
if(comm_a != MPI_COMM_NULL) 10
    MPI_Comm_free(&comm_a); 11
if(comm_b != MPI_COMM_NULL) 12
    MPI_Comm_free(&comm_b); 13
MPI_Group_free(&group_a); 14
MPI_Group_free(&group_b); 15
MPI_Group_free(&MPI_GROUP_WORLD); 16
MPI_Finalize(); 17
} 18
19
The library: 20
void lib_call(MPI_Comm comm) 21
{ 22
    int me, done = 0; 23
    MPI_Comm_rank(comm, &me); 24
    if(me == 0) 25
        while(!done) 26
        { 27
            MPI_Recv(..., MPI_ANY_SOURCE, MPI_ANY_TAG, comm); 28
            ... 29
        } 30
    else 31
    { 32
        /* work */ 33
        MPI_Send(..., 0, ARBITRARY_TAG, comm); 34
        .... 35
    } 36
#ifdef EXAMPLE_2C 37
    /* include (resp, exclude) for safety (resp, no safety): */ 38
    MPI_Barrier(comm); 39
#endif 40
} 41

```

The above example is really three examples, depending on whether or not one includes rank 3 in list_b, and whether or not a synchronize 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 the 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 synchronize is not needed to get safety from back masking.

1 Algorithms like “reduce” and “allreduce” have strong enough source selectivity prop-
2 erties so that they are inherently okay (no backmasking), provided that MPI provides basic
3 guarantees. So are multiple calls to a typical tree-broadcast algorithm with the same root
4 or different roots (see [28]). Here we rely on two guarantees of MPI: pairwise ordering of
5 messages between processes in the same context, and source selectivity — deleting either
6 feature removes the guarantee that backmasking cannot be required.

7 Algorithms that try to do non-deterministic broadcasts or other calls that include wild-
8 card operations will not generally have the good properties of the deterministic implemen-
9 tations of “reduce,” “allreduce,” and “broadcast.” Such algorithms would have to utilize
10 the monotonically increasing tags (within a communicator scope) to keep things straight.

11 All of the foregoing is a supposition of “collective calls” implemented with point-to-
12 point operations. MPI implementations may or may not implement collective calls using
13 point-to-point operations. These algorithms are used to illustrate the issues of correctness
14 and safety, independent of how MPI implements its collective calls. See also section 5.8.

16 5.6 Inter-Communication

18 This section introduces the concept of inter-communication and describes the portions of
19 MPI that support it. It describes support for writing programs that contain user-level
20 servers.

21 All point-to-point communication described thus far has involved communication be-
22 tween processes that are members of the same group. This type of communication is called
23 “intra-communication” and the communicator used is called an “intra-communicator,” as
24 we have noted earlier in the chapter.

25 In modular and multi-disciplinary applications, different process groups execute distinct
26 modules and processes within different modules communicate with one another in a pipeline
27 or a more general module graph. In these applications, the most natural way for a process
28 to specify a target process is by the rank of the target process within the target group. In
29 applications that contain internal user-level servers, each server may be a process group that
30 provides services to one or more clients, and each client may be a process group that uses
31 the services of one or more servers. It is again most natural to specify the target process
32 by rank within the target group in these applications. This type of communication is called
33 “inter-communication” and the communicator used is called an “inter-communicator,” as
34 introduced earlier.

35 An inter-communication is a point-to-point communication between processes in differ-
36 ent groups. The group containing a process that initiates an inter-communication operation
37 is called the “local group,” that is, the sender in a send and the receiver in a receive. The
38 group containing the target process is called the “remote group,” that is, the receiver in a
39 send and the sender in a receive. As in intra-communication, the target process is specified
40 using a (communicator, rank) pair. Unlike intra-communication, the rank is relative to a
41 second, remote group.

42 All inter-communicator constructors are blocking and require that the local and remote
43 groups be disjoint.

44
45 *Advice to users.* The groups must be disjoint for several reasons. Primarily, this is the
46 intent of the intercommunicators — to provide a communicator for communication be-
47 tween disjoint groups. This is reflected in the definition of `MPI_INTERCOMM_MERGE`,
48

which allows the user to control the ranking of the processes in the created intracommunicator; this ranking makes little sense if the groups are not disjoint. In addition, the natural extension of collective operations to intercommunicators makes the most sense when the groups are disjoint. (*End of advice to users.*)

Here is a summary of the properties of inter-communication and inter-communicators:

- The syntax of point-to-point communication is the same for both inter- and intracommunication. The same communicator can be used both for send and for receive operations.
- A target process is addressed by its rank in the remote group, both for sends and for receives.
- Communications using an inter-communicator are guaranteed not to conflict with any communications that use a different communicator.
- An inter-communicator cannot be used for collective communication.
- A communicator will provide either intra- or inter-communication, never both.

The routine `MPI_COMM_TEST_INTER` may be used to determine if a communicator is an inter- or intra-communicator. Inter-communicators can be used as arguments to some of the other communicator access routines. Inter-communicators cannot be used as input to some of the constructor routines for intra-communicators (for instance, `MPI_COMM_CREATE`).

Advice to implementors. For the purpose of point-to-point communication, communicators can be represented in each process by a tuple consisting of:

group
send_context
receive_context
source

For inter-communicators, **group** describes the remote group, and **source** is the rank of the process in the local group. For intra-communicators, **group** is the communicator group (remote=local), **source** is the rank of the process in this group, and **send context** and **receive context** are identical. A group is represented by a rank-to-absolute-address translation table.

The inter-communicator cannot be discussed sensibly without considering processes in both the local and remote groups. Imagine a process **P** in group \mathcal{P} , which has an inter-communicator $\mathbf{C}_{\mathcal{P}}$, and a process **Q** in group \mathcal{Q} , which has an inter-communicator $\mathbf{C}_{\mathcal{Q}}$. Then

- $\mathbf{C}_{\mathcal{P}}.\mathbf{group}$ describes the group \mathcal{Q} and $\mathbf{C}_{\mathcal{Q}}.\mathbf{group}$ describes the group \mathcal{P} .
- $\mathbf{C}_{\mathcal{P}}.\mathbf{send_context} = \mathbf{C}_{\mathcal{Q}}.\mathbf{receive_context}$ and the context is unique in \mathcal{Q} ;
 $\mathbf{C}_{\mathcal{P}}.\mathbf{receive_context} = \mathbf{C}_{\mathcal{Q}}.\mathbf{send_context}$ and this context is unique in \mathcal{P} .
- $\mathbf{C}_{\mathcal{P}}.\mathbf{source}$ is rank of **P** in \mathcal{P} and $\mathbf{C}_{\mathcal{Q}}.\mathbf{source}$ is rank of **Q** in \mathcal{Q} .

1 Assume that **P** sends a message to **Q** using the inter-communicator. Then **P** uses
 2 the **group** table to find the absolute address of **Q**; **source** and **send_context** are
 3 appended to the message.

4 Assume that **Q** posts a receive with an explicit source argument using the inter-
 5 communicator. Then **Q** matches **receive_context** to the message context and source
 6 argument to the message source.
 7

8 The same algorithm is appropriate for intra-communicators as well.

9 In order to support inter-communicator accessors and constructors, it is necessary to
 10 supplement this model with additional structures, that store information about the
 11 local communication group, and additional safe contexts. (*End of advice to imple-*
 12 *mentors.*)
 13

14 5.6.1 Inter-communicator Accessors

17 MPI_COMM_TEST_INTER(comm, flag)

18
 19 IN comm communicator (handle)
 20 OUT flag (logical)
 21

22 int MPI_Comm_test_inter(MPI_Comm comm, int *flag)

23
 24 MPI_COMM_TEST_INTER(COMM, FLAG, IERROR)
 25 INTEGER COMM, IERROR
 26 LOGICAL FLAG
 27

28 This local routine allows the calling process to determine if a communicator is an inter-
 29 communicator or an intra-communicator. It returns true if it is an inter-communicator,
 30 otherwise false.

31 When an inter-communicator is used as an input argument to the communicator ac-
 32 cessors described above under intra-communication, the following table describes behavior.
 33

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

39
 40 Furthermore, the operation MPI_COMM_COMPARE is valid for inter-communicators. Both
 41 communicators must be either intra- or inter-communicators, or else MPI_UNEQUAL results.
 42 Both corresponding local and remote groups must compare correctly to get the results
 43 MPI_CONGRUENT and MPI_SIMILAR. In particular, it is possible for MPI_SIMILAR to result
 44 because either the local or remote groups were similar but not identical.

45 The following accessors provide consistent access to the remote group of an inter-
 46 communicator:

47 The following are all local operations.
 48

```

MPI_COMM_REMOTE_SIZE(comm, size) 1
    IN      comm                    inter-communicator (handle) 2
    OUT     size                    number of processes in the remote group of comm 3
                                      (integer) 4
                                      5
                                      6
int MPI_Comm_remote_size(MPI_Comm comm, int *size) 7
MPI_COMM_REMOTE_SIZE(COMM, SIZE, IERROR) 8
    INTEGER COMM, SIZE, IERROR 9
                                      10
                                      11
MPI_COMM_REMOTE_GROUP(comm, group) 12
    IN      comm                    inter-communicator (handle) 13
    OUT     group                   remote group corresponding to comm (handle) 14
                                      15
                                      16
int MPI_Comm_remote_group(MPI_Comm comm, MPI_Group *group) 17
MPI_COMM_REMOTE_GROUP(COMM, GROUP, IERROR) 18
    INTEGER COMM, GROUP, IERROR 19
                                      20
                                      21

```

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.*)

5.6.2 Inter-communicator Operations

This section introduces four blocking inter-communicator operations. `MPI_INTERCOMM_CREATE` is used to bind two intra-communicators into an inter-communicator; 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. (If a process is multithreaded, and MPI calls block only a thread, rather than a process, then “dual membership” can be supported. It is then the user’s responsibility to make sure that calls on behalf of the two “roles” of a process are executed by two independent threads.)

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 a process in the local group and a process in the remote group.

In standard MPI implementations (with static process allocation at initialization), the `MPI_COMM_WORLD` communicator (or preferably a dedicated duplicate thereof) can be

1 this peer communicator. In dynamic MPI implementations, where, for example, a process
 2 may spawn new child processes during an MPI execution, the parent process may be the
 3 “bridge” between the old communication universe and the new communication world that
 4 includes the parent and its children.

5 The application topology functions described in chapter 6 do not apply to inter-
 6 communicators. Users that require this capability should utilize `MPI_INTERCOMM_MERGE`
 7 to build an intra-communicator, then apply the graph or cartesian topology capabilities to
 8 that intra-communicator, creating an appropriate topology-oriented intra-communicator.
 9 Alternatively, it may be reasonable to devise one’s own application topology mechanisms
 10 for this case, without loss of generality.

11
 12
 13 `MPI_INTERCOMM_CREATE(local_comm, local_leader, peer_comm, remote_leader, tag,`
 14 `newintercomm)`

15	IN	<code>local_comm</code>	local intra-communicator (handle)
16	IN	<code>local_leader</code>	rank of local group leader in <code>local_comm</code> (integer)
17			
18	IN	<code>peer_comm</code>	“peer” communicator; significant only at the <code>local_leader</code> 19 (handle)
20	IN	<code>remote_leader</code>	rank of remote group leader in <code>peer_comm</code> ; significant 21 only at the <code>local_leader</code> (integer)
22			
23	IN	<code>tag</code>	“safe” tag (integer)
24	OUT	<code>newintercomm</code>	new inter-communicator (handle)

25
 26 `int MPI_Intercomm_create(MPI_Comm local_comm, int local_leader,`
 27 `MPI_Comm peer_comm, int remote_leader, int tag,`
 28 `MPI_Comm *newintercomm)`

29 `MPI_INTERCOMM_CREATE(LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG,`
 30 `NEWINTERCOMM, IERROR)`
 31 `INTEGER LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG,`
 32 `NEWINTERCOMM, IERROR`
 33

34 This call creates an inter-communicator. It is collective over the union of the local and
 35 remote groups. Processes should provide identical `local_comm` and `local_leader` arguments
 36 within each group. Wildcards are not permitted for `remote_leader`, `local_leader`, and `tag`.

37 This call uses point-to-point communication with communicator `peer_comm`, and with
 38 tag `tag` between the leaders. Thus, care must be taken that there be no pending communi-
 39 cation on `peer_comm` that could interfere with this communication.

40
 41 *Advice to users.* We recommend using a dedicated peer communicator, such as a
 42 duplicate of `MPI_COMM_WORLD`, to avoid trouble with peer communicators. (*End*
 43 *of advice to users.*)
 44
 45
 46
 47
 48




Figure 5.1: Three-group pipeline.

```
MPI_INTERCOMM_MERGE(intercomm, high, newintracomm)
```

IN	intercomm	Inter-Communicator (handle)
IN	high	(logical)
OUT	newintracomm	new intra-communicator (handle)

```
int MPI_Intercomm_merge(MPI_Comm intercomm, int high,
                        MPI_Comm *newintracomm)
```

```
MPI_INTERCOMM_MERGE(INTERCOMM, HIGH, INTRACOMM, IERROR)
INTEGER INTERCOMM, INTRACOMM, IERROR
LOGICAL HIGH
```

This function creates an intra-communicator from the union of the two groups that are associated with `intercomm`. All processes should provide the same `high` value within each of the two groups. If processes in one group provided the value `high = false` and processes in the other group provided the value `high = true` then the union orders the “low” group before the “high” group. If all 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 intercommunicator in each process is inherited from the communicator that contributes the local group. Note that this can result in different 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.*)

5.6.3 Inter-Communication Examples

Example 1: Three-Group “Pipeline”

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.

```
main(int argc, char **argv)
{
    MPI_Comm    myComm;        /* intra-communicator of local sub-group */
```

```
1     MPI_Comm  myFirstComm; /* inter-communicator */
2     MPI_Comm  mySecondComm; /* second inter-communicator (group 1 only) */
3     int membershipKey;
4     int rank;
5
6     MPI_Init(&argc, &argv);
7     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
8
9     /* User code must generate membershipKey in the range [0, 1, 2] */
10    membershipKey = rank % 3;
11
12    /* Build intra-communicator for local sub-group */
13    MPI_Comm_split(MPI_COMM_WORLD, membershipKey, rank, &myComm);
14
15    /* Build inter-communicators. Tags are hard-coded. */
16    if (membershipKey == 0)
17    {
18        /* Group 0 communicates with group 1. */
19        MPI_Intercomm_create( myComm, 0, MPI_COMM_WORLD, 1,
20                             1, &myFirstComm);
21    }
22    else if (membershipKey == 1)
23    {
24        /* Group 1 communicates with groups 0 and 2. */
25        MPI_Intercomm_create( myComm, 0, MPI_COMM_WORLD, 0,
26                             1, &myFirstComm);
27        MPI_Intercomm_create( myComm, 0, MPI_COMM_WORLD, 2,
28                             12, &mySecondComm);
29    }
30    else if (membershipKey == 2)
31    {
32        /* Group 2 communicates with group 1. */
33        MPI_Intercomm_create( myComm, 0, MPI_COMM_WORLD, 1,
34                             12, &myFirstComm);
35    }
36
37    /* Do work ... */
38
39    switch(membershipKey) /* free communicators appropriately */
40    {
41    case 1:
42        MPI_Comm_free(&mySecondComm);
43    case 0:
44    case 2:
45        MPI_Comm_free(&myFirstComm);
46        break;
47    }
48 }
```



```

1     else if (membershipKey == 2)
2     {           /* Group 2 communicates with groups 0 and 1. */
3         MPI_Intercomm_create( myComm, 0, MPI_COMM_WORLD, 0,
4                               2, &myFirstComm);
5         MPI_Intercomm_create( myComm, 0, MPI_COMM_WORLD, 1,
6                               12, &mySecondComm);
7     }
8
9     /* Do some work ... */
10
11    /* Then free communicators before terminating... */
12    MPI_Comm_free(&myFirstComm);
13    MPI_Comm_free(&mySecondComm);
14    MPI_Comm_free(&myComm);
15    MPI_Finalize();
16 }
17

```

Example 3: Building Name Service for Intercommunication

The following procedures exemplify the process by which a user could create name service for building intercommunicators via a rendezvous involving a server communicator, and a tag name selected by both groups.

After all MPI processes execute `MPI_INIT`, every process calls the example function, `Init_server()`, defined below. Then, if the `new_world` returned is `NULL`, the process getting `NULL` is required to implement a server function, in a reactive loop, `Do_server()`. Everyone else just does their prescribed computation, using `new_world` as the new effective “global” communicator. One designated process calls `Undo_Server()` to get rid of the server when it is not needed any longer.

Features of this approach include:

- Support for multiple name servers
- Ability to scope the name servers to specific processes
- Ability to make such servers come and go as desired.

```

35
36 #define INIT_SERVER_TAG_1 666
37 #define UNDO_SERVER_TAG_1 777
38
39 static int server_key_val;
40
41 /* for attribute management for server_comm, copy callback: */
42 void handle_copy_fn(MPI_Comm *oldcomm, int *keyval, void *extra_state,
43 void *attribute_val_in, void **attribute_val_out, int *flag)
44 {
45     /* copy the handle */
46     *attribute_val_out = attribute_val_in;
47     *flag = 1; /* indicate that copy to happen */
48 }

```

```

1
2 int Init_server(peer_comm, rank_of_server, server_comm, new_world)
3 MPI_Comm peer_comm;
4 int rank_of_server;
5 MPI_Comm *server_comm;
6 MPI_Comm *new_world; /* new effective world, sans server */
7 {
8     MPI_Comm temp_comm, lone_comm;
9     MPI_Group peer_group, temp_group;
10    int rank_in_peer_comm, size, color, key = 0;
11    int peer_leader, peer_leader_rank_in_temp_comm;
12
13    MPI_Comm_rank(peer_comm, &rank_in_peer_comm);
14    MPI_Comm_size(peer_comm, &size);
15
16    if ((size < 2) || (0 > rank_of_server) || (rank_of_server >= size))
17        return (MPI_ERR_OTHER);
18
19    /* create two communicators, by splitting peer_comm
20     into the server process, and everyone else */
21
22    peer_leader = (rank_of_server + 1) % size; /* arbitrary choice */
23
24    if ((color = (rank_in_peer_comm == rank_of_server)))
25    {
26        MPI_Comm_split(peer_comm, color, key, &lone_comm);
27
28        MPI_Intercomm_create(lone_comm, 0, peer_comm, peer_leader,
29                             INIT_SERVER_TAG_1, server_comm);
30
31        MPI_Comm_free(&lone_comm);
32        *new_world = MPI_COMM_NULL;
33    }
34    else
35    {
36        MPI_Comm_Split(peer_comm, color, key, &temp_comm);
37
38        MPI_Comm_group(peer_comm, &peer_group);
39        MPI_Comm_group(temp_comm, &temp_group);
40        MPI_Group_translate_ranks(peer_group, 1, &peer_leader,
41 temp_group, &peer_leader_rank_in_temp_comm);
42
43        MPI_Intercomm_create(temp_comm, peer_leader_rank_in_temp_comm,
44                             peer_comm, rank_of_server,
45                             INIT_SERVER_TAG_1, server_comm);
46
47        /* attach new_world communication attribute to server_comm: */
48

```



```

{
    while (de_queue(queue, MPI_ANY_TAG, &pairs_rank_in_new_world,
                  &pairs_rank_in_server))
        ;

    MPI_Intercomm_free(&server_comm);
    break;
}

if (de_queue(queue, client_tag, &pairs_rank_in_new_world,
            &pairs_rank_in_server))
{
    /* matched pair with same tag, tell them
       about each other! */
    buffer[0] = pairs_rank_in_new_world;
    MPI_Send(buffer, 1, MPI_INT, client_src, client_tag,
             server_comm);

    buffer[0] = client_rank_in_new_world;
    MPI_Send(buffer, 1, MPI_INT, pairs_rank_in_server, client_tag,
             server_comm);
}
else
    en_queue(queue, client_tag, client_source,
            client_rank_in_new_world);
}
}

```

A particular process would be responsible for ending the server when it is no longer needed. Its call to `Undo_server` would terminate server function.

```

int Undo_server(server_comm)    /* example client that ends server */
MPI_Comm *server_comm;
{
    int buffer = 0;
    MPI_Send(&buffer, 1, MPI_INT, 0, UNDO_SERVER_TAG_1, *server_comm);
    MPI_Intercomm_free(server_comm);
}

```

The following is a blocking name-service for inter-communication, with same semantic restrictions as `MPI_Intercomm_create`, but simplified syntax. It uses the functionality just defined to create the name service.

```

int Intercomm_name_create(local_comm, server_comm, tag, comm)
MPI_Comm local_comm, server_comm;
int tag;
MPI_Comm *comm;
{

```



```

1   int error;
2   int found;   /* attribute acquisition mgmt for new_world */
3               /* comm in server_comm */
4   void *val;
5
6   MPI_Comm new_world;
7
8   int buffer[10], rank;
9   int local_leader = 0;
10
11  MPI_Attr_get(server_comm, server_keyval, &val, &found);
12  new_world = (MPI_Comm)val; /* retrieve cached handle */
13
14  MPI_Comm_rank(server_comm, &rank); /* rank in local group */
15
16  if (rank == local_leader)
17  {
18      buffer[0] = rank;
19      MPI_Send(&buffer, 1, MPI_INT, 0, tag, server_comm);
20      MPI_Recv(&buffer, 1, MPI_INT, 0, tag, server_comm);
21  }
22
23  error = MPI_Intercomm_create(local_comm, local_leader, new_world,
24                              buffer[0], tag, comm);
25
26  return(error);
27 }
28
29

```

5.7 Caching

MPI provides a “caching” facility that allows an application to attach arbitrary pieces of information, called **attributes**, to communicators. 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,
- quickly retrieve that information, and
- be guaranteed that out-of-date information is never retrieved, even if the communicator 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 process-local attributes, via this attributing-caching mechanism. (*End of advice to users.*)

5.7.1 Functionality

Attributes are attached to communicators. Attributes are local to the process and specific to the communicator to which they are attached. Attributes are not propagated by MPI from one communicator to another except when the communicator is duplicated using `MPI_COMM_DUP` (and even then the application must give specific permission through callback functions for the attribute to be copied).

Advice to users. Attributes in C are of type `void *`. Typically, such an attribute will be a pointer to a structure that contains further information, or a handle to an MPI object. In Fortran, attributes are of type `INTEGER`. Such attribute can be a handle to an MPI object, or just an integer-valued attribute. (*End of advice to users.*)

Advice to implementors. Attributes are scalar values, equal in size to, or larger than a C-language pointer. Attributes can always hold an MPI handle. (*End of advice to implementors.*)

The caching interface defined here represents that attributes be stored by MPI opaquely within a communicator. Accessor functions include the following:

- obtain a key value (used to identify an attribute); the user specifies “callback” functions by which MPI informs the application when the communicator is destroyed or copied.
- store and retrieve the value of an attribute;

Advice to implementors. Caching and callback functions are only called synchronously, in response to explicit application requests. This avoid problems that result from repeated crossings between user and system space. (This synchronous calling rule is a general property of MPI.)

The choice of key values is under control of MPI. This allows MPI to optimize its implementation of attribute sets. It also avoids conflict between independent modules caching information on the same communicators.

A much smaller interface, consisting of just a callback facility, would allow the entire caching facility to be implemented by portable code. However, with the minimal callback 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 process local.

```

1 MPI_KEYVAL_CREATE(copy_fn, delete_fn, keyval, extra_state)
2     IN      copy_fn           Copy callback function for keyval
3
4     IN      delete_fn        Delete callback function for keyval
5
6     OUT     keyval           key value for future access (integer)
7
8     IN      extra_state      Extra state for callback functions

```

```

9 int MPI_Keyval_create(MPI_Copy_function *copy_fn, MPI_Delete_function
10                      *delete_fn, int *keyval, void* extra_state)

```

```

11 MPI_KEYVAL_CREATE(COPY_FN, DELETE_FN, KEYVAL, EXTRA_STATE, IERROR)
12     EXTERNAL COPY_FN, DELETE_FN
13     INTEGER KEYVAL, EXTRA_STATE, IERROR

```

Generates a new attribute key. Keys are locally unique in a 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 `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:

```

20 typedef int MPI_Copy_function(MPI_Comm oldcomm, int keyval,
21                              void *extra_state, void *attribute_val_in,
22                              void *attribute_val_out, int *flag)

```

A Fortran declaration for such a function is as follows:

```

25 SUBROUTINE COPY_FUNCTION(OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
26                          ATTRIBUTE_VAL_OUT, FLAG, IERR)
27     INTEGER OLDCOMM, KEYVAL, EXTRA_STATE, ATTRIBUTE_VAL_IN,
28     ATTRIBUTE_VAL_OUT, IERR
29     LOGICAL FLAG

```

The copy callback function is invoked for each key value in `oldcomm` in arbitrary order. Each call to the copy callback is made with a key value and its corresponding attribute. If it returns `flag = 0`, then the attribute is deleted in the duplicated communicator. Otherwise (`flag = 1`), the new attribute value is set to the value returned in `attribute_val_out`. The function returns `MPI_SUCCESS` on success and an error code on failure (in which case `MPI_COMM_DUP` will fail).

`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 returning `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`.

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 `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:

```
typedef int MPI_Delete_function(MPI_Comm comm, int keyval,
void *attribute_val, void *extra_state);
```

A Fortran declaration for such a function is as follows:

```
SUBROUTINE DELETE_FUNCTION(COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR)
  INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR
```

This function is called by `MPI_COMM_FREE`, `MPI_ATTR_DELETE`, and `MPI_ATTR_PUT` 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).

`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 returning `MPI_SUCCESS`.

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_KEYVAL_CREATE`. Therefore, it can be used for static initialization of key values.

```
MPI_KEYVAL_FREE(keyval)
```

```
  INOUT  keyval                Frees the integer key value (integer)
```

```
int MPI_Keyval_free(int *keyval)
```

```
MPI_KEYVAL_FREE(KEYVAL, IERROR)
  INTEGER KEYVAL, IERROR
```

Frees an extant attribute key. This function sets the value of `keyval` to `MPI_KEYVAL_INVALID`. Note that it is not erroneous to free an attribute key that is in use, because the actual free does not transpire until after all references (in other communicators on the process) to the key have been freed. These references need to be explicitly freed by the program, either via calls to `MPI_ATTR_DELETE` that free one attribute instance, or by calls to `MPI_COMM_FREE` that free all attribute instances associated with the freed communicator.

1 MPI_ATTR_PUT(comm, keyval, attribute_val)

2 IN comm communicator to which attribute will be attached (handle)
 3
 4 IN keyval key value, as returned by
 5 MPI_KEYVAL_CREATE (integer)
 6
 7 IN attribute_val attribute value
 8

9 int MPI_Attr_put(MPI_Comm comm, int keyval, void* attribute_val)

10
 11 MPI_ATTR_PUT(COMM, KEYVAL, ATTRIBUTE_VAL, IERROR)

12 INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR

13
 14 This function stores the stipulated attribute value `attribute_val` for subsequent retrieval
 15 by `MPI_ATTR_GET`. If the value is already present, then the outcome is as if `MPI_ATTR_DELETE`
 16 was first called to delete the previous value (and the callback function `delete_fn` was exe-
 17 cuted), and a new value was next stored. The call is erroneous if there is no key with value
 18 `keyval`; in particular `MPI_KEYVAL_INVALID` is an erroneous key value. The call will fail if the
 19 `delete_fn` function returned an error code other than `MPI_SUCCESS`.
 20

21 MPI_ATTR_GET(comm, keyval, attribute_val, flag)

22 IN comm communicator to which attribute is attached (handle)
 23
 24 IN keyval key value (integer)
 25
 26 OUT attribute_val attribute value, unless `flag = false`
 27
 28 OUT flag true if an attribute value was extracted; false if no
 attribute is associated with the key

29 int MPI_Attr_get(MPI_Comm comm, int keyval, void *attribute_val, int *flag)

30
 31 MPI_ATTR_GET(COMM, KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)

32 INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR

33 LOGICAL FLAG

34
 35 Retrieves attribute value by key. The call is erroneous if there is no key with value
 36 `keyval`. On the other hand, the call is correct if the key value exists, but no attribute is
 37 attached on `comm` for that key; in such case, the call returns `flag = false`. In particular
 38 `MPI_KEYVAL_INVALID` is an erroneous key value.

39 *Advice to users.* The call to `MPI_Attr_put` passes in `attribute_val` the *value* of the
 40 attribute; the call to `MPI_Attr_get` passes in `attribute_val` the *address* of the the location
 41 where the attribute value is to be returned. Thus, if the attribute value itself is
 42 a pointer of type `void*`, the actual `attribute_val` parameter to `MPI_Attr_put` will be of
 43 type `void*` and the actual `attribute_val` parameter to `MPI_Attr_get` will be of type `void**`.
 44 (*End of advice to users.*)
 45

46 *Rationale.* The use of a formal parameter `attribute_val` or type `void*` (rather than
 47 `void**`) avoids the messy type casting that would be needed if the attribute value is
 48 declared with a type other than `void*`. (*End of rationale.*)

```

MPI_ATTR_DELETE(comm, keyval) 1
    IN      comm                communicator to which attribute is attached (handle) 2
    IN      keyval              The key value of the deleted attribute (integer) 3

```

```

int MPI_Attr_delete(MPI_Comm comm, int keyval) 4
MPI_ATTR_DELETE(COMM, KEYVAL, IERROR) 5
    INTEGER COMM, KEYVAL, IERROR 6

```

Delete attribute from cache by key. This function invokes the attribute delete function `delete_fn` specified when the `keyval` was created. The call will fail if the `delete_fn` function returns an error code other than `MPI_SUCCESS`.

Whenever a communicator is replicated using the function `MPI_COMM_DUP`, all callback 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.

5.7.2 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. The coding style assumes that MPI function results return only error statuses. (*End of advice to users.*)

```

/* key for this module's stuff: */ 23
static int gop_key = MPI_KEYVAL_INVALID; 24

typedef struct 25
{ 26
    int ref_count;          /* reference count */ 27
    /* other stuff, whatever else we want */ 28
} gop_stuff_type; 29

Efficient_Collective_Op (comm, ...) 30
MPI_Comm comm; 31
{ 32
    gop_stuff_type *gop_stuff; 33
    MPI_Group      group; 34
    int            foundflag; 35

    MPI_Comm_group(comm, &group); 36

    if (gop_key == MPI_KEYVAL_INVALID) /* get a key on first call ever */ 37
    { 38
        if ( ! MPI_Keyval_create( gop_stuff_copier, 39
                                gop_stuff_destructor, 40
                                &gop_key, (void *)0)); 41
        /* get the key while assigning its copy and delete callback 42
           behavior. */ 43

```

```
1
2     MPI_Abort (comm, 99);
3 }
4
5 MPI_Attr_get (comm, gop_key, &gop_stuff, &foundflag);
6 if (foundflag)
7 { /* This module has executed in this group before.
8     We will use the cached information */
9 }
10 else
11 { /* This is a group that we have not yet cached anything in.
12     We will now do so.
13     */
14
15     /* First, allocate storage for the stuff we want,
16        and initialize the reference count */
17
18     gop_stuff = (gop_stuff_type *) malloc (sizeof(gop_stuff_type));
19     if (gop_stuff == NULL) { /* abort on out-of-memory error */ }
20
21     gop_stuff -> ref_count = 1;
22
23     /* Second, fill in *gop_stuff with whatever we want.
24        This part isn't shown here */
25
26     /* Third, store gop_stuff as the attribute value */
27     MPI_Attr_put ( comm, gop_key, gop_stuff);
28 }
29 /* Then, in any case, use contents of *gop_stuff
30    to do the global op ... */
31 }
32
33 /* The following routine is called by MPI when a group is freed */
34
35 gop_stuff_destructor (comm, keyval, gop_stuff, extra)
36 MPI_Comm comm;
37 int keyval;
38 gop_stuff_type *gop_stuff;
39 void *extra;
40 {
41     if (keyval != gop_key) { /* abort -- programming error */ }
42
43     /* The group's being freed removes one reference to gop_stuff */
44     gop_stuff -> ref_count -= 1;
45
46     /* If no references remain, then free the storage */
47     if (gop_stuff -> ref_count == 0) {
48         free((void *)gop_stuff);
```

```

    }
  }
}
1
2
3
/* The following routine is called by MPI when a group is copied */
4
gop_stuff_copier (comm, keyval, extra, gop_stuff_in, gop_stuff_out, flag)
5
MPI_Comm comm;
6
int keyval;
7
gop_stuff_type *gop_stuff_in, *gop_stuff_out;
8
void *extra;
9
{
10
    if (keyval != gop_key) { /* abort -- programming error */ }
11
12
    /* The new group adds one reference to this gop_stuff */
13
    gop_stuff -> ref_count += 1;
14
    gop_stuff_out = gop_stuff_in;
15
}
16
17
18

```

5.8 Formalizing the Loosely Synchronous Model

In this section, we make further statements about the loosely synchronous model, with particular attention to intra-communication.

5.8.1 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 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.

5.8.2 Models of Execution

In the loosely synchronous model, transfer of control to a **parallel procedure** is effected by having each executing process invoke the procedure. The invocation is a collective operation: it is executed by all processes in the execution group, and invocations are similarly ordered at all processes. However, the invocation need not be synchronized.

We say that a parallel procedure is *active* in a process if the 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 a process, then this 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 process, and the group of executing processes is fixed. For example, all invocations of parallel procedures involve all processes, 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, 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 processes, then communicator creation be properly coordinated.

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

Chapter 6

Process Topologies

6.1 Introduction

This chapter discusses the MPI topology mechanism. A topology is an extra, optional attribute that one can give to an intra-communicator; topologies cannot be added to inter-communicators. A topology can provide a convenient naming mechanism for the processes of a group (within a communicator), and additionally, may assist the runtime system in mapping the processes onto hardware.

As stated in chapter 5, a process group in MPI is a collection of n processes. Each process in the group is assigned a rank between 0 and $n-1$. In many parallel applications a linear ranking of processes does not adequately reflect the logical communication pattern of the processes (which is usually determined by the underlying problem geometry and the numerical algorithm used). Often the processes are arranged in topological patterns such as two- or three-dimensional grids. More generally, the logical process arrangement is described by a graph. In this chapter we will refer to this logical process arrangement as the “virtual topology.”

A clear distinction must be made between the virtual process 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 proposed in this chapter deal only with machine-independent mapping.

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 [20]. 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 modern wormhole-routing architectures can be found in [10, 9].

Besides possible performance benefits, the virtual topology can function as a convenient, process-naming structure, with tremendous benefits for program readability

1 and notational power in message-passing programming. (*End of rationale.*)

3 6.2 Virtual Topologies

5 The communication pattern of a set of processes can be represented by a graph. The
6 nodes stand for the processes, and the edges connect processes that communicate with each
7 other. MPI provides message-passing between any pair of processes in a group. There
8 is no requirement for opening a channel explicitly. Therefore, a “missing link” in the
9 user-defined process graph does not prevent the corresponding processes from exchanging
10 messages. It means rather that this connection is neglected in the virtual topology. This
11 strategy implies that the topology gives no convenient way of naming this pathway of
12 communication. Another possible consequence is that an automatic mapping tool (if one
13 exists for the runtime environment) will not take account of this edge when mapping. Edges
14 in the communication graph are not weighted, so that processes are either simply connected
15 or not connected at all.

17 *Rationale.* Experience with similar techniques in PARMACS [5, 8] show that this
18 information is usually sufficient for a good mapping. Additionally, a more precise
19 specification is more difficult for the user to set up, and it would make the interface
20 functions substantially more complicated. (*End of rationale.*)

22 Specifying the virtual topology in terms of a graph is sufficient for all applications.
23 However, in many applications the graph structure is regular, and the detailed set-up of the
24 graph would be inconvenient for the user and might be less efficient at run time. A large frac-
25 tion of all parallel applications use process topologies like rings, two- or higher-dimensional
26 grids, or tori. These structures are completely defined by the number of dimensions and
27 the numbers of processes in each coordinate direction. Also, the mapping of grids and tori
28 is generally an easier problem than that of general graphs. Thus, it is desirable to address
29 these cases explicitly.

30 Process coordinates in a cartesian structure begin their numbering at 0. Row-major
31 numbering is always used for the processes in a cartesian structure. This means that, for
32 example, the relation between group rank and coordinates for four processes in a (2×2)
33 grid is as follows.

34
35 coord (0,0): rank 0
36 coord (0,1): rank 1
37 coord (1,0): rank 2
38 coord (1,1): rank 3

40 6.3 Embedding in MPI

42 The support for virtual topologies as defined in this chapter is consistent with other parts of
43 MPI, and, whenever possible, makes use of functions that are defined elsewhere. Topology
44 information is associated with communicators. It is added to communicators using the
45 caching mechanism described in Chapter 5.

6.4 Overview of the Functions

The functions `MPI_GRAPH_CREATE` and `MPI_CART_CREATE` are used to create general (graph) virtual topologies and cartesian topologies, respectively. 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 topology creation functions take as input an existing communicator `comm_old`, which defines the set of processes on which the topology is to be mapped. A new communicator `comm_topol` is created that carries the topological structure as cached information (see Chapter 5). 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 process structure is periodic or not. Note that an n -dimensional hypercube is an n -dimensional torus with 2 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 processes among a given number of dimensions.

Rationale. Similar functions are contained in EXPRESS [22] and PARMACS. (*End of rationale.*)

The function `MPI_TOPO_TEST` can be used to inquire about the topology associated with a communicator. The topological information can be extracted from the communicator using the functions `MPI_GRAPHDIMS_GET` and `MPI_GRAPH_GET`, for general graphs, and `MPI_CARTDIM_GET` and `MPI_CART_GET`, for cartesian topologies. Several additional functions are provided to manipulate cartesian topologies: the functions `MPI_CART_RANK` and `MPI_CART_COORDS` translate cartesian coordinates into a group rank, and vice-versa; the function `MPI_CART_SUB` can be used to extract a cartesian subspace (analogous to `MPI_COMM_SPLIT`). The function `MPI_CART_SHIFT` provides the information needed to communicate with neighbors in a cartesian dimension. The two functions `MPI_GRAPH_NEIGHBORS_COUNT` and `MPI_GRAPH_NEIGHBORS` can be used to extract the neighbors of a node in a graph. The function `MPI_CART_SUB` is collective over the input communicator's group; all other functions are local.

Two additional functions, `MPI_GRAPH_MAP` and `MPI_CART_MAP` are presented in the last section. In general these functions are not called by the user directly. However, together with the communicator manipulation functions presented in Chapter 5, they are sufficient to implement all other topology functions. Section 6.5.7 outlines such an implementation.

6.5 Topology Constructors

6.5.1 Cartesian Constructor

`MPI_CART_CREATE(comm_old, ndims, dims, periods, reorder, comm_cart)`

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>	ranking may be reordered (<code>true</code>) or not (<code>false</code>) (logical)
OUT	<code>comm_cart</code>	communicator with new cartesian topology (handle)

```
int MPI_Cart_create(MPI_Comm comm_old, int ndims, int *dims, int *periods,
                   int reorder, MPI_Comm *comm_cart)
```

```
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 process in the new group is identical to its rank in the old group. Otherwise, the function may reorder the 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`, then some processes are returned `MPI_COMM_NULL`, in analogy to `MPI_COMM_SPLIT`. The call is erroneous if it specifies a grid that is larger than the group size.

6.5.2 Cartesian Convenience Function: `MPI_DIMS_CREATE`

For cartesian topologies, the function `MPI_DIMS_CREATE` helps the user select a balanced distribution of processes per coordinate direction, depending on the number of processes in the group to be balanced and optional constraints that can be specified by the user. One use is to partition all the processes (the size of `MPI_COMM_WORLD`'s group) into an n -dimensional topology.

`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

```
int MPI_Dims_create(int nnodes, int ndims, int *dims)
```

```
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 non-increasing order. Array `dims` is suitable for use as input to routine `MPI_CART_CREATE`. `MPI_DIMS_CREATE` is local.

	dims before call	function call	dims on return
Example 6.1	(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

6.5.3 General (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>	ranking may be reordered (<code>true</code>) or not (<code>false</code>) (logical)
OUT	<code>comm_graph</code>	communicator with graph topology added (handle)

```
int MPI_Graph_create(MPI_Comm comm_old, int nnodes, int *index, int *edges,
    int reorder, MPI_Comm *comm_graph)
```

```
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 process in the new group is identical to its rank in the old group. Otherwise, the function may reorder the

processes. If the size, `nnodes`, of the graph is smaller than the size of the group of `comm`, then some processes are returned `MPI_COMM_NULL`, in analogy to `MPI_CART_CREATE` and `MPI_COMM_SPLIT`. 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 6.2 Assume there are four processes 0, 1, 2, 3 with the following adjacency matrix:

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, ..., nnodes-1`; the list of neighbors of node zero is stored in `edges[j]`, for $0 \leq j \leq \text{index}[0] - 1$ and the list of neighbors of node `i`, `i > 0`, is stored in `edges[j]`, $\text{index}[i-1] \leq j \leq \text{index}[i] - 1$.

In Fortran, `index(1)` is the degree of node zero, and `index(i+1) - index(i)` is the degree of node `i`, `i=1, ..., nnodes-1`; the list of neighbors of node zero is stored in `edges(j)`, for $1 \leq j \leq \text{index}(1)$ and the list of neighbors of node `i`, `i > 0`, is stored in `edges(j)`, $\text{index}(i) + 1 \leq j \leq \text{index}(i + 1)$.

Advice to implementors. The following topology information is likely to be stored with a communicator:

- Type of topology (cartesian/graph),
- For a cartesian topology:
 1. `ndims` (number of dimensions),
 2. `dims` (numbers of processes per coordinate direction),
 3. `periods` (periodicity information),
 4. `own_position` (own position in grid, could also be computed from rank and `dims`)
- For a graph topology:

1. `index`,
 2. `edges`,
- which are the vectors defining the graph structure.

For a graph structure the number of nodes is equal to the number of processes in the group. Therefore, the number of nodes does not have to be stored explicitly. An additional zero entry at the start of array `index` simplifies access to the topology information. (*End of advice to implementors.*)

6.5.4 Topology inquiry functions

If a 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	<code>comm</code>	communicator (handle)
OUT	<code>status</code>	topology type of communicator <code>comm</code> (state)

```
int MPI_Topo_test(MPI_Comm comm, int *status)
```

```
MPI_TOPO_TEST(COMM, STATUS, IERROR)
INTEGER COMM, STATUS, IERROR
```

The function `MPI_TOPO_TEST` returns the type of topology that is assigned to 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_UNDEFINED</code>	no topology

`MPI_GRAPHDIMS_GET(comm, nnodes, nedges)`

IN	<code>comm</code>	communicator for group with graph structure (handle)
OUT	<code>nnodes</code>	number of nodes in graph (integer) (same as number of processes in the group)
OUT	<code>nedges</code>	number of edges in graph (integer)

```
int MPI_Graphdims_get(MPI_Comm comm, int *nnodes, int *nedges)
```

```
MPI_GRAPHDIMS_GET(COMM, NNODES, NEDGES, IERROR)
INTEGER COMM, NNODES, NEDGES, IERROR
```

Functions `MPI_GRAPHDIMS_GET` and `MPI_GRAPH_GET` retrieve the graph-topology information that was associated with a communicator by `MPI_GRAPH_CREATE`.

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`.


```

1 MPI_GRAPH_GET(comm, maxindex, maxedges, index, edges)
2     IN      comm      communicator with graph structure (handle)
3
4     IN      maxindex  length of vector index in the calling program
5                        (integer)
6
7     IN      maxedges  length of vector edges in the calling program
8                        (integer)
9
10    OUT     index     array of integers containing the graph structure (for
11                        details see the definition of MPI_GRAPH_CREATE)
12
13    int MPI_Graph_get(MPI_Comm comm, int maxindex, int maxedges, int *index,
14                      int *edges)
15
16    MPI_GRAPH_GET(COMM, MAXINDEX, MAXEDGES, INDEX, EDGES, IERROR)
17      INTEGER COMM, MAXINDEX, MAXEDGES, INDEX(*), EDGES(*), IERROR
18
19
20 MPI_CARTDIM_GET(comm, ndims)
21     IN      comm      communicator with cartesian structure (handle)
22
23     OUT     ndims     number of dimensions of the cartesian structure (inte-
24                        ger)
25
26    int MPI_Cartdim_get(MPI_Comm comm, int *ndims)
27
28    MPI_CARTDIM_GET(COMM, NDIMS, IERROR)
29      INTEGER COMM, NDIMS, IERROR
30
31    The functions MPI_CARTDIM_GET and MPI_CART_GET return the cartesian topology
32    information that was associated with a communicator by MPI_CART_CREATE.
33
34 MPI_CART_GET(comm, maxdims, dims, periods, coords)
35     IN      comm      communicator with cartesian structure (handle)
36
37     IN      maxdims   length of vectors dims, periods, and coords in the
38                        calling program (integer)
39
40     OUT     dims      number of processes for each cartesian dimension (ar-
41                        ray of integer)
42
43     OUT     periods   periodicity (true/false) for each cartesian dimension
44                        (array of logical)
45
46     OUT     coords    coordinates of calling process in cartesian structure
47                        (array of integer)
48
49    int MPI_Cart_get(MPI_Comm comm, int maxdims, int *dims, int *periods,
50                    int *coords)

```

```

MPI_CART_GET(COMM, MAXDIMS, DIMS, PERIODS, COORDS, IERROR) 1
    INTEGER COMM, MAXDIMS, DIMS(*), COORDS(*), IERROR      2
    LOGICAL PERIODS(*)                                       3

```

```

MPI_CART_RANK(comm, coords, rank)                             6
    IN      comm      communicator with cartesian structure (handle) 8
    IN      coords    integer array (of size ndims) specifying the cartesian 9
                        coordinates of a process                    10
    OUT     rank      rank of specified process (integer)          11

```

```

int MPI_Cart_rank(MPI_Comm comm, int *coords, int *rank)     13

```

```

MPI_CART_RANK(COMM, COORDS, RANK, IERROR)                   15
    INTEGER COMM, COORDS(*), RANK, IERROR                    16

```

For a process group with cartesian structure, the function `MPI_CART_RANK` translates the logical process coordinates to process ranks as they are used by the point-to-point routines.

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 $0 \leq \text{coords}(i) < \text{dims}(i)$ automatically. Out-of-range coordinates are erroneous for non-periodic dimensions.

```

MPI_CART_COORDS(comm, rank, maxdims, coords)                 26
    IN      comm      communicator with cartesian structure (handle) 28
    IN      rank      rank of a process within group of comm (integer) 29
    IN      maxdims   length of vector coords in the calling program (inte- 30
                        ger)                                          31
    OUT     coords    integer array (of size ndims) containing the cartesian 32
                        coordinates of specified process (array of integers) 33

```

```

int MPI_Cart_coords(MPI_Comm comm, int rank, int maxdims, int *coords) 36

```

```

MPI_CART_COORDS(COMM, RANK, MAXDIMS, COORDS, IERROR)        37
    INTEGER COMM, RANK, MAXDIMS, COORDS(*), IERROR          38

```

The inverse mapping, rank-to-coordinates translation is provided by `MPI_CART_COORDS`.

1 MPI_GRAPH_NEIGHBORS_COUNT(comm, rank, nneighbors)

2 IN comm communicator with graph topology (handle)
 3
 4 IN rank rank of process in group of comm (integer)
 5 OUT nneighbors number of neighbors of specified process (integer)

6
 7 int MPI_Graph_neighbors_count(MPI_Comm comm, int rank, int *nneighbors)

8
 9 MPI_GRAPH_NEIGHBORS_COUNT(COMM, RANK, NNEIGHBORS, IERROR)
 10 INTEGER COMM, RANK, NNEIGHBORS, IERROR

11 MPI_GRAPH_NEIGHBORS_COUNT and MPI_GRAPH_NEIGHBORS provide adjacency
 12 information for a general, graph topology.

13
 14
 15 MPI_GRAPH_NEIGHBORS(comm, rank, maxneighbors, neighbors)

16 IN comm communicator with graph topology (handle)
 17
 18 IN rank rank of process in group of comm (integer)
 19 IN maxneighbors size of array neighbors (integer)
 20 OUT neighbors ranks of processes that are neighbors to specified pro-
 21 cess (array of integer)

22
 23
 24 int MPI_Graph_neighbors(MPI_Comm comm, int rank, int maxneighbors,
 25 int *neighbors)

26 MPI_GRAPH_NEIGHBORS(COMM, RANK, MAXNEIGHBORS, NEIGHBORS, IERROR)
 27 INTEGER COMM, RANK, MAXNEIGHBORS, NEIGHBORS(*), IERROR

28
 29 **Example 6.3** Suppose that `comm` is a communicator with a shuffle-exchange topology. The
 30 group has 2^n members. Each process is labeled by a_1, \dots, a_n with $a_i \in \{0, 1\}$, and has
 31 three neighbors: $\text{exchange}(a_1, \dots, a_n) = a_1, \dots, a_{n-1}, \bar{a}_n$ ($\bar{a} = 1 - a$), $\text{shuffle}(a_1, \dots, a_n) =$
 32 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
 33 illustrated below for $n = 3$.

34
 35

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

36
 37
 38
 39
 40
 41
 42
 43
 44
 45

46 Suppose that the communicator `comm` has this topology associated with it. The follow-
 47 ing code fragment cycles through the three types of neighbors and performs an appropriate
 48 permutation for each.

```

C  assume: each process has stored a real number A.           1
C  extract neighborhood information                             2
    CALL MPI_COMM_RANK(comm, myrank, ierr)                    3
    CALL MPI_GRAPH_NEIGHBORS(comm, myrank, 3, neighbors, ierr) 4
C  perform exchange permutation                               5
    CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, neighbors(1), 0, 6
    +   neighbors(1), 0, comm, status, ierr)                   7
C  perform shuffle permutation                                8
    CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, neighbors(2), 0, 9
    +   neighbors(3), 0, comm, status, ierr)                   10
C  perform unshuffle permutation                              11
    CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, neighbors(3), 0, 12
    +   neighbors(2), 0, comm, status, ierr)                   13

```

6.5.5 Cartesian Shift Coordinates

If the process topology is a cartesian structure, a `MPI_SENDRECV` operation is likely to be used along a coordinate direction to perform a shift of data. As input, `MPI_SENDRECV` takes the rank of a source process for the receive, and the rank of a destination process for the send. If the function `MPI_CART_SHIFT` is called for a cartesian process group, it provides the calling 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). The function is local.

`MPI_CART_SHIFT(comm, direction, disp, rank_source, rank_dest)`

IN	comm	communicator with cartesian structure (handle)	27
IN	direction	coordinate dimension of shift (integer)	28
IN	disp	displacement (> 0: upwards shift, < 0: downwards shift) (integer)	29 30 31
OUT	rank_source	rank of source process (integer)	32
OUT	rank_dest	rank of destination process (integer)	33

```

int MPI_Cart_shift(MPI_Comm comm, int direction, int disp, int *rank_source,
    int *rank_dest)

```

```

MPI_CART_SHIFT(COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR)
    INTEGER COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR

```

The `direction` argument indicates the dimension of the shift, i.e., the coordinate which value is modified by the shift. The coordinates are numbered from 0 to `ndims-1`, when `ndims` is the number of dimensions.

Depending on the periodicity of the cartesian group in the specified coordinate direction, `MPI_CART_SHIFT` provides the identifiers for a circular or an end-off shift. In the case of an end-off shift, the value `MPI_PROC_NULL` may be returned in `rank_source` or `rank_dest`, indicating that the source or the destination for the shift is out of range.

Example 6.4 The communicator, `comm`, has a two-dimensional, periodic, cartesian topology associated with it. A two-dimensional array of `REALs` is stored one element per process, in variable `A`. One wishes to skew this array, by shifting column `i` (vertically, i.e., along the column) by `i` steps.

```

1
2
3
4
5
6     ....
7 C find process rank
8     CALL MPI_COMM_RANK(comm, rank, ierr))
9 C find cartesian coordinates
10    CALL MPI_CART_COORDS(comm, rank, maxdims, coords, ierr)
11 C compute shift source and destination
12    CALL MPI_CART_SHIFT(comm, 0, coords(2), source, dest, ierr)
13 C skew array
14    CALL MPI_SENDRECV_REPLACE(A, 1, MPI_REAL, dest, 0, source, 0, comm,
15    +                          status, ierr)

```

Advice to users. In Fortran, the dimension indicated by `DIRECTION = i` has `DIMS(i+1)` nodes, where `DIMS` is the array that was used to create the grid. In C, the dimension indicated by `direction = i` is the dimension specified by `dims[i]`. (*End of advice to users.*)

6.5.6 Partitioning of Cartesian structures

```

24 MPI_CART_SUB(comm, remain_dims, newcomm)

```

25	IN	<code>comm</code>	communicator with cartesian structure (handle)
26			
27	IN	<code>remain_dims</code>	the <code>i</code> th entry of <code>remain_dims</code> specifies whether the <code>i</code> th dimension is kept in the subgrid (<code>true</code>) or is dropped (<code>false</code>) (logical vector)
28			
29			
30	OUT	<code>newcomm</code>	communicator containing the subgrid that includes the calling process (handle)
31			
32			

```

33 int MPI_Cart_sub(MPI_Comm comm, int *remain_dims, MPI_Comm *newcomm)

```

```

34 MPI_CART_SUB(COMM, REMAIN_DIMS, NEWCOMM, IERROR)

```

```

35     INTEGER COMM, NEWCOMM, IERROR

```

```

36     LOGICAL REMAIN_DIMS(*)

```

If a cartesian topology has been created with `MPI_CART_CREATE`, the function `MPI_CART_SUB` can be used to partition the communicator group into subgroups that form lower-dimensional cartesian subgrids, and to build for each subgroup a communicator with the associated subgrid cartesian topology. (This function is closely related to `MPI_COMM_SPLIT`.)

Example 6.5 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,

```

47     MPI_CART_SUB(comm, remain_dims, comm_new),

```

48

will create three communicators each with eight processes in a 2×4 cartesian topology. If `remain_dims = (false, false, true)` then the call to `MPI_CART_SUB(comm, remain_dims, comm_new)` will create six non-overlapping communicators, each with four processes, in a one-dimensional cartesian topology.

6.5.7 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, unless he or she is creating additional virtual topology capability other than that provided by MPI.

`MPI_CART_MAP(comm, ndims, dims, periods, newrank)`

IN	<code>comm</code>	input communicator (handle)
IN	<code>ndims</code>	number of dimensions of cartesian structure (integer)
IN	<code>dims</code>	integer array of size <code>ndims</code> specifying the number of processes in each coordinate direction
IN	<code>periods</code>	logical array of size <code>ndims</code> specifying the periodicity specification in each coordinate direction
OUT	<code>newrank</code>	reordered rank of the calling process; <code>MPI_UNDEFINED</code> if calling process does not belong to grid (integer)

```
int MPI_Cart_map(MPI_Comm comm, int ndims, int *dims, int *periods,
                int *newrank)
```

```
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 process on the physical machine. A possible implementation of this function is to always return the rank of the calling 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 ≠ MPI_UNDEFINED`, `color = MPI_UNDEFINED` otherwise, and `key = newrank`.

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 new function for general graph structures is as follows.

```

1 MPI_GRAPH_MAP(comm, nnodes, index, edges, newrank)
2     IN      comm      input communicator (handle)
3
4     IN      nnodes    number of graph nodes (integer)
5
6     IN      index     integer array specifying the graph structure, see
7                        MPI_GRAPH_CREATE
8     IN      edges     integer array specifying the graph structure
9     OUT     newrank   reordered rank of the calling process; MPI_UNDEFINED
10                        if the calling process does not belong to graph (inte-
11                        ger)

```

```

12
13 int MPI_Graph_map(MPI_Comm comm, int nnodes, int *index, int *edges,
14                  int *newrank)

```

```

15 MPI_GRAPH_MAP(COMM, NNODES, INDEX, EDGES, NEWRANK, IERROR)
16     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.*)

6.6 An Application Example

Example 6.6 The example in figure 6.1 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 processes organize themselves in a two-dimensional structure. Each 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 process computes new values for the solution grid function at all points owned by the process. Then the values at inter-process boundaries have to be exchanged with neighboring processes. For example, the exchange subroutine might contain a call like `MPI_SEND(...,neigh_rank(1),...)` to send updated values to the left-hand neighbor $(i-1, j)$.

```

integer ndims, num_neigh
logical reorder
parameter (ndims=2, num_neigh=4, reorder=.true.)
integer comm, comm_cart, dims(ndims), neigh_def(ndims), ierr
integer neigh_rank(num_neigh), own_position(ndims), i, j
logical periods(ndims)
real*8 u(0:101,0:101), f(0:101,0:101)
data dims / ndims * 0 /
comm = MPI_COMM_WORLD
C   Set process grid size and periodicity
call MPI_DIMS_CREATE(comm, ndims, dims,ierr)
periods(1) = .TRUE.
periods(2) = .TRUE.
C   Create a grid structure in WORLD group and inquire about own position
call MPI_CART_CREATE (comm, ndims, dims, periods, comm_cart,ierr)
call MPI_CART_GET (comm_cart, ndims, dims, periods, own_position,ierr)
C   Look up the ranks for the neighbors. Own process coordinates are (i,j).
C   Neighbors are (i-1,j), (i+1,j), (i,j-1), (i,j+1)
i = own_position(1)
j = own_position(2)
neigh_def(1) = i-1
neigh_def(2) = j
call MPI_CART_RANK (comm_cart, neigh_def, neigh_rank(1),ierr)
neigh_def(1) = i+1
neigh_def(2) = j
call MPI_CART_RANK (comm_cart, neigh_def, neigh_rank(2),ierr)
neigh_def(1) = i
neigh_def(2) = j-1
call MPI_CART_RANK (comm_cart, neigh_def, neigh_rank(3),ierr)
neigh_def(1) = i
neigh_def(2) = j+1
call MPI_CART_RANK (comm_cart, neigh_def, neigh_rank(4),ierr)
C   Initialize the grid functions and start the iteration
call init (u, f)
do 10 it=1,100
  call relax (u, f)
C   Exchange data with neighbor processes
  call exchange (u, comm_cart, neigh_rank, num_neigh)
10 continue
call output (u)
end

```

Figure 6.1: Set-up of process structure for two-dimensional parallel Poisson solver.

Chapter 7

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.

7.1 Implementation information

7.1.1 Version Inquiries

In order to cope with changes to the MPI Standard, there are both compile-time and runtime 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 and C++,

```
#define MPI_VERSION    1
#define MPI_SUBVERSION 2
```

in Fortran,

```
INTEGER MPI_VERSION, MPI_SUBVERSION
PARAMETER (MPI_VERSION    = 1)
PARAMETER (MPI_SUBVERSION = 2)
```

For runtime determination,

```
MPI_GET_VERSION( version, subversion )
```

```
OUT    version          version number (integer)
OUT    subversion       subversion number (integer)
```

```
int MPI_Get_version(int *version, int *subversion)
```

```
MPI_GET_VERSION(VERSION, SUBVERSION, IERROR)
INTEGER VERSION, SUBVERSION, IERROR
```

MPI_GET_VERSION is one of the few functions that can be called before MPI_INIT and after MPI_FINALIZE.

7.1.2 Environmental Inquiries

A set of attributes that describe the execution environment are attached to the communicator `MPI_COMM_WORLD` when MPI is initialized. The value of these attributes can be inquired by using the function `MPI_ATTR_GET` described in Chapter 5. 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_HOST` Host process rank, if such exists, `MPI_PROC_NULL`, otherwise.

`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.

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 legal value for `MPI_TAG_UB`.

The attribute `MPI_TAG_UB` has the same value on all processes of `MPI_COMM_WORLD`.

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 require that a `HOST` exists.

The attribute `MPI_HOST` 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 ANSI-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,

1 then its rank will be returned. Otherwise, if some process can provide language-standard
 2 I/O then the rank of one such process will be returned. The same value need not be
 3 returned by all processes. If no process can provide language-standard I/O, then the value
 4 MPI_PROC_NULL will be returned.

5
 6 *Advice to users.* Note that input is not collective, and this attribute does *not* indicate
 7 which process can or does provide input. (*End of advice to users.*)
 8

9 Clock synchronization

10 The value returned for MPI_WTIME_IS_GLOBAL is 1 if clocks at all processes in
 11 MPI_COMM_WORLD are synchronized, 0 otherwise. A collection of clocks is considered syn-
 12 chronized if explicit effort has been taken to synchronize them. The expectation is that
 13 the variation in time, as measured by calls to MPI_WTIME, will be less than one half the
 14 round-trip time for an MPI message of length zero. If time is measured at a process just
 15 before a send and at another process just after a matching receive, the second time should
 16 be always higher than the first one.
 17

18 The attribute MPI_WTIME_IS_GLOBAL need not be present when the clocks are not
 19 synchronized (however, the attribute key MPI_WTIME_IS_GLOBAL is always valid). This
 20 attribute may be associated with communicators other than MPI_COMM_WORLD.

21 The attribute MPI_WTIME_IS_GLOBAL has the same value on all processes of MPI_COMM_WORLD.
 22

23 MPI_GET_PROCESSOR_NAME(name, resultlen)

24	OUT	name	A unique specifier for the actual (as opposed to vir-
25			tual) node.
26			
27	OUT	resultlen	Length (in printable characters) of the result returned
28			in name
29			

30 `int MPI_Get_processor_name(char *name, int *resultlen)`

31 `MPI_GET_PROCESSOR_NAME(NAME, RESULTLEN, IERROR)`

32 CHARACTER*(*) NAME

33 INTEGER RESULTLEN, IERROR

34
 35 This routine returns the name of the processor on which it was called at the moment
 36 of the call. The name is a character string for maximum flexibility. From this value it
 37 must be possible to identify a specific piece of hardware; possible values include “processor
 38 9 in rack 4 of mpp.cs.org” and “231” (where 231 is the actual processor number in the
 39 running homogeneous system). The argument name must represent storage that is at least
 40 MPI_MAX_PROCESSOR_NAME characters long. MPI_GET_PROCESSOR_NAME may write up
 41 to this many characters into name.

42 The number of characters actually written is returned in the output argument, resultlen.
 43 In C, a null character is additionally stored at name[resultlen]. The resultlen cannot
 44 be larger than MPI_MAX_PROCESSOR_NAME-1. In Fortran, name is padded on the right with
 45 blank characters. The resultlen cannot be larger than MPI_MAX_PROCESSOR_NAME.
 46

47 *Rationale.* This function allows MPI implementations that do process migration
 48 to return the current processor. Note that nothing in MPI *requires* or defines pro-

cess 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.*)

The constant `MPI_BSEND_OVERHEAD` provides an upper bound on the fixed overhead per message buffered by a call to `MPI_BSEND` (see Section 3.6.1).

7.2 Error handling

An MPI implementation cannot or may choose not to handle some errors that occur during MPI calls. These can include errors that generate exceptions or traps, such as floating point errors or access violations. The set of errors that are handled by MPI is implementation-dependent. Each such error generates an **MPI exception**.

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.

A user can associate an error handler with a communicator. The specified error handling routine will be used for any MPI exception that occurs during a call to MPI for a communication with this communicator. MPI calls that are not related to any communicator are considered to be attached to the communicator `MPI_COMM_WORLD`. The attachment of error handlers to communicators is purely local: different processes may attach different error handlers to the same communicator.

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.

Several predefined error handlers are available in MPI:

`MPI_ERRORS_ARE_FATAL` The handler, when called, causes the program to abort on all executing processes. This has the same effect as if `MPI_ABORT` was called by the process that invoked the handler.

`MPI_ERRORS_RETURN` The handler has no effect other than returning the error code to the user.

Implementations may provide additional predefined error handlers and programmers can code their own error handlers.

The error handler `MPI_ERRORS_ARE_FATAL` is associated by default with `MPI_COMM_WORLD` after initialization. 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 non trivial MPI error handler.

1 After an error is detected, the state of MPI is undefined. That is, using a user-defined
 2 error handler, or `MPI_ERRORS_RETURN`, does *not* necessarily allow the user to continue to
 3 use MPI after an error is detected. The purpose of these error handlers is to allow a user to
 4 issue user-defined error messages and to take actions unrelated to MPI (such as flushing I/O
 5 buffers) before a program exits. An MPI implementation is free to allow MPI to continue
 6 after an error but is not required to do so.

7
 8 *Advice to implementors.* A good quality implementation will, to the greatest possible
 9 extent, circumscribe the impact of an error, so that normal processing can continue
 10 after an error handler was invoked. The implementation documentation will provide
 11 information on the possible effect of each class of errors. (*End of advice to implemen-*
 12 *tors.*)

13
 14 An MPI error handler is an opaque object, which is accessed by a handle. MPI calls
 15 are provided to create new error handlers, to associate error handlers with communicators,
 16 and to test which error handler is associated with a communicator.

17
 18 `MPI_ERRHANDLER_CREATE(function, errhandler)`

19	IN	function	user defined error handling procedure
20			
21	OUT	errhandler	MPI error handler (handle)

22
 23 `int MPI_Errhandler_create(MPI_Handler_function *function,`
 24 `MPI_Errhandler *errhandler)`

25
 26 `MPI_ERRHANDLER_CREATE(FUNCTION, ERRHANDLER, IERROR)`
 27 `EXTERNAL FUNCTION`
 28 `INTEGER ERRHANDLER, IERROR`

29 Register the user routine `function` for use as an MPI exception handler. Returns in
 30 `errhandler` a handle to the registered exception handler.

31 In the C language, the user routine should be a C function of type `MPI_Handler_function`,
 32 which is defined as:

33
 34 `typedef void (MPI_Handler_function)(MPI_Comm *, int *, ...);`
 35

36 The first argument is the communicator in use. The second is the error code to be re-
 37 turned by the MPI routine that raised the error. If the routine would have returned
 38 `MPI_ERR_IN_STATUS`, it is the error code returned in the status for the request that caused
 39 the error handler to be invoked. The remaining arguments are “`stdargs`” arguments whose
 40 number and meaning is implementation-dependent. An implementation should clearly doc-
 41 ument these arguments. Addresses are used so that the handler may be written in Fortran.

42 In the Fortran language, the user routine should be of the form:

43
 44 `SUBROUTINE HANDLER_FUNCTION(COMM, ERROR_CODE,)`
 45 `INTEGER COMM, ERROR_CODE`
 46

47 *Advice to users.* Users are discouraged from using a Fortran `HANDLER_FUNCTION`
 48 since the routine expects a variable number of arguments. Some Fortran systems

may allow this but some may fail to give the correct result or compile/link this code. Thus, it will not, in general, be possible to create portable code with a Fortran HANDLER_FUNCTION. (*End of advice to users.*)

Rationale. The variable argument list is provided because it provides an ANSI-standard hook for providing additional information to the error handler; without this hook, ANSI C prohibits additional arguments. (*End of rationale.*)

MPI_ERRHANDLER_SET(comm, errhandler)

IN	comm	communicator to set the error handler for (handle)
IN	errhandler	new MPI error handler for communicator (handle)

int MPI_Errhandler_set(MPI_Comm comm, MPI_Errhandler errhandler)

MPI_ERRHANDLER_SET(COMM, ERRHANDLER, IERROR)
 INTEGER COMM, ERRHANDLER, IERROR

Associates the new error handler `errhandler` with communicator `comm` at the calling process. Note that an error handler is always associated with the communicator.

MPI_ERRHANDLER_GET(comm, errhandler)

IN	comm	communicator to get the error handler from (handle)
OUT	errhandler	MPI error handler currently associated with communicator (handle)

int MPI_Errhandler_get(MPI_Comm comm, MPI_Errhandler *errhandler)

MPI_ERRHANDLER_GET(COMM, ERRHANDLER, IERROR)
 INTEGER COMM, ERRHANDLER, IERROR

Returns in `errhandler` (a handle to) the error handler that is currently associated with communicator `comm`.

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.

MPI_ERRHANDLER_FREE(errhandler)

INOUT	errhandler	MPI error handler (handle)
-------	------------	----------------------------

int MPI_Errhandler_free(MPI_Errhandler *errhandler)

MPI_ERRHANDLER_FREE(ERRHANDLER, IERROR)
 INTEGER ERRHANDLER, IERROR

1 Marks the error handler associated with `errhandler` for deallocation and sets `errhandler`
 2 to `MPI_ERRHANDLER_NULL`. The error handler will be deallocated after all communicators
 3 associated with it have been deallocated.

4
 5
 6 `MPI_ERROR_STRING(errorcode, string, resultlen)`

7	IN	<code>errorcode</code>	Error code returned by an MPI routine
8	OUT	<code>string</code>	Text that corresponds to the <code>errorcode</code>
9			
10	OUT	<code>resultlen</code>	Length (in printable characters) of the result returned 11 in <code>string</code>

12
 13 `int MPI_Error_string(int errorcode, char *string, int *resultlen)`

14 `MPI_ERROR_STRING(ERRORCODE, STRING, RESULTLEN, IERROR)`

15 `INTEGER ERRORCODE, RESULTLEN, IERROR`

16 `CHARACTER*(*) STRING`

17
 18 Returns the error string associated with an error code or class. The argument `string`
 19 must represent storage that is at least `MPI_MAX_ERROR_STRING` characters long.

20 The number of characters actually written is returned in the output argument, `resultlen`.

21
 22 *Rationale.* The form of this function was chosen to make the Fortran and C bindings
 23 similar. A version that returns a pointer to a string has two difficulties. First, the
 24 return string must be statically allocated and different for each error message (allowing
 25 the pointers returned by successive calls to `MPI_ERROR_STRING` to point to the correct
 26 message). Second, in Fortran, a function declared as returning `CHARACTER*(*)` can
 27 not be referenced in, for example, a `PRINT` statement. (*End of rationale.*)

28 29 7.3 Error codes and classes

30
 31 The error codes returned by MPI are left entirely to the implementation (with the exception
 32 of `MPI_SUCCESS`). This is done to allow an implementation to provide as much information
 33 as possible in the error code (for use with `MPI_ERROR_STRING`).

34 To make it possible for an application to interpret an error code, the routine `MPI_ERROR_CLASS`
 35 converts any error code into one of a small set of standard error codes, called *error classes*.

36 Valid error classes include

37
 38
 39
 40
 41
 42
 43
 44
 45
 46
 47
 48

MPI_SUCCESS	No error	1
MPI_ERR_BUFFER	Invalid buffer pointer	2
MPI_ERR_COUNT	Invalid count argument	3
MPI_ERR_TYPE	Invalid datatype argument	4
MPI_ERR_TAG	Invalid tag argument	5
MPI_ERR_COMM	Invalid communicator	6
MPI_ERR_RANK	Invalid rank	7
MPI_ERR_REQUEST	Invalid request (handle)	8
MPI_ERR_ROOT	Invalid root	9
MPI_ERR_GROUP	Invalid group	10
MPI_ERR_OP	Invalid operation	11
MPI_ERR_TOPOLOGY	Invalid topology	12
MPI_ERR_DIMS	Invalid dimension argument	13
MPI_ERR_ARG	Invalid argument of some other kind	14
MPI_ERR_UNKNOWN	Unknown error	15
MPI_ERR_TRUNCATE	Message truncated on receive	16
MPI_ERR_OTHER	Known error not in this list	17
MPI_ERR_INTERN	Internal MPI (implementation) error	18
MPI_ERR_IN_STATUS	Error code is in status	19
MPI_ERR_PENDING	Pending request	20
MPI_ERR_LASTCODE	Last error code	21

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. An MPI error class is a valid MPI error code. Specifically, 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.*)

```
MPI_ERROR_CLASS( errorcode, errorclass )
```

IN	errorcode	Error code returned by an MPI routine
OUT	errorclass	Error class associated with errorcode

```
int MPI_Error_class(int errorcode, int *errorclass)
```

```
MPI_ERROR_CLASS(ERRORCODE, ERRORCLASS, IERROR)
INTEGER ERRORCODE, ERRORCLASS, IERROR
```

The function `MPI_ERROR_CLASS` maps each standard error code (error class) onto itself.

7.4 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()

```
double MPI_Wtime(void)
```

```
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”), it allows high-resolution, and carries no unnecessary baggage. One would use it like this:

```
{
  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).

MPI_WTICK()

```
double MPI_Wtick(void)
```

```
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} .

7.5 Startup

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 to be performed before other MPI routines may be called. To provide for this, MPI includes an initialization routine `MPI_INIT`.

`MPI_INIT()`

```
int MPI_Init(int *argc, char ***argv)
```

```
MPI_INIT(IERROR)
    INTEGER IERROR
```

This routine must be called before any other MPI routine. It must be called at most once; subsequent calls are erroneous (see `MPI_INITIALIZED`).

All MPI programs must contain a call to `MPI_INIT`; this routine must be called before any other MPI routine (apart from `MPI_INITIALIZED`) is called. The version for ANSI C accepts the `argc` and `argv` that are provided by the arguments to `main`:

```
int main(argc, argv)
int argc;
char **argv;
{
    MPI_Init(&argc, &argv);

    /* parse arguments */
    /* main program */

    MPI_Finalize();    /* see below */
}
```

The Fortran version takes only `IERROR`.

An MPI implementation is free to require that the arguments in the C binding must be the arguments to `main`.

Rationale. The command line arguments are provided to `MPI_Init` to allow an MPI implementation to use them in initializing the MPI environment. They are passed by reference to allow an MPI implementation to *provide* them in environments where the command-line arguments are not provided to `main`. (*End of rationale.*)

`MPI_FINALIZE()`

```
int MPI_Finalize(void)
```

```
MPI_FINALIZE(IERROR)
    INTEGER IERROR
```

This routine cleans up all MPI state. Each process must call `MPI_FINALIZE` before it exits. Unless there has been a call to `MPI_ABORT`, each process must ensure that all pending non-blocking communications are (locally) complete before calling `MPI_FINALIZE`.

Further, at the instant at which the last process calls `MPI_FINALIZE`, all pending sends must be matched by a receive, and all pending receives must be matched by a send.

For example, the following program is correct:

```

Process 0                Process 1
-----                -----
MPI_Init();              MPI_Init();
MPI_Send(dest=1);        MPI_Recv(src=0);
MPI_Finalize();          MPI_Finalize();

```

Without the matching receive, the program is erroneous:

```

Process 0                Process 1
-----                -----
MPI_Init();              MPI_Init();
MPI_Send (dest=1);       MPI_Finalize();
MPI_Finalize();          MPI_Finalize();

```

A successful return from a blocking communication operation or from `MPI_WAIT` or `MPI_TEST` tells the user that the buffer can be reused and means that the communication is completed by the user, but does not guarantee that the local process has no more work to do. A successful return from `MPI_REQUEST_FREE` with a request handle generated by an `MPI_ISEND` nullifies the handle but provides no assurance of operation completion. The `MPI_ISEND` is complete only when it is known by some means that a matching receive has completed. `MPI_FINALIZE` guarantees that all local actions required by communications the user has completed will, in fact, occur before it returns.

`MPI_FINALIZE` guarantees nothing about pending communications that have not been completed (completion is assured only by `MPI_WAIT`, `MPI_TEST`, or `MPI_REQUEST_FREE` combined with some other verification of completion).

Example 7.1 This program is correct:

```

rank 0                rank 1
=====
...
MPI_Isend();           MPI_Recv();
MPI_Request_free();    MPI_Barrier();
MPI_Barrier();         MPI_Finalize();
MPI_Finalize();        exit();
exit();

```

Example 7.2 This program is erroneous and its behavior is undefined:

```

rank 0                rank 1
=====
...
MPI_Isend();           MPI_Recv();
MPI_Request_free();    MPI_Finalize();
MPI_Finalize();        exit();
exit();

```

If no MPI_BUFFER_DETACH occurs between an MPI_BSEND (or other buffered send) and MPI_FINALIZE, the MPI_FINALIZE implicitly supplies the MPI_BUFFER_DETACH.

Example 7.3 This program is correct, and after the MPI_Finalize, it is as if the buffer had been detached.

```

rank 0                                rank 1
=====
...
buffer = malloc(1000000);             MPI_Recv();
MPI_Buffer_attach();                  MPI_Finalize();
MPI_Bsend();                           exit();
MPI_Finalize();
free(buffer);
exit();

```

Example 7.4 In this example, MPI_Iprobe() must return a FALSE flag. MPI_Test_cancelled() must return a TRUE flag, independent of the relative order of execution of MPI_Cancel() in process 0 and MPI_Finalize() in process 1.

The MPI_Iprobe() call is there to make sure the implementation knows that the “tag1” message exists at the destination, without being able to claim that the user knows about it.

```

rank 0                                rank 1
=====
MPI_Init();                           MPI_Init();
MPI_Isend(tag1);                       MPI_Barrier();
MPI_Barrier();                         MPI_Iprobe(tag2);
MPI_Barrier();                         MPI_Barrier();
                                        MPI_Finalize();
                                        exit();

MPI_Cancel();
MPI_Wait();
MPI_Test_cancelled();
MPI_Finalize();
exit();

```

Advice to implementors. An implementation may need to delay the return from MPI_FINALIZE until all potential future message cancellations have been processed. One possible solution is to place a barrier inside MPI_FINALIZE (*End of advice to implementors.*)

Once MPI_FINALIZE returns, no MPI routine (not even MPI_INIT) may be called, except for MPI_GET_VERSION, MPI_INITIALIZED, and the MPI-2 function MPI_FINALIZED. Each process must complete any pending communication it initiated before it calls MPI_FINALIZE. If the call returns, each process may continue local computations, or exit, without participating in further MPI communication with other processes. MPI_FINALIZE is collective on MPI_COMM_WORLD.

1 *Advice to implementors.* Even though a process has completed all the communication
 2 it initiated, such communication may not yet be completed from the viewpoint of the
 3 underlying MPI system. E.g., a blocking send may have completed, even though the
 4 data is still buffered at the sender. The MPI implementation must ensure that a
 5 process has completed any involvement in MPI communication before MPI_FINALIZE
 6 returns. Thus, if a process exits after the call to MPI_FINALIZE, this will not cause
 7 an ongoing communication to fail. (*End of advice to implementors.*)
 8

9 Although it is not required that all processes return from MPI_FINALIZE, it is required
 10 that at least process 0 in MPI_COMM_WORLD return, so that users can know that the MPI
 11 portion of the computation is over. In addition, in a POSIX environment, they may desire
 12 to supply an exit code for each process that returns from MPI_FINALIZE.
 13

14 **Example 7.5** The following illustrates the use of requiring that at least one process return
 15 and that it be known that process 0 is one of the processes that return. One wants code
 16 like the following to work no matter how many processes return.
 17

```
18     ...
19     MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
20     ...
21     MPI_Finalize();
22     if (myrank == 0) {
23         resultfile = fopen("outfile","w");
24         dump_results(resultfile);
25         fclose(resultfile);
26     }
27     exit(0);
```

28
 29
 30 MPI_INITIALIZED(flag)

31 OUT flag Flag is true if MPI_INIT has been called and false
 32 otherwise.
 33

34
 35 int MPI_Initialized(int *flag)

36 MPI_INITIALIZED(FLAG, IERROR)

37 LOGICAL FLAG
 38 INTEGER IERROR
 39

40 This routine may be used to determine whether MPI_INIT has been called. MPI_INITIALIZED
 41 returns true if the calling process has called MPI_INIT. Whether MPI_FINALIZE has been
 42 called does not affect the behavior of MPI_INITIALIZED. It is one of the few routines that
 43 may be called before MPI_INIT is called.
 44
 45
 46
 47
 48

MPI_ABORT(comm, errorcode)			1
IN	comm	communicator of tasks to abort	2
			3
IN	errorcode	error code to return to invoking environment	4
			5

int MPI_Abort(MPI_Comm comm, int errorcode)	6
---	---

MPI_ABORT(COMM, ERRORCODE, IERROR)	7
INTEGER COMM, ERRORCODE, IERROR	8
	9

This routine makes a “best attempt” to abort all tasks 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 .

It may not be possible for an MPI implementation to abort only the processes represented by `comm` if this is a subset of the processes. In this case, the MPI implementation should attempt to abort all the connected processes but should not abort any unconnected processes. If no processes were spawned, accepted or connected then this has the effect of aborting all the processes associated with `MPI_COMM_WORLD`.

Rationale. The communicator argument is provided to allow for future extensions of MPI to environments with, for example, dynamic process management. In particular, it allows but does not require an MPI implementation to abort a subset of `MPI_COMM_WORLD`. (*End of rationale.*)

Advice to users. Whether the errorcode is returned from the executable or from the MPI process startup mechanism (e.g., `mpiexec`), is an aspect of quality of the MPI library but not mandatory. (*End of advice to users.*)

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.*)

Chapter 8

Profiling Interface

8.1 Requirements

To meet the MPI profiling interface, an implementation of the MPI functions *must*

1. provide a mechanism through which all of the MPI defined functions may be accessed with a name shift. Thus all of the MPI functions (which normally start with the prefix “MPI_”) should also be accessible with the prefix “PMPI_”.
2. ensure that those MPI functions which 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 she must implement the profile interface for each binding, or can economise 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 which call the C implementation), ensure that these wrapper functions are separable from the rest of the library.

This is necessary to allow a separate profiling library to be correctly implemented, since (at least with Unix linker semantics) the profiling library must contain these wrapper functions if it is to perform as expected. This requirement allows the person who builds the profiling library to extract these functions from the original MPI library and add them into the profiling library without bringing along any other unnecessary code.

5. provide a no-op routine MPI_PCONTROL in the MPI library.

8.2 Discussion

The objective of the MPI profiling interface is to ensure that it is relatively easy for authors of profiling (and other similar) tools to interface their codes to MPI implementations on different machines.

Since MPI is a machine independent standard with many different implementations, it is unreasonable to expect that the authors of profiling tools for MPI will have access to

the source code which implements MPI on any particular machine. It is therefore necessary to provide a mechanism by which the implementors of such tools can collect whatever performance information they wish *without* access to the underlying implementation.

We believe that having such an interface is important if MPI is to be attractive to end users, since the availability of many different tools will be a significant factor in attracting users to the MPI standard.

The profiling interface is just that, an interface. It says *nothing* about the way in which it is used. There is therefore no attempt to lay down what information is collected through the interface, or how the collected information is saved, filtered, or displayed.

While the initial impetus for the development of this interface arose from the desire to permit the implementation of profiling tools, it is clear that an interface like that specified may also prove useful for other purposes, such as “internetworking” multiple MPI implementations. Since all that is defined is an interface, there is no objection to its being used wherever it is useful.

As the issues being addressed here are intimately tied up with the way in which executable images are built, which may differ greatly on different machines, the examples given below should be treated solely as one way of implementing the objective of the MPI profiling interface. The actual requirements made of an implementation are those detailed in the Requirements section above, the whole of the rest of this chapter is only present as justification and discussion of the logic for those requirements.

The examples below show one way in which an implementation could be constructed to meet the requirements on a Unix system (there are doubtless others which would be equally valid).

8.3 Logic of the design

Provided that an MPI implementation meets the requirements above, it is possible for the implementor of the profiling system to intercept all of the MPI calls which are made by the user program. She can then collect whatever information she requires before calling the underlying MPI implementation (through its name shifted entry points) to achieve the desired effects.

8.3.1 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 is normally used for (at least) the purposes of

- Enabling and disabling profiling depending on the state of the calculation.
- Flushing trace buffers at non-critical points in the calculation
- Adding user events to a trace file.

These requirements are met by use of the MPI_PCONTROL.

```
MPI_PCONTROL(level, ...)
```

```
IN          level          Profiling level
```

```
int MPI_Pcontrol(const int level, ...)
```



```

1 MPI_PCONTROL(LEVEL)
2     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 which 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. (This 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_INIT` has been called 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 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 allows them to modify their source code to obtain more detailed profiling information, but still be able to link exactly the same code against the standard MPI library.

8.4 Examples

8.4.1 Profiler implementation

Suppose that the profiler wishes to accumulate the total amount of data sent by the `MPI_SEND` function, along with the total elapsed time spent in the function. This could trivially be achieved thus

```

34 static int totalBytes;
35 static double totalTime;
36
37
38 int MPI_SEND(void * buffer, const int count, MPI_Datatype datatype,
39             int dest, int tag, MPI_comm comm)
40 {
41     double tstart = MPI_Wtime();      /* Pass on all the arguments */
42     int extent;
43     int result    = PMPI_Send(buffer, count, datatype, dest, tag, comm);
44
45     MPI_Type_size(datatype, &extent); /* Compute size */
46     totalBytes += count*extent;
47
48     totalTime += MPI_Wtime() - tstart; /* and time */

```

```

    return result;
}

```

8.4.2 MPI library implementation

On a Unix system, in which the MPI library is implemented in C, then there are various possible options, of which two of the most obvious are presented here. Which is better depends on whether the linker and compiler support weak symbols.

Systems with weak symbols

If the compiler and linker support weak external symbols (e.g. Solaris 2.x, other system V.4 machines), then only a single library is required through the use of `#pragma weak` thus

```

#pragma weak MPI_Example = PMPI_Example

int PMPI_Example(/* appropriate args */)
{
    /* Useful content */
}

```

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.

Systems without weak symbols

In the absence of weak symbols then one possible solution would be to use the C macro pre-processor thus

```

#ifdef PROFILELIB
#   ifdef __STDC__
#       define FUNCTION(name) P##name
#   else
#       define FUNCTION(name) P/**/name
#   endif
#else
#   define FUNCTION(name) name
#endif

```

Each of the user visible functions in the library would then be declared thus

```

int FUNCTION(MPI_Example)(/* appropriate args */)
{
    /* Useful content */
}

```

1 The same source file can then be compiled to produce both versions of the library,
2 depending on the state of the `PROFILELIB` macro symbol.

3 It is required that the standard MPI library be built in such a way that the inclusion of
4 MPI functions can be achieved one at a time. This is a somewhat unpleasant requirement,
5 since it may mean that each external function has to be compiled from a separate file.
6 However this is necessary so that the author of the profiling library need only define those
7 MPI functions which she wishes to intercept, references to any others being fulfilled by the
8 normal MPI library. Therefore the link step can look something like this

```
9  
10 % cc ... -lmyprof -lpmpi -lmpi
```

11
12 Here `libmyprof.a` contains the profiler functions which intercept some of the MPI
13 functions. `libpmpi.a` contains the “name shifted” MPI functions, and `libmpi.a` contains
14 the normal definitions of the MPI functions.

15 16 8.4.3 Complications

17 Multiple counting

18
19 Since parts of the MPI library may themselves be implemented using more basic MPI func-
20 tions (e.g. a portable implementation of the collective operations implemented using point
21 to point communications), there is potential for profiling functions to be called from within
22 an MPI function which was called from a profiling function. This could lead to “double
23 counting” of the time spent in the inner routine. Since this effect could actually be useful
24 under some circumstances (e.g. it might allow one to answer the question “How much time
25 is spent in the point to point routines when they’re called from collective functions ?”), we
26 have decided not to enforce any restrictions on the author of the MPI library which would
27 overcome this. Therefore the author of the profiling library should be aware of this problem,
28 and guard against it herself. In a single threaded world this is easily achieved through use of
29 a static variable in the profiling code which remembers if you are already inside a profiling
30 routine. It becomes more complex in a multi-threaded environment (as does the meaning
31 of the times recorded !)

32 33 Linker oddities

34
35 The Unix linker traditionally operates in one pass : the effect of this is that functions from
36 libraries are only included in the image if they are needed at the time the library is scanned.
37 When combined with weak symbols, or multiple definitions of the same function, this can
38 cause odd (and unexpected) effects.

39 Consider, for instance, an implementation of MPI in which the Fortran binding is
40 achieved by using wrapper functions on top of the C implementation. The author of the
41 profile library then assumes that it is reasonable only to provide profile functions for the C
42 binding, since Fortran will eventually call these, and the cost of the wrappers is assumed
43 to be small. However, if the wrapper functions are not in the profiling library, then none
44 of the profiled entry points will be undefined when the profiling library is called. Therefore
45 none of the profiling code will be included in the image. When the standard MPI library
46 is scanned, the Fortran wrappers will be resolved, and will also pull in the base versions of
47 the MPI functions. The overall effect is that the code will link successfully, but will not be
48 profiled.

To overcome this we must ensure that the Fortran wrapper functions are included in the profiling version of the library. We ensure that this is possible by requiring that these be separable from the rest of the base MPI library. This allows them to be aared out of the base library and into the profiling one.

8.5 Multiple levels of interception

The scheme given here does not directly support the nesting of profiling functions, since it provides only a single alternative name for each MPI function. Consideration was given to an implementation which would allow multiple levels of call interception, however we were unable to construct an implementation of this which did not have the following disadvantages

- assuming a particular implementation language.
- imposing a run time cost even when no profiling was taking place.

Since one of the objectives of MPI is to permit efficient, low latency implementations, and it is not the business of a standard to require a particular implementation language, we decided to accept the scheme outlined above.

Note, however, that it is possible to use the scheme above to implement a multi-level system, since the function called by the user may call many different profiling functions before calling the underlying MPI function.

Unfortunately such an implementation may require more cooperation between the different profiling libraries than is required for the single level implementation detailed above.

Bibliography

- [1] 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. [1.1](#)
- [2] 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. [1.1](#)
- [3] Purushotham V. Bangalore, Nathan E. Doss, and Anthony Skjellum. MPI++: Issues and Features. In *OON-SKI '94*, page in press, 1994. [5.1](#)
- [4] 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. [1.1](#)
- [5] 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. [1.1](#), [6.2](#)
- [6] R. Butler and E. Lusk. User’s guide to the p4 programming system. Technical Report TM-ANL-92/17, Argonne National Laboratory, 1992. [1.1](#)
- [7] Ralph Butler and Ewing Lusk. Monitors, messages, and clusters: the p4 parallel programming system. *Journal of Parallel Computing*, 1994. to appear (Also Argonne National Laboratory Mathematics and Computer Science Division preprint P362-0493). [1.1](#)
- [8] Robin Calkin, Rolf Hempel, Hans-Christian Hoppe, and Peter Wypior. Portable programming with the parmacs message-passing library. *Parallel Computing, Special issue on message-passing interfaces*, to appear. [1.1](#), [6.2](#)
- [9] 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. [6.1](#)
- [10] 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. [6.1](#)
- [11] 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. [1.1](#)

- [12] 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. [1.1](#)
- [13] Nathan Doss, William Gropp, Ewing Lusk, and Anthony Skjellum. A model implementation of MPI. Technical report, Argonne National Laboratory, 1993. [1.3](#)
- [14] Edinburgh Parallel Computing Centre, University of Edinburgh. *CHIMP Concepts*, June 1991. [1.1](#)
- [15] Edinburgh Parallel Computing Centre, University of Edinburgh. *CHIMP Version 1.0 Interface*, May 1992. [1.1](#)
- [16] D. Feitelson. Communicators: Object-based multiparty interactions for parallel programming. Technical Report 91-12, Dept. Computer Science, The Hebrew University of Jerusalem, November 1991. [5.1.2](#)
- [17] Hubertus Franke, Peter Hochschild, Pratap Pattnaik, and Marc Snir. An efficient implementation of MPI. In *1994 International Conference on Parallel Processing*, 1994. [1.3](#)
- [18] G. A. Geist, M. T. Heath, B. W. Peyton, and P. H. Worley. A user's guide to PICL: a portable instrumented communication library. Technical Report TM-11616, Oak Ridge National Laboratory, October 1990. [1.1](#)
- [19] William D. Gropp and Barry Smith. Chameleon parallel programming tools users manual. Technical Report ANL-93/23, Argonne National Laboratory, March 1993. [1.1](#)
- [20] O. Krämer and H. Mühlenbein. Mapping strategies in message-based multiprocessor systems. *Parallel Computing*, 9:213–225, 1989. [6.1](#)
- [21] nCUBE Corporation. *nCUBE 2 Programmers Guide, r2.0*, December 1990. [1.1](#)
- [22] Parasoft Corporation, Pasadena, CA. *Express User's Guide*, version 3.2.5 edition, 1992. [1.1](#), [6.4](#)
- [23] 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. [1.1](#)
- [24] 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. [1.1](#), [5.1.2](#)
- [25] A. Skjellum, S. Smith, C. Still, A. Leung, and M. Morari. The Zipcode message passing system. Technical report, Lawrence Livermore National Laboratory, September 1992. [1.1](#)
- [26] Anthony Skjellum, Nathan E. Doss, and Purushotham V. Bangalore. Writing Libraries in MPI. In Anthony Skjellum and Donna S. Reese, editors, *Proceedings of the Scalable Parallel Libraries Conference*, pages 166–173. IEEE Computer Society Press, October 1993. [5.1](#)

- 1 [27] Anthony Skjellum, Steven G. Smith, Nathan E. Doss, Alvin P . Leung, and Manfred
2 Morari. The Design and Evolution of Zipcode. *Parallel Computing*, 1994. (Invited
3 Paper, to appear in Special Issue on Message Passing). [5.1.2](#)
4
- 5 [28] Anthony Skjellum, Steven G. Smith, Nathan E. Doss, Charles H. Still, Alvin P. Le-
6 ung, and Manfred Morari. Zipcode: A Portable Communication Layer for High Per-
7 formance Multicomputing. Technical Report UCRL-JC-106725 (revised 9/92, 12/93,
8 4/94), Lawrence Livermore National Laboratory, March 1991. To appear in *Concur-*
9 *rency: Practice & Experience*. [5.5.6](#)
- 10 [29] D. Walker. Standards for message passing in a distributed memory environment. Tech-
11 nical Report TM-12147, Oak Ridge National Laboratory, August 1992. [1.1](#)
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

Annex A

Language Binding

A.1 Introduction

In this section we summarize the specific bindings for both Fortran and C. We present first the C bindings, then the Fortran bindings. Listings are alphabetical within chapter.

A.2 Defined Constants for C and Fortran

These are required defined constants, to be defined in the files `mpi.h` (for C) and `mpif.h` (for Fortran).

```
/* return codes (both C and Fortran) */
```

```
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
```

```
MPI_ERR_IN_STATUS
```

```
MPI_ERR_LASTCODE
```

```
/* assorted constants (both C and Fortran) */
```

```
MPI_BOTTOM
```

```
MPI_PROC_NULL
```

```
MPI_ANY_SOURCE
```

```
MPI_ANY_TAG
```



```
1  MPI_UNDEFINED
2  MPI_BSEND_OVERHEAD
3  MPI_KEYVAL_INVALID
4
5  /* status size and reserved index values (Fortran) */
6  MPI_STATUS_SIZE
7  MPI_SOURCE
8  MPI_TAG
9  MPI_ERROR
10
11 /* Error-handling specifiers (C and Fortran) */
12 MPI_ERRORS_ARE_FATAL
13 MPI_ERRORS_RETURN
14
15 /* Maximum sizes for strings */
16 MPI_MAX_PROCESSOR_NAME
17 MPI_MAX_ERROR_STRING
18
19 /* elementary datatypes (C) */
20 MPI_CHAR
21 MPI_SHORT
22 MPI_INT
23 MPI_LONG
24 MPI_UNSIGNED_CHAR
25 MPI_UNSIGNED_SHORT
26 MPI_UNSIGNED
27 MPI_UNSIGNED_LONG
28 MPI_FLOAT
29 MPI_DOUBLE
30 MPI_LONG_DOUBLE
31 MPI_BYTE
32 MPI_PACKED
33
34
35
36 /* elementary datatypes (Fortran) */
37 MPI_INTEGER
38 MPI_REAL
39 MPI_DOUBLE_PRECISION
40 MPI_COMPLEX
41 MPI_LOGICAL
42 MPI_CHARACTER
43 MPI_BYTE
44 MPI_PACKED
45
46 /* datatypes for reduction functions (C) */
47 MPI_FLOAT_INT
48 MPI_DOUBLE_INT
```

MPI_LONG_INT	1
MPI_2INT	2
MPI_SHORT_INT	3
MPI_LONG_DOUBLE_INT	4
	5
/* datatypes for reduction functions (Fortran) */	6
MPI_2REAL	7
MPI_2DOUBLE_PRECISION	8
MPI_2INTEGER	9
	10
/* optional datatypes (Fortran) */	11
MPI_DOUBLE_COMPLEX	12
MPI_INTEGER1	13
MPI_INTEGER2	14
MPI_INTEGER4	15
MPI_REAL2	16
MPI_REAL4	17
MPI_REAL8	18
etc.	19
	20
/* optional datatypes (C) */	21
MPI_LONG_LONG_INT	22
etc.	23
	24
/* special datatypes for constructing derived datatypes	25
MPI_UB	26
MPI_LB	27
	28
/* reserved communicators (C and Fortran) */	29
MPI_COMM_WORLD	30
MPI_COMM_SELF	31
	32
/* results of communicator and group comparisons */	33
	34
MPI_IDENT	35
MPI_CONGRUENT	36
MPI_SIMILAR	37
MPI_UNEQUAL	38
	39
/* environmental inquiry keys (C and Fortran) */	40
MPI_TAG_UB	41
MPI_IO	42
MPI_HOST	43
MPI_WTIME_IS_GLOBAL	44
	45
/* collective operations (C and Fortran) */	46
MPI_MAX	47
MPI_MIN	48

```
1  MPI_SUM
2  MPI_PROD
3  MPI_MAXLOC
4  MPI_MINLOC
5  MPI_BAND
6  MPI_BOR
7  MPI_BXOR
8  MPI_LAND
9  MPI_LOR
10 MPI_LXOR
11
12 /* Null handles */
13 MPI_GROUP_NULL
14 MPI_COMM_NULL
15 MPI_DATATYPE_NULL
16 MPI_REQUEST_NULL
17 MPI_OP_NULL
18 MPI_ERRHANDLER_NULL
19
20 /* Empty group */
21 MPI_GROUP_EMPTY
22
23 /* topologies (C and Fortran) */
24 MPI_GRAPH
25 MPI_CART
26
27 /* Predefined functions in C and Fortran */
28 MPI_NULL_COPY_FN
29 MPI_NULL_DELETE_FN
30 MPI_DUP_FN
31
32
33     The following are defined C type definitions, also included in the file mpi.h.
34
35 /* opaque types (C) */
36 MPI_Aint
37 MPI_Status
38
39 /* handles to assorted structures (C) */
40 MPI_Group
41 MPI_Comm
42 MPI_Datatype
43 MPI_Request
44 MPI_Op
45 MPI_Errhandler
46
47 /* prototypes for user-defined functions (C) */
48 typedef int MPI_Copy_function(MPI_Comm oldcomm, int keyval,
```

```

        void *extra_state, void *attribute_val_in,
        void *attribute_val_out, int *flag);
typedef int MPI_Delete_function(MPI_Comm comm, int keyval,
        void *attribute_val, void *extra_state)
typedef void MPI_Handler_function(MPI_Comm *, int *, ...);
typedef void MPI_User_function( void *invec, void *inoutvec, int *len,
        MPI_Datatype *datatype);

```

For Fortran, here are examples of how each of the user-defined functions should be declared.

The user-function argument to MPI_OP_CREATE should be declared like this:

```

SUBROUTINE USER_FUNCTION( INVEC, INOUTVEC, LEN, TYPE)
<type> INVEC(LEN), INOUTVEC(LEN)
INTEGER LEN, TYPE

```

The copy-function argument to MPI_KEYVAL_CREATE should be declared like this:

```

SUBROUTINE COPY_FUNCTION(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

```

The delete-function argument to MPI_KEYVAL_CREATE should be declared like this:

```

SUBROUTINE DELETE_FUNCTION(COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR)
INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, EXTRA_STATE, IERR

```

The handler-function for error handlers should be declared like this:

```

SUBROUTINE HANDLER_FUNCTION(COMM, ERROR_CODE, ..... )
INTEGER COMM, ERROR_CODE

```

A.3 C bindings for Point-to-Point Communication

These are presented here in the order of their appearance in the chapter.

```

int MPI_Send(void* buf, int count, MPI_Datatype datatype, int dest,
        int tag, MPI_Comm comm)
int MPI_Recv(void* buf, int count, MPI_Datatype datatype, int source,
        int tag, MPI_Comm comm, MPI_Status *status)
int MPI_Get_count(MPI_Status *status, MPI_Datatype datatype, int *count)
int MPI_Bsend(void* buf, int count, MPI_Datatype datatype, int dest,
        int tag, MPI_Comm comm)
int MPI_Ssend(void* buf, int count, MPI_Datatype datatype, int dest,
        int tag, MPI_Comm comm)
int MPI_Rsend(void* buf, int count, MPI_Datatype datatype, int dest,
        int tag, MPI_Comm comm)

```

```
1 int MPI_Buffer_attach( void* buffer, int size)
2
3 int MPI_Buffer_detach( void* buffer, int* size)
4
5 int MPI_Isend(void* buf, int count, MPI_Datatype datatype, int dest,
6               int tag, MPI_Comm comm, MPI_Request *request)
7
8 int MPI_Ibsend(void* buf, int count, MPI_Datatype datatype, int dest,
9               int tag, MPI_Comm comm, MPI_Request *request)
10
11 int MPI_Issend(void* buf, int count, MPI_Datatype datatype, int dest,
12               int tag, MPI_Comm comm, MPI_Request *request)
13
14 int MPI_Irsend(void* buf, int count, MPI_Datatype datatype, int dest,
15               int tag, MPI_Comm comm, MPI_Request *request)
16
17 int MPI_Irecv(void* buf, int count, MPI_Datatype datatype, int source,
18               int tag, MPI_Comm comm, MPI_Request *request)
19
20 int MPI_Wait(MPI_Request *request, MPI_Status *status)
21
22 int MPI_Test(MPI_Request *request, int *flag, MPI_Status *status)
23
24 int MPI_Request_free(MPI_Request *request)
25
26 int MPI_Waitany(int count, MPI_Request *array_of_requests, int *index,
27                MPI_Status *status)
28
29 int MPI_Testany(int count, MPI_Request *array_of_requests, int *index,
30                int *flag, MPI_Status *status)
31
32 int MPI_Waitall(int count, MPI_Request *array_of_requests,
33                MPI_Status *array_of_statuses)
34
35 int MPI_Testall(int count, MPI_Request *array_of_requests, int *flag,
36                MPI_Status *array_of_statuses)
37
38 int MPI_Waitsome(int incount, MPI_Request *array_of_requests, int *outcount,
39                 int *array_of_indices, MPI_Status *array_of_statuses)
40
41 int MPI_Testsome(int incount, MPI_Request *array_of_requests, int *outcount,
42                 int *array_of_indices, MPI_Status *array_of_statuses)
43
44 int MPI_Iprobe(int source, int tag, MPI_Comm comm, int *flag,
45                MPI_Status *status)
46
47 int MPI_Probe(int source, int tag, MPI_Comm comm, MPI_Status *status)
48
49 int MPI_Cancel(MPI_Request *request)
50
51 int MPI_Test_cancelled(MPI_Status *status, int *flag)
52
53 int MPI_Send_init(void* buf, int count, MPI_Datatype datatype, int dest,
54                  int tag, MPI_Comm comm, MPI_Request *request)
55
56 int MPI_Bsend_init(void* buf, int count, MPI_Datatype datatype, int dest,
57                   int tag, MPI_Comm comm, MPI_Request *request)
```



```
1 int MPI_Get_elements(MPI_Status *status, MPI_Datatype datatype, int *count)
2
3 int MPI_Pack(void* inbuf, int incount, MPI_Datatype datatype, void *outbuf,
4             int outsize, int *position, MPI_Comm comm)
5
6 int MPI_Unpack(void* inbuf, int insize, int *position, void *outbuf,
7              int outcount, MPI_Datatype datatype, MPI_Comm comm)
8
9 int MPI_Pack_size(int incount, MPI_Datatype datatype, MPI_Comm comm,
10                 int *size)
```

A.4 C Bindings for Collective Communication

```
13 int MPI_Barrier(MPI_Comm comm )
14
15 int MPI_Bcast(void* buffer, int count, MPI_Datatype datatype, int root,
16             MPI_Comm comm )
17
18 int MPI_Gather(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_Gatherv(void* sendbuf, int sendcount, MPI_Datatype sendtype,
23               void* recvbuf, int *recvcounts, int *displs,
24               MPI_Datatype recvtype, int root, MPI_Comm comm)
25
26 int MPI_Scatter(void* sendbuf, int sendcount, MPI_Datatype sendtype,
27               void* recvbuf, int recvcount, MPI_Datatype recvtype, int root,
28               MPI_Comm comm)
29
30 int MPI_Scatterv(void* sendbuf, int *sendcounts, int *displs,
31                 MPI_Datatype sendtype, void* recvbuf, int recvcount,
32                 MPI_Datatype recvtype, int root, MPI_Comm comm)
33
34 int MPI_Allgather(void* sendbuf, int sendcount, MPI_Datatype sendtype,
35                  void* recvbuf, int recvcount, MPI_Datatype recvtype,
36                  MPI_Comm comm)
37
38 int MPI_Allgatherv(void* sendbuf, int sendcount, MPI_Datatype sendtype,
39                   void* recvbuf, int *recvcounts, int *displs,
40                   MPI_Datatype recvtype, MPI_Comm comm)
41
42 int MPI_Alltoall(void* sendbuf, int sendcount, MPI_Datatype sendtype,
43                 void* recvbuf, int recvcount, MPI_Datatype recvtype,
44                 MPI_Comm comm)
45
46 int MPI_Alltoallv(void* sendbuf, int *sendcounts, int *sdispls,
47                  MPI_Datatype sendtype, void* recvbuf, int *recvcounts,
48                  int *rdispls, MPI_Datatype recvtype, MPI_Comm comm)
```



```

1  int MPI_Comm_free(MPI_Comm *comm)
2
3  int MPI_Comm_test_inter(MPI_Comm comm, int *flag)
4
5  int MPI_Comm_remote_size(MPI_Comm comm, int *size)
6
7  int MPI_Comm_remote_group(MPI_Comm comm, MPI_Group *group)
8
9  int MPI_Intercomm_create(MPI_Comm local_comm, int local_leader,
10                          MPI_Comm peer_comm, int remote_leader, int tag,
11                          MPI_Comm *newintercomm)
12
13 int MPI_Intercomm_merge(MPI_Comm intercomm, int high,
14                          MPI_Comm *newintracomm)
15
16 int MPI_Keyval_create(MPI_Copy_function *copy_fn, MPI_Delete_function
17                       *delete_fn, int *keyval, void* extra_state)
18
19 int MPI_Keyval_free(int *keyval)
20
21 int MPI_Attr_put(MPI_Comm comm, int keyval, void* attribute_val)
22
23 int MPI_Attr_get(MPI_Comm comm, int keyval, void* attribute_val, int *flag)
24
25 int MPI_Attr_delete(MPI_Comm comm, int keyval)

```

A.6 C Bindings for Process Topologies

```

25 int MPI_Cart_create(MPI_Comm comm_old, int ndims, int *dims, int *periods,
26                    int reorder, MPI_Comm *comm_cart)
27
28 int MPI_Dims_create(int nnodes, int ndims, int *dims)
29
30 int MPI_Graph_create(MPI_Comm comm_old, int nnodes, int *index, int *edges,
31                     int reorder, MPI_Comm *comm_graph)
32
33 int MPI_Topo_test(MPI_Comm comm, int *status)
34
35 int MPI_Graphdims_get(MPI_Comm comm, int *nnodes, int *nedges)
36
37 int MPI_Graph_get(MPI_Comm comm, int maxindex, int maxedges, int *index,
38                  int *edges)
39
40 int MPI_Cartdim_get(MPI_Comm comm, int *ndims)
41
42 int MPI_Cart_get(MPI_Comm comm, int maxdims, int *dims, int *periods,
43                 int *coords)
44
45 int MPI_Cart_rank(MPI_Comm comm, int *coords, int *rank)
46
47 int MPI_Cart_coords(MPI_Comm comm, int rank, int maxdims, int *coords)
48
49 int MPI_Graph_neighbors_count(MPI_Comm comm, int rank, int *nneighbors)
50
51 int MPI_Graph_neighbors(MPI_Comm comm, int rank, int maxneighbors,
52                          int *neighbors)

```

```

int MPI_Cart_shift(MPI_Comm comm, int direction, int disp, int *rank_source,
                  int *rank_dest)
int MPI_Cart_sub(MPI_Comm comm, int *remain_dims, MPI_Comm *newcomm)
int MPI_Cart_map(MPI_Comm comm, int ndims, int *dims, int *periods,
                 int *newrank)
int MPI_Graph_map(MPI_Comm comm, int nnodes, int *index, int *edges,
                  int *newrank)

```

A.7 C bindings for Environmental Inquiry

```

int MPI_Get_processor_name(char *name, int *resultlen)
int MPI_Errhandler_create(MPI_Handler_function *function,
                           MPI_Errhandler *errhandler)
int MPI_Errhandler_set(MPI_Comm comm, MPI_Errhandler errhandler)
int MPI_Errhandler_get(MPI_Comm comm, MPI_Errhandler *errhandler)
int MPI_Errhandler_free(MPI_Errhandler *errhandler)
int MPI_Error_string(int errorcode, char *string, int *resultlen)
int MPI_Error_class(int errorcode, int *errorclass)
double MPI_Wtime(void)
double MPI_Wtick(void)
int MPI_Init(int *argc, char ***argv)
int MPI_Finalize(void)
int MPI_Initialized(int *flag)
int MPI_Abort(MPI_Comm comm, int errorcode)
int MPI_Get_version(int *version, int *subversion)

```

A.8 C Bindings for Profiling

```

int MPI_Pcontrol(const int level, ...)

```

A.9 Fortran Bindings for Point-to-Point Communication

```

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

```

```
1      <type> BUF(*)
2      INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE),
3      IERROR
4
5      MPI_GET_COUNT(STATUS, DATATYPE, COUNT, IERROR)
6      INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR
7
8      MPI_BSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
9      <type> BUF(*)
10     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR
11
12     MPI_SSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
13     <type> BUF(*)
14     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR
15
16     MPI_RSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, IERROR)
17     <type> BUF(*)
18     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, IERROR
19
20     MPI_BUFFER_ATTACH( BUFFER, SIZE, IERROR)
21     <type> BUFFER(*)
22     INTEGER SIZE, IERROR
23
24     MPI_BUFFER_DETACH( BUFFER, SIZE, IERROR)
25     <type> BUFFER(*)
26     INTEGER SIZE, IERROR
27
28     MPI_ISEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
29     <type> BUF(*)
30     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
31
32     MPI_IBSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
33     <type> BUF(*)
34     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
35
36     MPI_ISSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
37     <type> BUF(*)
38     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
39
40     MPI_IRSEND(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
41     <type> BUF(*)
42     INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
43
44     MPI_WAIT(REQUEST, STATUS, IERROR)
45     INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR
46
47     MPI_TEST(REQUEST, FLAG, STATUS, IERROR)
48     LOGICAL FLAG
49     INTEGER REQUEST, STATUS(MPI_STATUS_SIZE), IERROR
```

MPI_REQUEST_FREE(REQUEST, IERROR)	1
INTEGER REQUEST, IERROR	2
	3
MPI_WAITANY(COUNT, ARRAY_OF_REQUESTS, INDEX, STATUS, IERROR)	4
INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE),	5
IERROR	6
	7
MPI_TESTANY(COUNT, ARRAY_OF_REQUESTS, INDEX, FLAG, STATUS, IERROR)	8
LOGICAL FLAG	9
INTEGER COUNT, ARRAY_OF_REQUESTS(*), INDEX, STATUS(MPI_STATUS_SIZE),	10
IERROR	11
	12
MPI_WAITALL(COUNT, ARRAY_OF_REQUESTS, ARRAY_OF_STATUSES, IERROR)	13
INTEGER COUNT, ARRAY_OF_REQUESTS(*),	14
ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), IERROR	15
	16
MPI_TESTALL(COUNT, ARRAY_OF_REQUESTS, FLAG, ARRAY_OF_STATUSES, IERROR)	17
LOGICAL FLAG	18
INTEGER COUNT, ARRAY_OF_REQUESTS(*),	19
ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), IERROR	20
	21
MPI_WAITSSOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT, ARRAY_OF_INDICES,	22
ARRAY_OF_STATUSES, IERROR)	23
INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),	24
ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), IERROR	25
	26
MPI_TESTSSOME(INCOUNT, ARRAY_OF_REQUESTS, OUTCOUNT, ARRAY_OF_INDICES,	27
ARRAY_OF_STATUSES, IERROR)	28
INTEGER INCOUNT, ARRAY_OF_REQUESTS(*), OUTCOUNT, ARRAY_OF_INDICES(*),	29
ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), IERROR	30
	31
MPI_IPROBE(SOURCE, TAG, COMM, FLAG, STATUS, IERROR)	32
LOGICAL FLAG	33
INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR	34
	35
MPI_PROBE(SOURCE, TAG, COMM, STATUS, IERROR)	36
INTEGER SOURCE, TAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR	37
	38
MPI_CANCEL(REQUEST, IERROR)	39
INTEGER REQUEST, IERROR	40
	41
MPI_TEST_CANCELLED(STATUS, FLAG, IERROR)	42
LOGICAL FLAG	43
INTEGER STATUS(MPI_STATUS_SIZE), IERROR	44
	45
MPI_SEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	46
<type> BUF(*)	47
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR	48
	49
MPI_BSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	50
<type> BUF(*)	51
INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR	52
	53
MPI_SSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)	54

```
1      <type> BUF(*)
2      INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
3
4      MPI_RSEND_INIT(BUF, COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR)
5      <type> BUF(*)
6      INTEGER COUNT, DATATYPE, DEST, TAG, COMM, REQUEST, IERROR
7
8      MPI_RECV_INIT(BUF, COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR)
9      <type> BUF(*)
10     INTEGER COUNT, DATATYPE, SOURCE, TAG, COMM, REQUEST, IERROR
11
12     MPI_START(REQUEST, IERROR)
13     INTEGER REQUEST, IERROR
14
15     MPI_STARTALL(COUNT, ARRAY_OF_REQUESTS, IERROR)
16     INTEGER COUNT, ARRAY_OF_REQUESTS(*), IERROR
17
18     MPI_SENDRECV(SENDBUF, SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVBUF,
19     RECVCOUNT, RECVMODE, SOURCE, RECVTAG, COMM, STATUS, IERROR)
20     <type> SENDBUF(*), RECVBUF(*)
21     INTEGER SENDCOUNT, SENDTYPE, DEST, SENDTAG, RECVCOUNT, RECVMODE,
22     SOURCE, RECVTAG, COMM, STATUS(MPI_STATUS_SIZE), IERROR
23
24     MPI_SENDRECV_REPLACE(BUF, COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG,
25     COMM, STATUS, IERROR)
26     <type> BUF(*)
27     INTEGER COUNT, DATATYPE, DEST, SENDTAG, SOURCE, RECVTAG, COMM,
28     STATUS(MPI_STATUS_SIZE), IERROR
29
30     MPI_TYPE_CONTIGUOUS(COUNT, OLDTYPE, NEWTYPE, IERROR)
31     INTEGER COUNT, OLDTYPE, NEWTYPE, IERROR
32
33     MPI_TYPE_VECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR)
34     INTEGER COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR
35
36     MPI_TYPE_HVECTOR(COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR)
37     INTEGER COUNT, BLOCKLENGTH, STRIDE, OLDTYPE, NEWTYPE, IERROR
38
39     MPI_TYPE_INDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
40     OLDTYPE, NEWTYPE, IERROR)
41     INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*),
42     OLDTYPE, NEWTYPE, IERROR
43
44     MPI_TYPE_HINDEXED(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
45     OLDTYPE, NEWTYPE, IERROR)
46     INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*),
47     OLDTYPE, NEWTYPE, IERROR
48
49     MPI_TYPE_STRUCT(COUNT, ARRAY_OF_BLOCKLENGTHS, ARRAY_OF_DISPLACEMENTS,
50     ARRAY_OF_TYPES, NEWTYPE, IERROR)
51     INTEGER COUNT, ARRAY_OF_BLOCKLENGTHS(*), ARRAY_OF_DISPLACEMENTS(*),
52     ARRAY_OF_TYPES(*), NEWTYPE, IERROR
53
54     MPI_ADDRESS(LOCATION, ADDRESS, IERROR)
```

<type> LOCATION(*)	1
INTEGER ADDRESS, IERROR	2
	3
MPI_TYPE_EXTENT(DATATYPE, EXTENT, IERROR)	4
INTEGER DATATYPE, EXTENT, IERROR	5
	6
MPI_TYPE_SIZE(DATATYPE, SIZE, IERROR)	7
INTEGER DATATYPE, SIZE, IERROR	8
	9
MPI_TYPE_LB(DATATYPE, DISPLACEMENT, IERROR)	10
INTEGER DATATYPE, DISPLACEMENT, IERROR	11
	12
MPI_TYPE_UB(DATATYPE, DISPLACEMENT, IERROR)	13
INTEGER DATATYPE, DISPLACEMENT, IERROR	14
	15
MPI_TYPE_COMMIT(DATATYPE, IERROR)	16
INTEGER DATATYPE, IERROR	17
	18
MPI_TYPE_FREE(DATATYPE, IERROR)	19
INTEGER DATATYPE, IERROR	20
	21
MPI_GET_ELEMENTS(STATUS, DATATYPE, COUNT, IERROR)	22
INTEGER STATUS(MPI_STATUS_SIZE), DATATYPE, COUNT, IERROR	23
	24
MPI_PACK(INBUF, INCOUNT, DATATYPE, OUTBUF, OUTSIZE, POSITION, COMM, IERROR)	25
<type> INBUF(*), OUTBUF(*)	26
INTEGER INCOUNT, DATATYPE, OUTSIZE, POSITION, COMM, IERROR	27
	28
MPI_UNPACK(INBUF, INSIZE, POSITION, OUTBUF, OUTCOUNT, DATATYPE, COMM,	29
IERROR)	30
<type> INBUF(*), OUTBUF(*)	31
INTEGER INSIZE, POSITION, OUTCOUNT, DATATYPE, COMM, IERROR	32
	33
MPI_PACK_SIZE(INCOUNT, DATATYPE, COMM, SIZE, IERROR)	34
INTEGER INCOUNT, DATATYPE, COMM, SIZE, IERROR	35
	36

A.10 Fortran Bindings for Collective Communication

MPI_BARRIER(COMM, IERROR)	37
INTEGER COMM, IERROR	38
	39
MPI_BCAST(BUFFER, COUNT, DATATYPE, ROOT, COMM, IERROR)	40
<type> BUFFER(*)	41
INTEGER COUNT, DATATYPE, ROOT, COMM, IERROR	42
	43
MPI_GATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPER,	44
ROOT, COMM, IERROR)	45
<type> SENDBUF(*), RECVBUF(*)	46
INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPER, ROOT, COMM, IERROR	47
	48
MPI_GATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,	49
RECVTYPER, ROOT, COMM, IERROR)	50
<type> SENDBUF(*), RECVBUF(*)	51

```
1     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, ROOT,
2     COMM, IERROR
3
4     MPI_SCATTER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
5     ROOT, COMM, IERROR)
6     <type> SENDBUF(*), RECVBUF(*)
7     INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, ROOT, COMM, IERROR
8
9     MPI_SCATTERV(SENDBUF, SENDCOUNTS, DISPLS, SENDTYPE, RECVBUF, RECVCOUNT,
10    RECVTYPE, ROOT, COMM, IERROR)
11    <type> SENDBUF(*), RECVBUF(*)
12    INTEGER SENDCOUNTS(*), DISPLS(*), SENDTYPE, RECVCOUNT, RECVTYPE, ROOT,
13    COMM, IERROR
14
15    MPI_ALLGATHER(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
16    COMM, IERROR)
17    <type> SENDBUF(*), RECVBUF(*)
18    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, IERROR
19
20    MPI_ALLGATHERV(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNTS, DISPLS,
21    RECVTYPE, COMM, IERROR)
22    <type> SENDBUF(*), RECVBUF(*)
23    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNTS(*), DISPLS(*), RECVTYPE, COMM,
24    IERROR
25
26    MPI_ALLTOALL(SENDBUF, SENDCOUNT, SENDTYPE, RECVBUF, RECVCOUNT, RECVTYPE,
27    COMM, IERROR)
28    <type> SENDBUF(*), RECVBUF(*)
29    INTEGER SENDCOUNT, SENDTYPE, RECVCOUNT, RECVTYPE, COMM, IERROR
30
31    MPI_ALLTOALLV(SENDBUF, SENDCOUNTS, SDISPLS, SENDTYPE, RECVBUF, RECVCOUNTS,
32    RDISPLS, RECVTYPE, COMM, IERROR)
33    <type> SENDBUF(*), RECVBUF(*)
34    INTEGER SENDCOUNTS(*), SDISPLS(*), SENDTYPE, RECVCOUNTS(*), RDISPLS(*),
35    RECVTYPE, COMM, IERROR
36
37    MPI_REDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, ROOT, COMM, IERROR)
38    <type> SENDBUF(*), RECVBUF(*)
39    INTEGER COUNT, DATATYPE, OP, ROOT, COMM, IERROR
40
41    MPI_OP_CREATE( FUNCTION, COMMUTE, OP, IERROR)
42    EXTERNAL FUNCTION
43    LOGICAL COMMUTE
44    INTEGER OP, IERROR
45
46    MPI_OP_FREE( OP, IERROR)
47    INTEGER OP, IERROR
48
49    MPI_ALLREDUCE(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)
50    <type> SENDBUF(*), RECVBUF(*)
51    INTEGER COUNT, DATATYPE, OP, COMM, IERROR
52
53    MPI_REDUCE_SCATTER(SENDBUF, RECVBUF, RECVCOUNTS, DATATYPE, OP, COMM,
```

IERROR)	1
<type> SENDBUF(*), RECVBUF(*)	2
INTEGER RECVCOUNTS(*), DATATYPE, OP, COMM, IERROR	3
MPI_SCAN(SENDBUF, RECVBUF, COUNT, DATATYPE, OP, COMM, IERROR)	4
<type> SENDBUF(*), RECVBUF(*)	5
INTEGER COUNT, DATATYPE, OP, COMM, IERROR	6
	7
	8
	9
A.11 Fortran Bindings for Groups, Contexts, etc.	10
	11
MPI_GROUP_SIZE(GROUP, SIZE, IERROR)	12
INTEGER GROUP, SIZE, IERROR	13
MPI_GROUP_RANK(GROUP, RANK, IERROR)	14
INTEGER GROUP, RANK, IERROR	15
MPI_GROUP_TRANSLATE_RANKS(GROUP1, N, RANKS1, GROUP2, RANKS2, IERROR)	16
INTEGER GROUP1, N, RANKS1(*), GROUP2, RANKS2(*), IERROR	17
	18
MPI_GROUP_COMPARE(GROUP1, GROUP2, RESULT, IERROR)	19
INTEGER GROUP1, GROUP2, RESULT, IERROR	20
	21
MPI_COMM_GROUP(COMM, GROUP, IERROR)	22
INTEGER COMM, GROUP, IERROR	23
MPI_GROUP_UNION(GROUP1, GROUP2, NEWGROUP, IERROR)	24
INTEGER GROUP1, GROUP2, NEWGROUP, IERROR	25
	26
MPI_GROUP_INTERSECTION(GROUP1, GROUP2, NEWGROUP, IERROR)	27
INTEGER GROUP1, GROUP2, NEWGROUP, IERROR	28
MPI_GROUP_DIFFERENCE(GROUP1, GROUP2, NEWGROUP, IERROR)	29
INTEGER GROUP1, GROUP2, NEWGROUP, IERROR	30
	31
MPI_GROUP_INCL(GROUP, N, RANKS, NEWGROUP, IERROR)	32
INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR	33
MPI_GROUP_EXCL(GROUP, N, RANKS, NEWGROUP, IERROR)	34
INTEGER GROUP, N, RANKS(*), NEWGROUP, IERROR	35
	36
MPI_GROUP_RANGE_INCL(GROUP, N, RANGES, NEWGROUP, IERROR)	37
INTEGER GROUP, N, RANGES(3,*), NEWGROUP, IERROR	38
MPI_GROUP_RANGE_EXCL(GROUP, N, RANGES, NEWGROUP, IERROR)	39
INTEGER GROUP, N, RANGES(3,*), NEWGROUP, IERROR	40
	41
MPI_GROUP_FREE(GROUP, IERROR)	42
INTEGER GROUP, IERROR	43
	44
MPI_COMM_SIZE(COMM, SIZE, IERROR)	45
INTEGER COMM, SIZE, IERROR	46
MPI_COMM_RANK(COMM, RANK, IERROR)	47
INTEGER COMM, RANK, IERROR	48


```
1 MPI_COMM_COMPARE(COMM1, COMM2, RESULT, IERROR)
2     INTEGER COMM1, COMM2, RESULT, IERROR
3
4 MPI_COMM_DUP(COMM, NEWCOMM, IERROR)
5     INTEGER COMM, NEWCOMM, IERROR
6
7 MPI_COMM_CREATE(COMM, GROUP, NEWCOMM, IERROR)
8     INTEGER COMM, GROUP, NEWCOMM, IERROR
9
10 MPI_COMM_SPLIT(COMM, COLOR, KEY, NEWCOMM, IERROR)
11     INTEGER COMM, COLOR, KEY, NEWCOMM, IERROR
12
13 MPI_COMM_FREE(COMM, IERROR)
14     INTEGER COMM, IERROR
15
16 MPI_COMM_TEST_INTER(COMM, FLAG, IERROR)
17     INTEGER COMM, IERROR
18     LOGICAL FLAG
19
20 MPI_COMM_REMOTE_SIZE(COMM, SIZE, IERROR)
21     INTEGER COMM, SIZE, IERROR
22
23 MPI_COMM_REMOTE_GROUP(COMM, GROUP, IERROR)
24     INTEGER COMM, GROUP, IERROR
25
26 MPI_INTERCOMM_CREATE(LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG,
27     NEWINTERCOMM, IERROR)
28     INTEGER LOCAL_COMM, LOCAL_LEADER, PEER_COMM, REMOTE_LEADER, TAG,
29     NEWINTERCOMM, IERROR
30
31 MPI_INTERCOMM_MERGE(INTERCOMM, HIGH, NEWINTRACOMM, IERROR)
32     INTEGER INTERCOMM, NEWINTRACOMM, IERROR
33     LOGICAL HIGH
34
35 MPI_KEYVAL_CREATE(COPY_FN, DELETE_FN, KEYVAL, EXTRA_STATE, IERROR)
36     EXTERNAL COPY_FN, DELETE_FN
37     INTEGER KEYVAL, EXTRA_STATE, IERROR
38
39 MPI_KEYVAL_FREE(KEYVAL, IERROR)
40     INTEGER KEYVAL, IERROR
41
42 MPI_ATTR_PUT(COMM, KEYVAL, ATTRIBUTE_VAL, IERROR)
43     INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR
44
45 MPI_ATTR_GET(COMM, KEYVAL, ATTRIBUTE_VAL, FLAG, IERROR)
46     INTEGER COMM, KEYVAL, ATTRIBUTE_VAL, IERROR
47     LOGICAL FLAG
48
49 MPI_ATTR_DELETE(COMM, KEYVAL, IERROR)
50     INTEGER COMM, KEYVAL, IERROR
```

A.12 Fortran Bindings for Process Topologies

	1
	2
MPI_CART_CREATE(COMM_OLD, NDIMS, DIMS, PERIODS, REORDER, COMM_CART, IERROR)	3
INTEGER COMM_OLD, NDIMS, DIMS(*), COMM_CART, IERROR	4
LOGICAL PERIODS(*), REORDER	5
	6
MPI_DIMS_CREATE(NNODES, NDIMS, DIMS, IERROR)	7
INTEGER NNODES, NDIMS, DIMS(*), IERROR	8
	9
MPI_GRAPH_CREATE(COMM_OLD, NNODES, INDEX, EDGES, REORDER, COMM_GRAPH,	10
IERROR)	11
INTEGER COMM_OLD, NNODES, INDEX(*), EDGES(*), COMM_GRAPH, IERROR	12
LOGICAL REORDER	13
	14
MPI_TOPO_TEST(COMM, STATUS, IERROR)	15
INTEGER COMM, STATUS, IERROR	16
	17
MPI_GRAPHDIMS_GET(COMM, NNODES, NEDGES, IERROR)	18
INTEGER COMM, NNODES, NEDGES, IERROR	19
	20
MPI_GRAPH_GET(COMM, MAXINDEX, MAXEDGES, INDEX, EDGES, IERROR)	21
INTEGER COMM, MAXINDEX, MAXEDGES, INDEX(*), EDGES(*), IERROR	22
	23
MPI_CARTDIM_GET(COMM, NDIMS, IERROR)	24
INTEGER COMM, NDIMS, IERROR	25
	26
MPI_CART_GET(COMM, MAXDIMS, DIMS, PERIODS, COORDS, IERROR)	27
INTEGER COMM, MAXDIMS, DIMS(*), COORDS(*), IERROR	28
LOGICAL PERIODS(*)	29
	30
MPI_CART_RANK(COMM, COORDS, RANK, IERROR)	31
INTEGER COMM, COORDS(*), RANK, IERROR	32
	33
MPI_CART_COORDS(COMM, RANK, MAXDIMS, COORDS, IERROR)	34
INTEGER COMM, RANK, MAXDIMS, COORDS(*), IERROR	35
	36
MPI_GRAPH_NEIGHBORS_COUNT(COMM, RANK, NNEIGHBORS, IERROR)	37
INTEGER COMM, RANK, NNEIGHBORS, IERROR	38
	39
MPI_GRAPH_NEIGHBORS(COMM, RANK, MAXNEIGHBORS, NEIGHBORS, IERROR)	40
INTEGER COMM, RANK, MAXNEIGHBORS, NEIGHBORS(*), IERROR	41
	42
MPI_CART_SHIFT(COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR)	43
INTEGER COMM, DIRECTION, DISP, RANK_SOURCE, RANK_DEST, IERROR	44
	45
MPI_CART_SUB(COMM, REMAIN_DIMS, NEWCOMM, IERROR)	46
INTEGER COMM, NEWCOMM, IERROR	47
LOGICAL REMAIN_DIMS(*)	48
MPI_CART_MAP(COMM, NDIMS, DIMS, PERIODS, NEWRANK, IERROR)	
INTEGER COMM, NDIMS, DIMS(*), NEWRANK, IERROR	
LOGICAL PERIODS(*)	
MPI_GRAPH_MAP(COMM, NNODES, INDEX, EDGES, NEWRANK, IERROR)	
INTEGER COMM, NNODES, INDEX(*), EDGES(*), NEWRANK, IERROR	

1 A.13 Fortran Bindings for Environmental Inquiry

```
2
3 MPI_GET_PROCESSOR_NAME(NAME, RESULTLEN, IERROR)
4     CHARACTER*(*) NAME
5     INTEGER RESULTLEN, IERROR
6
7 MPI_ERRHANDLER_CREATE(FUNCTION, ERRHANDLER, IERROR)
8     EXTERNAL FUNCTION
9     INTEGER ERRHANDLER, IERROR
10
11 MPI_ERRHANDLER_SET(COMM, ERRHANDLER, IERROR)
12     INTEGER COMM, ERRHANDLER, IERROR
13
14 MPI_ERRHANDLER_GET(COMM, ERRHANDLER, IERROR)
15     INTEGER COMM, ERRHANDLER, IERROR
16
17 MPI_ERRHANDLER_FREE(ERRHANDLER, IERROR)
18     INTEGER ERRHANDLER, IERROR
19
20 MPI_ERROR_STRING(ERRORCODE, STRING, RESULTLEN, IERROR)
21     INTEGER ERRORCODE, RESULTLEN, IERROR
22     CHARACTER*(*) STRING
23
24 MPI_ERROR_CLASS(ERRORCODE, ERRORCLASS, IERROR)
25     INTEGER ERRORCODE, ERRORCLASS, IERROR
26
27 DOUBLE PRECISION MPI_WTIME()
28
29 DOUBLE PRECISION MPI_WTICK()
30
31 MPI_INIT(IERROR)
32     INTEGER IERROR
33
34 MPI_FINALIZE(IERROR)
35     INTEGER IERROR
36
37 MPI_INITIALIZED(FLAG, IERROR)
38     LOGICAL FLAG
39     INTEGER IERROR
40
41 MPI_ABORT(COMM, ERRORCODE, IERROR)
42     INTEGER COMM, ERRORCODE, IERROR
43
44 MPI_GET_VERSION(VERSION, SUBVERSION, IERROR)
45     INTEGER VERSION, SUBVERSION, IERROR
46
47
48
```

42 A.14 Fortran Bindings for Profiling

```
43
44 MPI_PCONTROL(LEVEL)
45     INTEGER LEVEL, ...
46
47
48
```

MPI Function Index

MPIABORT, 206
MPIADDRESS, 70
MPIALLGATHER, 111
MPIALLGATHERV, 112
MPIALLREDUCE, 126
MPIALLTOALL, 113
MPIALLTOALLV, 114
MPIATTR_DELETE, 174
MPIATTR_GET, 173
MPIATTR_PUT, 173
MPIBARRIER, 96
MPIBCAST, 96
MPIBSEND, 28
MPIBSEND_INIT, 57
MPIBUFFER_ATTACH, 34
MPIBUFFER_DETACH, 34
MPIOCANCEL, 54
MPIOCART_COORDS, 186
MPIOCART_CREATE, 181
MPIOCART_GET, 185
MPIOCART_MAP, 190
MPIOCART_RANK, 186
MPIOCART_SHIFT, 188
MPIOCART_SUB, 189
MPIOCARTDIM_GET, 185
MPIOCOMM_COMPARE, 146
MPIOCOMM_CREATE, 148
MPIOCOMM_DUP, 147
MPIOCOMM_FREE, 150
MPIOCOMM_GROUP, 141
MPIOCOMM_RANK, 146
MPIOCOMM_REMOTE_GROUP, 160
MPIOCOMM_REMOTE_SIZE, 160
MPIOCOMM_SIZE, 145
MPIOCOMM_SPLIT, 149
MPIOCOMM_TEST_INTER, 159
MPIODIMS_CREATE, 181
MPIERRHANDLER_CREATE, 197
MPIERRHANDLER_FREE, 198
MPIERRHANDLER_GET, 198
MPIERRHANDLER_SET, 198
MPIERROR_CLASS, 200
MPIERROR_STRING, 199
MPIFINALIZE, 202
MPIGATHER, 97
MPIGATHERV, 99
MPIGET_COUNT, 22
MPIGET_ELEMENTS, 76
MPIGET_PROCESSOR_NAME, 195
MPIGET_VERSION, 193
MPIGRAPH_CREATE, 182
MPIGRAPH_GET, 185
MPIGRAPH_MAP, 191
MPIGRAPH_NEIGHBORS, 187
MPIGRAPH_NEIGHBORS_COUNT, 187
MPIGRAPHDIMS_GET, 184
MPIGROUP_COMPARE, 140
MPIGROUP_DIFFERENCE, 142
MPIGROUP_EXCL, 143
MPIGROUP_FREE, 144
MPIGROUP_INCL, 142
MPIGROUP_INTERSECTION, 141
MPIGROUP_RANGE_EXCL, 144
MPIGROUP_RANGE_INCL, 143
MPIGROUP_RANK, 139
MPIGROUP_SIZE, 139
MPIGROUP_TRANSLATE_RANKS, 140
MPIGROUP_UNION, 141
MPIIBSEND, 39
MPIINIT, 202
MPIINITIALIZED, 205
MPIINTERCOMM_CREATE, 161
MPIINTERCOMM_MERGE, 162
MPIIPROBE, 52
MPIIRECV, 40
MPIIRSEND, 40
MPIISEND, 38
MPIISSEND, 39
MPIKEYVAL_CREATE, 171
MPIKEYVAL_FREE, 172

- 1 MPIOP_CREATE, [122](#)
 - 2 MPIOP_FREE, [124](#)
 - 3 MPI_PACK, [87](#)
 - 4 MPI_PACK_SIZE, [89](#)
 - 5 MPI_PCONTROL, [208](#)
 - 6 MPI_PROBE, [53](#)
 - 7 MPI_RECV, [20](#)
 - 8 MPI_RECV_INIT, [58](#)
 - 9 MPI_REDUCE, [115](#)
 - 10 MPI_REDUCE_SCATTER, [127](#)
 - 11 MPI_REQUEST_FREE, [43](#)
 - 12 MPI_RSEND, [29](#)
 - 13 MPI_RSEND_INIT, [58](#)
 - 14 MPI_SCAN, [128](#)
 - 15 MPI_SCATTER, [106](#)
 - 16 MPI_SCATTERV, [107](#)
 - 17 MPI_SEND, [17](#)
 - 18 MPI_SEND_INIT, [56](#)
 - 19 MPI_SENDRECV, [60](#)
 - 20 MPI_SENDRECV_REPLACE, [61](#)
 - 21 MPI_SSEND, [29](#)
 - 22 MPI_SSEND_INIT, [57](#)
 - 23 MPI_START, [59](#)
 - 24 MPI_STARTALL, [59](#)
 - 25 MPI_TEST, [42](#)
 - 26 MPI_TEST_CANCELLED, [55](#)
 - 27 MPI_TESTALL, [48](#)
 - 28 MPI_TESTANY, [46](#)
 - 29 MPI_TESTSOME, [50](#)
 - 30 MPI_TOPO_TEST, [184](#)
 - 31 MPI_TYPE_COMMIT, [73](#)
 - 32 MPI_TYPE_CONTIGUOUS, [64](#)
 - 33 MPI_TYPE_EXTENT, [71](#)
 - 34 MPI_TYPE_FREE, [74](#)
 - 35 MPI_TYPE_HINDEXED, [68](#)
 - 36 MPI_TYPE_HVECTOR, [66](#)
 - 37 MPI_TYPE_INDEXED, [67](#)
 - 38 MPI_TYPE_LB, [73](#)
 - 39 MPI_TYPE_SIZE, [71](#)
 - 40 MPI_TYPE_STRUCT, [69](#)
 - 41 MPI_TYPE_UB, [73](#)
 - 42 MPI_TYPE_VECTOR, [65](#)
 - 43 MPI_UNPACK, [87](#)
 - 44 MPI_WAIT, [41](#)
 - 45 MPI_WAITALL, [47](#)
 - 46 MPI_WAITANY, [46](#)
 - 47 MPI_WAITSOME, [49](#)
 - 48 MPI_WTICK, [201](#)
- MPI_WTIME, [201](#)