IBM Parallel Environment for AIX: MPI Programming and Subroutine Reference
Version 2, Release 3

Document Number GC23-3894-02

Note

Before using this information and the product is supports, be sure to read the general information under “Notices” on page v

| **Third Edition (August 1997)**

| This edition applies to the IBM Parallel Environment for AIX, Version 2, Release 3, Modification 0, program number 5765-543, and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

IBM welcomes your comments. A form for your comments appears at the back of this publication. If the form has been removed, address your comments to:

IBM Corporation, Department 55JA, Mail Station P384
522 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX: (United States and Canada): 914+432-9405
FAX: (Other Countries)
Your International Access Code +1+914+432-9405

IBMLink (United States customers only): KGNVMC(MHVRCFS)
IBM Mail Exchange: USIB6TC9 at IBMMAIL
Internet: mhvrcfs@vnet.ibm.com

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:

- Title and order number of this book
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Copyright International Business Machines Corporation 1995, 1997. All rights reserved. Note to U.S. Government Users - Documentation related to restricted rights - Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corporation.

Permission to copy without fee all or part of *MPI: A Message Passing Interface Standard, Version 1.1*, Message Passing Interface Forum is granted, provided the University of Tennessee copyright notice and the title of the document appear, and notice is given that copying is by permission of the University of Tennessee. ©1993, 1995 University of Tennessee, Knoxville, Tennessee.

Contents

Notices v
Trademarks v
About This Book vi
Who Should Use This Book vi
How This Book Is Organized vi
Overview of Contents vi
Typographic Conventions vii
How FORTRAN and C Are Documented vii
Related Publications vii
IBM Parallel Environment for AIX Publications vii
Related IBM Publications viii
Related Non-IBM Publications viii
National Language Support viii
Accessing Online Information ix
Obtaining Documentation ix

Chapter 1. Subroutines Listed Alphabetically 1

Appendix A. MPI Subroutine Bindings Quick Reference 253

Appendix B. Profiling Message Passing 269

Appendix C. MPI Size Limits 275

Appendix D. Reduction Operations 277

Appendix E. Parallel Utility Functions 281

Appendix F. Tracing Routines 311

Appendix G. Programming Considerations for User Applications in POE 329

Appendix H. Using Signals and the IBM PE Programs 345

Appendix I. Predefined Datatypes 349

Index 353

© Copyright IBM Corp. 1997

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
USA

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
522 South Road
Poughkeepsie, NY 12601-5400
USA
Attention: Information Request

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

AIX
IBM
RISC System/6000
Scalable POWERparallel Systems

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PostScript is a trademark of Adobe Systems, Incorporated.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

© Copyright IBM Corp. 1997

Other company, product, and service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others.

About This Book

This book lists the subroutines a programmer can use when writing parallel applications along with the associated parameters, and syntax. The IBM Message Passing Interface implementation intends to comply with the requirements of the Message Passing Interface Forum, *MPI: A Message-Passing Interface Standard, Version 1.1*, University of Tennessee, Knoxville, Tennessee, June 6, 1995. In addition, this book provides brief introductory information regarding parallel programming tasks.

Who Should Use This Book

This book is intended for experienced programmers who want to write parallel applications using either the C or FORTRAN programming languages. Readers of this book should know C and FORTRAN and should be familiar with AIX and UNIX commands, file formats, and special files. They should also be familiar with the Message-Passing Interface(MPI) concepts. In addition, readers should be familiar with distributed-memory machines.

How This Book Is Organized

Overview of Contents

This book is divided into the following sections:

- Chapter 1, “Subroutines Listed Alphabetically” on page 1 lists the subroutines alphabetically along with their descriptions, type, syntax and so on. MPI_SAMPLE is included which is not an MPI function but a brief description of how each routine is structured.
- Appendix A, “MPI Subroutine Bindings Quick Reference” on page 253 briefly lists the subroutines and their arguments. Use it as a quick reference. For detailed information on the subroutines refer to Chapter 1, “Subroutines Listed Alphabetically” on page 1.
- Appendix B, “Profiling Message Passing” on page 269 gives information about the name-shifted interface for Message Passing Interface(MPI).
- Appendix C, “MPI Size Limits” on page 275 gives information about the MPI size limits.
- Appendix D, “Reduction Operations” on page 277 gives additional information about reduction functions.
- Appendix E, “Parallel Utility Functions” on page 281 contains the syntax man pages of the user-callable functions that take advantage of the Parallel Operating Environment(POE).
- Appendix F, “Tracing Routines” on page 311 contains the syntax man pages for modifying trace generation for the visualization tool.
- Appendix G, “Programming Considerations for User Applications in POE” on page 329 contains various information for user applications written to run under

the IBM Parallel Environment for AIX. This includes specific considerations for FORTRAN, threaded, and signal-handling library applications.

- Appendix H, "Using Signals and the IBM PE Programs" on page 345 contains information for understanding how the IBM Parallel Environment for AIX(PE) calls use timer signals to manage message traffic. Sample programs are included. This section applies to the signal-handling version of the Message Passing library.
- Appendix I, "Predefined Datatypes" on page 349 contains a list of the various MPI predefined datatypes that you can use with the signal-handling library.

Typographic Conventions

This book uses the following typographic conventions:

Type Style	Used For
Bold	**Bold** words or characters represent system elements that you must use literally, such as command names, program names, file names, flag names, and path names.
	Bold words also indicate the first use of a term included in the glossary.
Italic	*Italic* words or characters represent variable values that you must supply.
	Italics are also used for book titles and for general emphasis in text.
`Constant width`	Examples and information that the system displays appear in `constant width` typeface.

How FORTRAN and C Are Documented

FORTRAN is not case sensitive while the reverse is true for C. This means that unless you use the XLF complier option **-qmixed**, case does not matter in FORTRAN subroutine names. However, to ensure MPI standard compliant code, it is suggested that all FORTRAN subroutine names use upper case. The C subroutines must be entered exactly as specified.

For the purpose of distinguishing between the FORTRAN and C syntax in this document, FORTRAN subroutines are documented in all upper case and are referred to as FORTRAN throughout the book. C is documented in mixed case.

For both C and FORTRAN, the Message-Passing Interface(MPI) uses the same spelling for function names. The only distinction is in the capitalization. For the purpose of clarity, when referring to a function without specifying FORTRAN or C version, the function is in all upper case.

Related Publications

IBM Parallel Environment for AIX Publications

- *IBM Parallel Environment for AIX: General Information*, GC23-3906
- *IBM Parallel Environment for AIX: Hitchhiker's Guide*, GC23-3895
- *IBM Parallel Environment for AIX: Installation*, GC28-1981
- *IBM Parallel Environment for AIX: Messages*, GC28-1982
- *IBM Parallel Environment for AIX: Operation and Use, Volume 1*, SC28-1979

- *IBM Parallel Environment for AIX: Operation and Use, Volume 2*, SC28-1980
- *IBM Parallel Environment for AIX: Programming and Subroutine Reference*, GC23-3894
- *IBM Parallel Environment for AIX: MPL Programming and Subroutine Reference*, GC23-3893
- *IBM Parallel Environment for AIX: Licensed Program Specification*, GC23-3896

As an alternative to ordering the individual books, you can use SBOF-8588 to order the entire IBM Parallel Environment for AIX library.

Related IBM Publications

- *IBM AIX Version 4.1 Technical References*, SBOF-1852
- *IBM XL FORTRAN Compiler Version 3.2 for AIX Language Reference*, SC09-1611

Related Non-IBM Publications

- Snir, M., Otto, Steve W., Huss-Lederman, Steven, Walker, David W., Dongarra, Jack, *MPI: The Complete Reference*, The MIT Press, 1995, ISBN 0-262-69184-1.
- Gropp, W., Lusk, E., Skejellum, A., *Using MPI*, The MIT Press, 1994.

  As an alternative, you can use SR28-5757-00 to order this book through your IBM representative or IBM branch office serving your locality.
- Message Passing Interface Forum, *MPI: A Message-Passing Interface Standard, Version 1.1*, University of Tennessee, Knoxville, Tennessee, June 6, 1995.

  Permission to copy without fee all or part of *MPI: A Message-Passing Interface Standard, Version 1.1*, Message Passing Interface Forum is granted, provided the University of Tennessee copyright notice and the title of the document appear, and notice is given that copying is by permission of the University of Tennessee. ©1993, 1995 University of Tennessee, Knoxville, Tennessee.

National Language Support

For National Language Support (NLS), all PE components and tools display messages located in externalized message catalogs. English versions of the message catalogs are shipped with the PE program product, but your site may be using its own translated message catalogs. The AIX environment variable **NLSPATH** is used by the various PE components to find the appropriate message catalog. **NLSPATH** specifies a list of directories to search for message catalogs. The directories are searched, in the order listed, to locate the message catalog. In resolving the path to the message catalog, **NLSPATH** is affected by the values of the environment variables **LC_MESSAGES** and **LANG**. If you get an error saying that a message catalog is not found, and want the default message catalog:

ENTER **export NLSPATH=/usr/lib/nls/msg/%L/%N**

export LANG=C

The PE message catalogs are in English, and are located in the following directories:

/usr/lib/nls/msg/C
/usr/lib/nls/msg/En_US
/usr/lib/nls/msg/en_US

If your site is using its own translations of the message catalogs, consult your system administrator for the appropriate value of **NLSPATH** or **LANG**. For additional information on NLS and message catalogs, see *IBM Parallel Environment for AIX: Messages* and *IBM AIX Version 4 General Programming Concepts: Writing and Debugging Programs*.

Accessing Online Information

The PE program product provides manual pages as well as online versions of the publications that can be viewed using an HTML browser (such as Netscape, WebExplorer, or Mosaic).

To access a list of manual pages for PE:

ENTER man pe

Note: In order to use the PE MPI subroutine manual pages or access the PE online (HTML) publications, the **ppe.pedocs** fileset must first be installed. To view the PE online publications, you need the following:

- Access to a common HTML document browser (such as Netscape, WebExplorer, or Mosaic)
- The location of the HTML index file, provided with the **ppe.pedocs** fileset. Contact your system administrator or installer for this location.

Obtaining Documentation

This book is available in PostScript format. You can obtain the PostScript files for the product documentation by:

- Downloading them from the **lscftp.pok.ibm.com** anonymous FTP server. The files for *PE_prod_name* can be found in the directory called /**pub**/**pps**/**PE**.
- Printing the PostScript documentation files that are shipped on the product media.

You can also view or download this book from the IBM RISC System/6000 web site, whose current URL is listed in the PE filesets' README files.

Chapter 1. Subroutines Listed Alphabetically

The following lists the subroutines alphabetically. Refer to the appropriate section for information related to subroutine purpose, syntax, and so on.

Subroutine C/FORTRAN	Pg	Type	Description
MPE_Iallgather MPE_IALLGATHER	12	Nonblocking Collective Communication	Nonblocking allgather operation.
MPE_Iallgatherv MPE_IALLGATHERV	14	Nonblocking Collective Communication	Nonblocking allgatherv operation.
MPE_Iallreduce MPE_IALLREDUCE	16	Nonblocking Collective Communication	Nonblocking allreduce operation.
MPE_Ialltoall MPE_IALLTOALL	18	Nonblocking Collective Communication	Nonblocking alltoall operation.
MPE_Ialltoallv MPE_IALLTOALLV	20	Nonblocking Collective Communication	Nonblocking alltoallv operation.
MPE_Ibarrier MPE_IBARRIER	22	Nonblocking Collective Communication	Nonblocking barrier operation.
MPE_Ibcast MPE_IBCAST	24	Nonblocking Collective Communication	Nonblocking broadcast operation.
MPE_Igather MPE_IGATHER	26	Nonblocking Collective Communication	Nonblocking gather operation.
MPE_Igatherv MPE_IGATHERV	28	Nonblocking Collective Communication	Nonblocking gatherv operation.
MPE_Ireduce_scatter MPE_IREDUCE_SCATTER	30	Nonblocking Collective Communication	Nonblocking reduce_scatter operation.
MPE_Ireduce MPE_IREDUCE	32	Nonblocking Collective Communication	Nonblocking reduce operation.
MPE_Iscan MPE_ISCAN	34	Nonblocking Collective Communication	Nonblocking scan operation.
MPE_Iscatter MPE_ISCATTER	36	Nonblocking Collective Communication	Nonblocking scatter operation.
MPE_Iscatterv MPE_ISCATTERV	38	Nonblocking Collective Communication	Nonblocking scatterv operation.

Subroutine C/FORTRAN	Page	Type	Description
MPI_Abort MPI_ABORT	40	Environmental Management	Forces all tasks of an MPI job to terminate.
MPI_Address MPI_ADDRESS	41	Derived Datatype	Returns address of a location in memory.

© Copyright IBM Corp. 1997

Subroutine C/FORTRAN	Page	Type	Description
MPI_Allgather MPI_ALLGATHER	42	Collective Communication	Collects messages from each process and distributes the resulting message to each.
MPI_Allgatherv MPI_ALLGATHERV	44	Collective Communication	Collects messages from each process and distributes the resulting message to all processes. Messages can have variable sizes and displacements.
MPI_Allreduce MPI_ALLREDUCE	46	Collective Communication	Applies a reduction operation to the vector **sendbuf** over the set of processes specified by **comm** and places the result in **recvbuf** on all of the processes.
MPI_Alltoall MPI_ALLTOALL	48	Collective Communication	Sends a distinct message from each process to every process.
MPI_Alltoallv MPI_ALLTOALLV	50	Collective Communication	Sends a distinct message from each process to every process. Messages can have different sizes and displacements.
MPI_Attr_delete MPI_ATTR_DELETE	52	Communicator	Removes an attribute value from a communicator.
MPI_Attr_get MPI_ATTR_GET	53	Communicator	Retrieves an attribute value from a communicator.
MPI_Attr_put MPI_ATTR_PUT	55	Communicator	Associates an attribute value with a communicator.
MPI_Barrier MPI_BARRIER	57	Collective Communication	Blocks each process until all processes have called it.
MPI_Bcast MPI_BCAST	58	Collective Communication	Broadcasts a message from **root** to all processes in the group.
MPI_Bsend MPI_BSEND	60	Point-to-Point	Blocking buffered mode send.
MPI_Bsend_init MPI_BSEND_INIT	62	Point-to-Point	Creates a persistent buffered mode send request.
MPI_Buffer_attach MPI_BUFFER_ATTACH	64	Point-to-Point	Provides MPI with a message buffer for sending.
MPI_Buffer_detach MPI_BUFFER_DETACH	65	Point-to-Point	Detaches the current buffer.
MPI_Cancel MPI_CANCEL	67	Point-to-Point	Marks a nonblocking point-to-point operation for cancellation.
MPI_Cart_coords MPI_CART_COORDS	69	Topology	Translates process rank in a communicator into cartesian process coordinates.

Subroutine C/FORTRAN	Page	Type	Description
MPI_Cart_create MPI_CART_CREATE	71	Topology	Creates a communicator containing topology information.
MPI_Cart_get MPI_CART_GET	73	Topology	Retrieves cartesian topology information from a communicator.
MPI_Cart_map MPI_CART_MAP	75	Topology	Computes placement of processes on the physical machine.
MPI_Cart_rank MPI_CART_RANK	77	Topology	Translates process coordinates into a process rank.
MPI_Cart_shift MPI_CART_SHIFT	79	Topology	Returns shifted source and destination ranks for a process.
MPI_Cart_sub MPI_CART_SUB	81	Topology	Partitions a cartesian communicator into lower-dimensional subgroups.
MPI_Cartdim_get MPI_CARTDIM_GET	83	Topology	Retrieves the number of cartesian dimensions from a communicator.
MPI_Comm_compare MPI_COMM_COMPARE	84	Communicator	Compares the groups and contexts of two communicators.
MPI_Comm_create MPI_COMM_CREATE	86	Communicator	Creates a new intracommunicator with a given group.
MPI_Comm_dup MPI_COMM_DUP	88	Communicator	Creates a new communicator that is a duplicate of an existing communicator.
MPI_Comm_free MPI_COMM_FREE	90	Communicator	Marks a communicator for deallocation.
MPI_Comm_group MPI_COMM_GROUP	91	Process Group	Returns the group handle associated with a communicator.
MPI_Comm_rank MPI_COMM_RANK	92	Communicator	Returns the rank of the local process in the group associated with a communicator.
MPI_Comm_remote_group MPI_COMM_REMOTE_GROUP	94	Communicator	Returns the handle of the remote group of an intercommunicator.
MPI_Comm_remote_size MPI_COMM_REMOTE_SIZE	95	Communicator	Returns the size of the remote group of an intercommunicator.
MPI_Comm_size MPI_COMM_SIZE	96	Communicator	Returns the size of the group associated with a communicator.
MPI_Comm_split MPI_COMM_SPLIT	98	Communicator	Splits a communicator into multiple communicators based on **color** and **key**.
MPI_Comm_test_inter MPI_COMM_TEST_INTER	100	Communicator	Returns the type of a communicator (intra or inter).
MPI_Dims_create MPI_DIMS_CREATE	101	Topology	Defines a cartesian grid to balance processes.

Subroutine C/FORTRAN	Page	Type	Description
MPI_Errhandler_create MPI_ERRHANDLER_CREATE	103	Environmental Management	Registers a user defined error handler.
MPI_Errhandler_free MPI_ERRHANDLER_FREE	104	Environmental Management	Marks an error handler for deallocation.
MPI_Errhandler_get MPI_ERRHANDLER_GET	105	Environmental Management	Gets an error handler associated with a communicator.
MPI_Errhandler_set MPI_ERRHANDLER_SET	106	Environmental Management	Associates a new error handler with a communicator.
MPI_Error_class MPI_ERROR_CLASS	108	Environmental Management	Returns the error class for the corresponding error code.
MPI_Error_string MPI_ERROR_STRING	110	Environmental Management	Returns the error string for a given error code.
MPI_Finalize MPI_FINALIZE	111	Environmental Management	Terminates all MPI processing.
MPI_Gather MPI_GATHER	113	Collective Communication	Collects individual messages from each process in a group at the **root** process.
MPI_Gatherv MPI_GATHERV	115	Collective Communication	Collects individual messages from each process in **comm** at the **root** process. Messages can have different sizes and displacements.
MPI_Get_count MPI_GET_COUNT	117	Point-to-Point	Returns the number of elements in a message.
MPI_Get_elements MPI_GET_ELEMENTS	119	Derived Datatype	Returns the number of basic elements in a message.
MPI_Get_processor_name MPI_GET_PROCESSOR_NAME	121	Environmental Management	Returns the name of the local processor.
MPI_Get_version MPI_GET_VERSION	122	Environmental Management	Returns the version of MPI standard supported.
MPI_Graph_create MPI_GRAPH_CREATE	123	Topology	Creates a new communicator containing graph topology information.
MPI_Graph_get MPI_GRAPH_GET	125	Topology	Retrieves graph topology information from a communicator.
MPI_Graph_map MPI_GRAPH_MAP	75	Topology	Computes placement of processes on the physical machine.
MPI_Graph_neighbors MPI_GRAPH_NEIGHBORS	129	Topology	Returns the neighbors of the given process.
MPI_Graph_neighbors_count MPI_GRAPH_NEIGHBORS_COUNT	130	Topology	Returns the number of neighbors of the given process.
MPI_Graphdims_get MPI_GRAPHDIMS_GET	131	Topology	Retrieves graph topology information from a communicator.

Subroutine C/FORTRAN	Page	Type	Description
MPI_Group_compare MPI_GROUP_COMPARE	132	Process Group	Compares the contents of two process groups.
MPI_Group_difference MPI_GROUP_DIFFERENCE	133	Process Group	Creates a new group that is the difference of two existing groups.
MPI_Group_excl MPI_GROUP_EXCL	134	Process Group	Removes selected processes from an existing group to create a new group.
MPI_Group_free MPI_GROUP_FREE	136	Process Group	Marks a group for deallocation.
MPI_Group_incl MPI_GROUP_INCL	137	Process Group	Creates a new group consisting of selected processes from an existing group.
MPI_Group_intersection MPI_GROUP_INTERSECTION	139	Process Group	Creates a new group that is the intersection of two existing groups.
MPI_Group_range_excl MPI_GROUP_RANGE_EXCL	140	Process Group	Creates a new group by excluding selected processes of an existing group.
MPI_Group_range_incl MPI_GROUP_RANGE_INCL	142	Process Group	Creates a new group consisting of selected ranges of processes from an existing group.
MPI_Group_rank MPI_GROUP_RANK	144	Process Group	Returns the rank of the local process with respect to group.
MPI_Group_size MPI_GROUP_SIZE	145	Process Group	Returns the number of processes in a group.
MPI_Group_translate_ranks MPI_GROUP_TRANSLATE_RANKS	146	Process Group	Converts process ranks of one group into ranks of another group.
MPI_Group_union MPI_GROUP_UNION	148	Process Group	Creates a new group that is the union of two existing groups.
MPI_Ibsend MPI_IBSEND	149	Point-to-Point	Nonblocking buffered send.
MPI_Init MPI_INIT	151	Environmental Management	Initializes MPI.
MPI_Initialized MPI_INITIALIZED	153	Environmental Management	Determines if MPI is initialized.
MPI_Intercomm_create MPI_INTERCOM_CREATE	154	Communicator	Returns the handle of the remote group of an intercommunicator.
MPI_Intercomm_merge MPI_INTERCOMM_MERGE	156	Communicator	Creates an intracommunicator by merging the local and the remote groups of an intercommunicator.
MPI_Iprobe MPI_IPROBE	157	Point-to-Point	Checks if a message matching **source**, **tag**, and **comm** has arrived.

Subroutine C/FORTRAN	Page	Type	Description
MPI_Irecv MPI_IRECV	159	Point-to-Point	Nonblocking receive.
MPI_Irsend MPI_IRSEND	161	Point-to-Point	Nonblocking ready send.
MPI_Isend MPI_ISEND	163	Point-to-Point	Nonblocking standard mode send.
MPI_Issend MPI_ISSEND	165	Point-to-Point	Nonblocking synchronous mode send.
MPI_Keyval_create MPI_KEYVAL_CREATE	167	Communicator	Generates a new attribute key.
MPI_Keyval_free MPI_KEYVAL_FREE	169	Communicator	Marks an attribute key for deallocation.
MPI_Op_create MPI_OP_CREATE	170	Collective Communication	Binds a user defined reduction operation to an **op** handle.
MPI_Op_free MPI_OP_FREE	172	Collective Communication	Marks a user defined reduction operation for deallocation.
MPI_Pack MPI_PACK	173	Derived Datatype	Packs the message in the specified send buffer into the specified buffer space.
MPI_Pack_size MPI_PACK_SIZE	175	Dervived Datatype	Returns the number of bytes required to hold the data.
MPI_Pcontrol MPI_PCONTROL	176	Environmental Management	Provides profile control.
MPI_Probe MPI_PROBE	177	Point-to-Point	Waits until a message matching **source**, **tag**, and **comm** arrives.
MPI_Recv MPI_RECV	179	Point-to-Point	Blocking receive
MPI_Recv_init MPI_RECV_INIT	181	Point-to-Point	Creates a persistent receive request.
MPI_Reduce MPI_REDUCE	183	Collective Communication	Reduces processes specified and places the result in **recvbuf** on **root**.
MPI_Reduce_scatter MPI_REDUCE_SCATTER	185	Collective Communication	Applies a reduction operation to the vector **sendbuf** over the set of processes specified by **comm** and scatters the result according to the values in **recvcounts**.
MPI_Request_free MPI_REQUEST_FREE	187	Point-to-Point	Marks a request for deallocation.
MPI_Rsend MPI_RSEND	188	Point-to-Point	Blocking ready mode send.
MPI_Rsend_init MPI_RSEND_INIT	190	Point-to-Point	Creates a persistent ready mode send request.

Subroutine C/FORTRAN	Page	Type	Description
MPI_Sample MPI_SAMPLE	10	Sample	This is not an MPI function but a brief description of how each routine is structured.
MPI_Scan MPI_SCAN	192	Collective Communication	Performs a parallel prefix reduction on data distributed across a group.
MPI_Scatter MPI_SCATTER	194	Collective Communication	Distributes individual messages from **root** to each process in **comm**.
MPI_Scatterv MPI_SCATTERV	196	Collective Communication	Distributes individual messages from **root** to each process in **comm**. Messages can have different sizes and displacements.
MPI_Send MPI_SEND	198	Point-to-Point	Blocking standard mode send.
MPI_Send_init MPI_SEND_INIT	200	Point-to-Point	Creates a persistent standard mode send request.
MPI_Sendrecv MPI_SENDRECV	202	Point-to-Point	A blocking send and receive operation.
MPI_Sendrecv_replace MPI_SENDRECV_REPLACE	204	Point-to-Point	Blocking send and receive operation using a common buffer.
MPI_Ssend MPI_SSEND	206	Point-to-Point	Blocking synchronous mode send.
MPI_Ssend_init MPI_SSEND_INIT	208	Point-to-Point	Creates a persistent synchronous mode send request.
MPI_Start MPI_START	210	Point-to-Point	Activates a persistent request operation.
MPI_Startall MPI_STARTALL	211	Point-to-Point	Activates a collection of persistent request operations.
MPI_Test MPI_TEST	213	Point-to-Point	Checks to see if a nonblocking send or receive operation has completed.
MPI_Test_cancelled MPI_TEST_CANCELLED	215	Point-to-Point	Tests whether a nonblocking operation was cancelled.
MPI_Testall MPI_TESTALL	216	Point-to-Point	Tests a collection of nonblocking operations for completion.
MPI_Testany MPI_TESTANY	218	Point-to-Point	Tests for the completion of any specified request.
MPI_Testsome MPI_TESTSOME	220	Point-to-Point	Tests a collection of nonblocking operations for completion.
MPI_Topo_test MPI_TOPO_TEST	222	Topology	Returns the type of virtual topology associated with a communicator.
MPI_Type_commit MPI_TYPE_COMMIT	223	Derived Datatype	Makes a datatype ready for use in communications.

Subroutine C/FORTRAN	Page	Type	Description
MPI_Type_contiguous MPI_TYPE_CONTIGUOUS	224	Derived Datatype	Returns a new datatype that represents the concatenation of **count** instances of **oldtype**.
MPI_Type_extent MPI_TYPE_EXTENT	225	Derived Datatype	Returns the extent of any defined datatype.
MPI_Type_free MPI_TYPE_FREE	226	Derived Datatype	Marks the datatype for deallocation.
MPI_Type_hindexed MPI_TYPE_HINDEXED	227	Derived Datatype	Returns a new datatype that represents **count** distinct blocks with offsets expressed in bytes.
MPI_Type_hvector MPI_TYPE_HVECTOR	229	Derived Datatype	Returns a new datatype of **count** blocks with **stride** expressed in bytes.
MPI_Type_indexed MPI_TYPE_INDEXED	231	Derived Datatype	Returns a new datatype that represents **count** blocks with stride in terms of defining type.
MPI_Type_lb MPI_TYPE_LB	233	Derived Datatype	Returns the lower bound of a datatype.
MPI_Type_size MPI_TYPE_SIZE	234	Derived Datatype	Returns the number of bytes represented by any defined datatype.
MPI_Type_struct MPI_TYPE_STRUCT	235	Derived Datatype	Returns a new datatype that represents **count** blocks each with a distinct format and offset.
MPI_Type_ub MPI_TYPE_UB	237	Derived Datatype	Returns the upper bound of a datatype.
MPI_Type_vector MPI_TYPE_VECTOR	238	Derived Datatype	Returns a new datatype that represents equally spaced blocks of replicated data.
MPI_Unpack MPI_UNPACK	240	Derived Datatype	Unpacks the message into the specified receive buffer from the specified packed buffer.
MPI_Wait MPI_WAIT	242	Point-to-Point	Waits for a nonblocking operation to complete.
MPI_Waitall MPI_WAITALL	244	Point-to-Point	Waits for a collection of nonblocking operations to complete.
MPI_Waitany MPI_WAITANY	246	Point-to-Point	Waits for any specified request to complete.
MPI_Waitsome MPI_WAITSOME	248	Point-to-Point	Waits for at least one of a list of nonblocking operations to complete.
MPI_Wtick MPI_WTICK	250	Environmental Management	Returns the resolution of MPI_Wtime in seconds.
MPI_Wtime MPI_WTIME	251	Environmental Management	Returns the current value of **time** as a floating point value.

The following pages list the subroutines available for parallel programming. For each subroutine, a purpose, C and FORTRAN synopsis, description, notes, and error conditions are provided. Review the following sample subroutine before proceeding to better understand how the subroutines are structured.

NAME

MPI_Sample, MPI_SAMPLE-The text presented in this section briefly describes the use of the subroutine.

C SYNOPSIS

- Header file *mpi.h* supplies ANSI-C prototypes for every function described in the message passing subroutine section of this manual.

```
#include <mpi.h>
int MPI_Sample (parameter(s));
```

In the C prototype:

- A declaration of void * indicates that a pointer to any data type is allowable.

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SAMPLE (parameter(s))
```

In the FORTRAN routines, formal parameters are described using a subroutine prototype format, even though FORTRAN does not support prototyping. The term 'choice' is used to indicate that any FORTRAN datatype is valid.

PARAMETERS

Variable or parameter definitions appear below:

parameter1 parameter description (type)

...

parameter4 parameter description (type)

Parameter types:

IN - call uses but does not update an argument

OUT - call returns information via an argument but does not use its input value

INOUT - call uses and updates an argument

DESCRIPTION

This section contains a more detailed description of the subroutine or function.

NOTES

If applicable, this section contains notes about the IBM MPI implementation and its relationship to the requirements of the MPI Standard. The IBM implementation intends to comply fully with the requirements of the MPI Standard. There are issues, however, which the Standard leaves open to the implementation's choice.

ERRORS

Error conditions appear here.

In almost every routine, the C version is invoked as a function returning integer. In the FORTRAN version, the routine is CALLed as a subroutine; that is, it has no return value. The FORTRAN version includes a return code parameter **IERROR** as the last parameter.

RELATED INFORMATION

This section contains a list of related functions or routines in this book.

For both C and FORTRAN, the Message-Passing Interface(MPI) uses the same spelling for function names. The only distinction is the capitalization. For the purpose of clarity, when referring to a function without specifying FORTRAN or C version, all upper case is used.

FORTRAN refers to FORTRAN 77(F77) bindings which are officially supported for MPI. However, F77 bindings for MPI can be used by Fortran 90. Fortran 90 and High Performance Fortran(HPF) offer array section and assumed shape arrays as parameters on calls. These are not safe with MPI.

NAME

MPE_Iallgather, MPE_IALLGATHER-Nonblocking allgather.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Iallgather(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,MPI_Comm comm,
    MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IALLGATHER(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,
    INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements in the send buffer (integer) (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcount	is the number of elements received from any process (integer) (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_ALLGATHER. It performs the same function as MPI_ALLGATHER except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of your applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective communication routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Unequal message length

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent message length

RELATED INFORMATION

MPI_ALLGATHER

NAME

MPE_Iallgatherv, MPE_IALLGATHERV-Nonblocking allgatherv.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Iallgatherv(void* sendbuf,int sendcount,
    MPI_Datatype sendtype,void* recvbuf,int recvcounts,
    int *displs,MPI_Datatype recvtype,
    MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IALLGATHERV(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
     CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),
     INTEGER RECVTYPE,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements in the send buffer (integer) (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcounts	integer array (of length group size) that contains the number of elements received from each process (IN)
displs	integer array (of length group size). Entry **i** specifies the displacement (relative to **recvbuf**) at which to place the incoming data from process **i** (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communictor (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_ALLGATHERV. It performs the same function as MPI_ALLGATHERV except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Unequal message length

MPI not initialized

MPI already finalized

Develop mode error if:

None

RELATED INFORMATION

MPI_ALLGATHERV

NAME

MPE_Iallreduce, MPE_IALLREDUCE-Nonblocking allreduce operation.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Iallreduce(void* sendbuf,void* recvbuf,int count,
       MPI_Datatype datatype,MPI_Op op,MPI_Comm comm,
       MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IALLREDUCE(CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,
    INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER REQUEST,
    INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
recvbuf	is the starting address of the receive buffer (choice) (OUT)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of elements in the send buffer (handle) (IN)
op	is the reduction operation (handle) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_ALLREDUCE. It performs the same function as MPI_ALLREDUCE except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count **count** < 0

Invalid datatype

Type not committed

Invalid op

Invalid communicator

Invalid communicator type must be intracommunicator

Unequal message length

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent op

Inconsistent datatype

Inconsistent message length

RELATED INFORMATION

MPI_ALLREDUCE

NAME

MPE_Ialltoall, MPE_IALLTOALL-Nonblocking alltoall.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Ialltoall(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,MPI_Comm comm,
    MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IALLTOALL(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,
    INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements sent to each process (integer) (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcount	is the number of elements received from any process (integer) (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_ALLTOALL. It performs the same function as MPI_ALLTOALL except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

Nonblocking collective function can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count(s)	**count** < 0
Invalid datatype(s)	
Type not committed	
Invalid communicator	
Invalid communicator type	must be intracommunicator
Unequal message lengths	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent message lengths

RELATED INFORMATION

MPI_ALLTOALL

NAME

MPE_Ialltoallv, MPE_IALLTOALLV-Nonblocking alltoallv

C SYNOPSIS

```
#include <mpi.h>
int MPE_Ialltoallv(void* sendbuf,int *sendcounts,int *sdispls,
    MPI_Datatype sendtype,void* recvbuf,int *recvcounts,int *rdispls,
    MPI_Datatype recvtype,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_ALLTOALLV(CHOICE SENDBUF,INTEGER SENDCOUNTS(*),
    INTEGER SDISPLS(*),INTEGER SENDTYPE,CHOICE RECVBUF,
    INTEGER RECVCOUNTS(*),INTEGER RDISPLS(*),INTEGER RECVTYPE,
    INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcounts	integer array (of length group size) specifying the number of elements to send to each process (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**. (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcounts	integer array (of length group size) specifying the number of elements that can be received from each process (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**. (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_ALLTOALLV. It performs the same function as MPI_ALLTOALLV except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid communicator

Invalid communicator type must be intracommunicator

A send and receive have unequal message lengths

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ALLTOALLV

NAME

MPE_Ibarrier, MPE_IBARRIER-Nonblocking barrier.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Ibarrier(MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IBARRIER(INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

comm is a communicator (handle) (IN)

request is the communication request (handle) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_BARRIER. It returns immediately, without blocking, but will not complete (via MPI_WAIT or MPI_TEST) until all group members have called it.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

A typical use of MPE_IBARRIER is to make a call to it, and then periodically test for completion with MPI_TEST. Completion indicates that all tasks in **comm** have arrived at the barrier. Until then, computation can continue.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_BARRIER

NAME

MPE_Ibcast, MPE_IBCAST-nonblocking broadcast operation

C SYNOPSIS

```
#include <mpi.h>
int MPE_Ibcast(void* buffer,int count,MPI_Datatype datatype,
    int root,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IBCAST(CHOICE BUFFER,INTEGER COUNT,INTEGER DATATYPE,INTEGER ROOT,
    INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buffer	is the starting address of the buffer (choice) (INOUT)
count	is the number of elements in the buffer (integer) (IN)
datatype	is the datatype of the buffer elements (handle) (IN)
root	is the rank of the root task (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_BCAST. It performs the same function as MPI_BCAST except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, "Programming Considerations for User Applications in POE" on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Error Conditions:

Invalid communicator	
Invalid communicator type	must be intracommunicator
Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid root	**root** < 0 or **root** >= groupsize
Unequal message length	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent root	
Inconsistent message length	

RELATED INFORMATION

MPI_BCAST

NAME

MPE_Igather, MPE_IGATHER-Nonblocking gather.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Igather(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,
    MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IGATHER(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,
    INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements in the send buffer (integer) (IN)
sendtype	is datatype of the send buffer elements (integer) (IN)
recvbuf	is the address of the receive buffer (choice, significant only at **root**) (OUT)
recvcount	is the number of elements for any single receive (integer, significant only at **root**) (IN)
recvtype	is the datatype of the receive buffer elements (handle, significant at **root**) (IN)
root	is the rank of the receiving process (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_GATHER. It performs the same function as MPI_GATHER except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid root **root** <0 or **root** >= groupsize

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

Inconsistent message lengths

RELATED INFORMATION

MPI_GATHER

NAME

MPE_Igatherv, MPE_IGATHERV-Nonblocking gatherv.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Igatherv(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcounts,int *displs,MPI_Datatype recvtype,
    int root,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IGATHERV(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),
    INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,
    INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements to be sent (integer) (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice, significant only at **root**) (OUT)
recvcounts	integer array (of length group size) that contains the number of elements received from each process (significant only at **root**) (IN)
displs	integer array (of length group size). Entry **i** specifies the displacement relative to **recvbuf** at which to place the incoming data from process **i** (significant only at **root**) (IN)
recvtype	is the datatype of the receive buffer elements (handle, significant only at **root**) (IN)
root	is the rank of the receiving process (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_GATHERV. It performs the same function as MPI_GATHERV except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s)

Invalid datatype(s)

Type not committed

Invalid root **root** < 0 or **root** >= groupsize

A send and receive have unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

RELATED INFORMATION

MPI_GATHERV

NAME

MPE_Ireduce_scatter, MPE_IREDUCE_SCATTER-Nonblocking reduce_scatter operation.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Ireduce_scatter(void* sendbuf,void* recvbuf,int *recvcounts,
    MPI_Datatype datatype,MPI_Op op,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IREDUCE_SCATTER(CHOICE SENDBUF,CHOICE RECVBUF,
     INTEGER RECVCOUNTS(*),INTEGER DATATYPE,INTEGER OP,
     INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
recvbuf	is the starting address of the receive buffer (choice) (OUT)
recvcounts	integer array specifying the number of elements in result distributed to each process. Must be identical on all calling processes. (IN)
datatype	is the datatype of elements in the input buffer (handle) (IN)
op	is the reduction operation (handle) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_REDUCE_SCATTER. It performs the same function as MPI_REDUCE_SCATTER except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid recvcount(s) **recvcounts(i)** < 0

Invalid datatype

Type not committed

Invalid op

Invalid communicator

Invalid communicator type must be intracommunicator

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent op

Inconsistent datatype

RELATED INFORMATION

MPI_REDUCE_SCATTER

NAME

MPE_Ireduce, MPE_IREDUCE-Nonblocking reduce operation.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Ireduce(void* sendbuf,void* recvbuf,int count,
    MPI_Datatype datatype,MPI_Op op,int root,MPI_Comm comm,
    MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_IREDUCE(CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,
    INTEGER DATATYPE,INTEGER OP,INTEGER ROOT,INTEGER COMM,
    INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the address of the send buffer (choice) (IN)
recvbuf	is the address of the receive buffer (choice, significant only at root) (OUT)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of elements of the send buffer (handle) (IN)
op	is the reduction operation (handle) (IN)
root	is the rank of the root process (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_REDUCE. It performs the same function as MPI_REDUCE except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid op	
Invalid root	**root** < 0 or **root** > = groupsize
Invalid communicator	
Invalid communicator type	must be intracommunicator
Unequal message lengths	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent op

Inconsistent datatype

Inconsistent root

Inconsistent message length

RELATED INFORMATION

MPI_REDUCE

NAME

MPE_Iscan, MPE_ISCAN-Nonblocking scan operation.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Iscan(void* sendbuf,void* recvbuf,int count,
    MPI_Datatype datatype,MPI_Op op,MPI_Comm comm,
    MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_ISCAN(CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,
    INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER REQUEST,
    INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
recvbuf	is the starting address of the receive buffer (choice) (OUT)
count	is the number of elements in **sendbuf** (integer) (IN)
datatype	is the datatype of elements in **sendbuf** (handle) (IN)
op	is the reduction operation (handle) (IN)
comm	is the communicator (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_SCAN. It performs the same function as MPI_SCAN except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count **count** < 0

Invalid datatype

Type not committed

Invalid op

Invalid communicator

Invalid communicator type must be intracommunicator

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent op

Inconsistent datatype

Inconsistent message length

RELATED INFORMATION

MPI_SCAN

NAME

MPE_Iscatter, MPE_ISCATTER-Nonblocking scatter.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Iscatter(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,
    MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_ISCATTER(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,
    INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf is the address of the send buffer (choice, significant only at **root**) (IN)

sendcount is the number of elements to be sent to each process (integer, significant only at **root**) (IN)

sendtype is the datatype of the send buffer elements (handle, significant only at **root**) (IN)

recvbuf is the address of the receive buffer (choice) (OUT)

recvcount is the number of elements in the receive buffer (integer) (IN)

recvtype is the datatype of the receive buffer elements (handle) (IN)

root is the rank of the sending process (integer) (IN)

comm is the communicator (handle) (IN)

request communication request (handle) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_SCATTER. It performs the same function as MPI_SCATTER except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid root **root** < 0 or **root** >= groupsize

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

Inconsistent message length

RELATED INFORMATION

MPI_SCATTER

NAME

MPE_Iscatterv, MPE_ISCATTERV-Nonblocking scatterv.

C SYNOPSIS

```
#include <mpi.h>
int MPE_Iscatterv(void* sendbuf,int *sendcounts,int *displs,
    MPI_Datatype sendtype,void* recvbuf,int recvcount,
    MPI_Datatype recvtype,int root,MPI_Comm comm,MPI_Comm *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPE_ISCATTERV(CHOICE SENDBUF,INTEGER SENDCOUNTS(*),INTEGER DISPLS(*),
    INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,
    INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

sendbuf is the address of the send buffer (choice, significant only at **root**) (IN)

sendcounts integer array (of length group size) that contains the number of elements to send to each process (significant only at **root**) (IN)

displs integer array (of length group size). Entry **i** specifies the displacement relative to **sendbuf** from which to take the outgoing data to process **i** (significant only at **root**) (IN)

sendtype is the datatype of the send buffer elements (handle, significant only at **root**) (IN)

recvbuf is the address of the receive buffer (choice) (OUT)

recvcount is the number of elements in the receive buffer (integer) (IN)

recvtype is the datatype of the receive buffer elements (handle) (IN)

root is the rank of the sending process (integer) (IN)

comm is the communicator (handle) (IN)

request is the communication request (handle) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a nonblocking version of MPI_SCATTERV. It performs the same function as MPI_SCATTERV except that it returns a **request** handle that must be explicitly completed by using one of the MPI wait or test operations.

NOTES

The "MPE" prefix used with this routine indicates that it is an IBM extension to the MPI standard and is not part of the standard itself. MPE routines are provided to enhance the function and the performance of user applications, but applications that use them will not be directly portable to other MPI implementations.

Nonblocking collective communication routines allow for increased efficiency and flexibility in some applications. Because these routines do not synchronize the participating tasks like blocking collective routines generally do, tasks running at different speeds do not waste time waiting for each other.

When it is expected that tasks will be reasonably synchronized, the blocking collective communication routines provided by standard MPI will commonly give better performance then the nonblocking versions.

The nonblocking collective routines can be used in conjunction with the MPI blocking collective routines and can be completed by any of the MPI wait or test functions. Use of MPI_REQUST_FREE and MPI_CANCEL is not supported.

Applications using nonblocking collective calls often provide their best performance when run in interrupt mode.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator are started in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator	
Invalid communicator type	must be intracommunicator
Invalid count(s)	**count** < 0
Invalid datatype(s)	
Type not committed	
Invalid root	**root** < 0 or **root** >= groupsize
Unequal message lengths	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent root

RELATED INFORMATION

MPI_SCATTERV

NAME

MPI_Abort, MPI_ABORT-Forces all tasks of an MPI job to terminate

C SYNOPSIS

```
#include <mpi.h>
int MPI_Abort(MPI_Comm comm,int errorcode);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ABORT(INTEGER COMM,INTEGER ERRORCODE,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator of the tasks to abort. (IN)
errorcode	is the error code returned to the invoking environment. (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine forces an MPI program to terminate all tasks in the job. **comm** currently is not used. All tasks in the job are aborted. The low order 8 bits of **errorcode** are returned as an AIX return code.

NOTES

MPI_ABORT causes *all* tasks to exit immediately.

ERRORS

MPI already finalized

MPI not initialized

NAME

MPI_Address, MPI_ADDRESS-Returns address of a variable in memory.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Address(void* location,MPI_Aint *address);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ADDRESS(CHOICE LOCATION,INTEGER ADDRESS,INTEGER IERROR)
```

PARAMETERS

location	is the location in caller memory (choice) (IN)
address	is the address of location (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the byte address of **location**.

NOTES

On the RS/6000 Scalable POWERparallel Systems(SP), this is equivalent to **address**= **(MPI_Aint) location** in C, but the MPI_ADDRESS routine is portable to machines with less straightforward addressing.

ERRORS

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_INDEXED
MPI_TYPE_HINDEXED
MPI_TYPE_STRUCT

NAME

MPI_Allgather, MPI_ALLGATHER-Gathers individual messages from each process in **comm** and distributes the resulting message to each process.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Allgather(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ALLGATHER(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,
    INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements in the send buffer (integer) (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcount	is the number of elements received from any process (integer) (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_ALLGATHER is similar to MPI_GATHER except that all processes receive the result instead of just the **root**.

The block of data sent from process **j** 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.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator	
Invalid communicator type	must be intracommunicator
Invalid count(s)	**count** < 0
Invalid datatype(s)	
Type not committed	
Unequal message length	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent message length

RELATED INFORMATION

MPE_IALLGATHER
MPI_ALLGATHER
MPI_GATHER

NAME

MPI_Allgatherv, MPI_ALLGATHERV-Collects individual messages from each process in **comm** and distributes the resulting message to all processes. Messages can have different sizes and displacements.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Allgatherv(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int *recvcounts,int *displs,MPI_Datatype recvtype,
    MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ALLGATHERV(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),
    INTEGER RECVTYPE,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements in the send buffer (integer) (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcounts	integer array (of length group size) that contains the number of elements received from each process (IN)
displs	integer array (of length group size). Entry **i** specifies the displacement (relative to **recvbuf**) at which to place the incoming data from process **i** (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communictor (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine collects individual messages from each process in **comm** and distributes the resulting message to all processes. Messages can have different sizes and displacements.

The block of data sent from process **j** is **recvcounts[j]** elements long, and is received by every process and placed in **recvbuf** at offset **displs[j]**.

The type signature associated with **sendcount**, **sendtype** at process **j** must be equal to the type signature of **recvcounts[j]**, **recvtype** at any other process.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

None

RELATED INFORMATION

MPE_IALLGATHERV
MPI_ALLGATHER

NAME

MPI_Allreduce, MPI_ALLREDUCE-Applies a reduction operation to the vector **sendbuf** over the set of processes specified by **comm** and places the result in **recvbuf** on all of the processes in **comm**

C SYNOPSIS

```
#include <mpi.h>
int MPI_Allreduce(void* sendbuf,void* recvbuf,int count,
    MPI_Datatype datatype,MPI_Op op,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ALLREDUCE(CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,
    INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
recvbuf	is the starting address of the receive buffer (choice) (OUT)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of elements in the send buffer (handle) (IN)
op	is the reduction operation (handle) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine applies a reduction operation to the vector **sendbuf** over the set of processes specified by **comm** and places the result in **recvbuf** on all of the processes.

This routine is similar to MPI_REDUCE except the result is returned to the receive buffer of all the group members.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

NOTES

See Appendix D, “Reduction Operations” on page 277.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid op	
Invalid communicator	
Invalid communicator type	must be intracommunicator
Unequal message lengths	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent op

Inconsistent datatype

Inconsistent message length

RELATED INFORMATION

MPE_IALLREDUCE
MPI_REDUCE
MPI_REDUCE_SCATTER
MPI_OP_CREATE

NAME

MPI_Alltoall, MPI_ALLTOALL-Sends a distinct message from each process to every process.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Alltoall(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,MPI_Comm comm):
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ALLTOALL(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,
    INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements sent to each process (integer) (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcount	is the number of elements received from any process (integer) (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_ALLTOALL sends a distinct message from each process to every process.

The **j**th block of data sent from process **i** is received by process **j** and placed in the **i**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. This means the amount of data sent must be equal to the amount of data received, pair wise between every pair of processes. The type maps can be different.

All arguments on all processes are significant.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on

page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Unequal lengths

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid communicator

Invalid communicator type must be intracommunicator

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent message lengths

RELATED INFORMATION

MPE_IALLTOALL
MPI_ALLTOALLV

NAME

MPI_Alltoallv, MPI_ALLTOALLV-Sends a distinct message from each process to every process. Messages can have different sizes and displacements.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Alltoallv(void* sendbuf,int *sendcounts,int *sdispls,
    MPI_Datatype sendtype,void* recvbuf,int *recvcounts,int *rdispls,
    MPI_Datatype recvtype,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ALLTOALLV(CHOICE SENDBUF,INTEGER SENDCOUNTS(*),
    INTEGER SDISPLS(*),INTEGER SENDTYPE,CHOICE RECVBUF,
    INTEGER RECVCOUNTS(*),INTEGER RDISPLS(*),INTEGER RECVTYPE,
    INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcounts	integer array (of length group size) specifying the number of elements to send to each process (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**. (IN)
sendtype	is the datatype of the send buffer elements (handle) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcounts	integer array (of length group size) specifying the number of elements to be received from each process (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**. (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_ALLTOALLV sends a distinct message from each process to every process. Messages can have different sizes and displacements.

This routine is similar to MPI_ALLTOALL with the following differences. MPI_ALLTOALLV allows you the flexibility to specify the location of the data for the send with **sdispls** and the location of where the data will be placed on the receive with **rdispls**.

The block of data sent from process **i** is **sendcounts[j]** elements long, and is received by process **j** and placed in **recvbuf** at offset offset **rdispls[i]**. These blocks do not have to be the same size.

The type signature associated with **sendcount[j]**, **sendtype** at process **i** must be equal to the type signature associated with **recvcounts[i]**, **recvtype** at process **j**. This means the amount of data sent must be equal to the amount of data received, pair wise between every pair of processes. Distinct type maps between sender and receiver are allowed.

All arguments on all processes are significant.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid communicator

Invalid communicator type must be intracommunicator

A send and receive hand unequal message lengths

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPE_IALLTOALLV
MPI_ALLTOALL

NAME

MPI_Attr_delete, MPI_ATTR_DELETE-Removes an attribute value from a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Attr_delete(MPI_Comm comm,int keyval);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ATTR_DELETE(INTEGER COMM,INTEGER KEYVAL,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator that the attribute is attached (handle) (IN)
keyval	is the key value of the deleted attribute (integer) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine deletes an attribute from cache by key. MPI_ATTR_DELETE also invokes the attribute delete function **delete_fn** specified when the **keyval** is created.

ERRORS

A delete_fn did not return MPI_SUCCESS

Invalid communicator

Invalid keyval keyval is undefined

Invalid keyval keyval is predefined

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_KEYVAL_CREATE

NAME

MPI_Attr_get, MPI_ATTR_GET-Retrieves an attribute value from a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Attr_get(MPI_Comm comm,int keyval,void *attribute_val,
    int *flag);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ATTR_GET(INTEGER COMM,INTEGER KEYVAL,INTEGER ATTRIBUTE_VAL,
    LOGICAL FLAG,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator to which attribute is attached (handle) (IN)
keyval	is the key value (integer) (IN)
attribute_val	is the attribute value unless **flag** = **false** (OUT)
flag	is **true** if an attribute value was extracted and *false* if no attribute is associated with the key. (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This function retrieves an attribute value by key. If there is no key with value **keyval**, the call is erroneous. However, the call is valid if there is a key value **keyval**, but no attribute is attached on **comm** for that key. In this case, the call returns **flag** = **false**.

NOTES

The implementation of the MPI_ATTR_PUT and MPI_ATTR_GET involves saving a single word of information in the communicator. The languages C and FORTRAN have different approaches to using this capability:

In C: As the programmer, you normally define a struct which holds arbitrary "attribute" information. Before calling MPI_ATTR_PUT, you allocate some storage for the attribute structure and then call MPI_ATTR_PUT to record the address of this structure. You must assure that the structure remains intact as long as it may be useful. As the programmer, you will also declare a variable of type "pointer to attribute structure" and pass the address of this variable when calling MPI_ATTR_GET. Both MPI_ATTR_PUT and MPI_ATTR_GET take a void* parameter but this does not imply the same parameter is passed to either one.

In FORTRAN: MPI_ATTR_PUT records an INTEGER*4 and MPI_ATTR_GET returns the INTEGER*4. As the programmer, you may choose to encode all attribute information in this integer or maintain a some kind of database in which the

integer can index. Either of these approaches will port to other MPI implementations.

XL FORTRAN has an additional feature which will allow some of the same function a C programmer would use. This is the POINTER type which is described in the *IBM XL FORTRAN Compiler V3.2 for AIX Language Reference*. Use of this will impact the program's portability.

ERRORS

Invalid communicator	
Invalid keyval	keyval is undefined
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_ATTR_PUT

NAME

MPI_Attr_put, MPI_ATTR_PUT-Stores an attribute value in a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Attr_put(MPI_Comm comm,int keyval,void* attribute_val);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ATTR_PUT(INTEGER COMM,INTEGER KEYVAL,INTEGER ATTRIBUTE_VAL,INTEGER IERROR)
```

PARAMETERS

comm is the communicator to which attribute will be attached (handle) (IN)

keyval is the key value as returned by MPI_KEYVAL_CREATE (integer) (IN)

attribute_val is the attribute value (IN)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine stores the attribute value for retrieval by MPI_ATTR_GET. Any previous value is deleted with the attribute **delete_fn** being called and the new value is stored. If there is no key with value **keyval**, the call is erroneous.

NOTES

The implementation of the MPI_ATTR_PUT and MPI_ATTR_GET involves saving a single word of information in the communicator. The languages C and FORTRAN have different approaches to using this capability:

In C: As the programmer, you normally define a struct which holds arbitrary "attribute" information. Before calling MPI_ATTR_PUT, you allocate some storage for the attribute structure and then call MPI_ATTR_PUT to record the address of this structure. You must assure that the structure remains intact as long as it may be useful. As the programmer, you will also declare a variable of type "pointer to attribute structure" and pass the address of this variable when calling MPI_ATTR_GET. Both MPI_ATTR_PUT and MPI_ATTR_GET take a void* parameter, but this does not imply the same parameter is passed to either one.

In FORTRAN: MPI_ATTR_PUT records an INTEGER*4 and MPI_ATTR_GET returns the INTEGER*4. As the programmer, you may choose to encode all attribute information in this integer or maintain a some kind of database in which the integer can index. Either of these approaches will port to other MPI implementations.

XL FORTRAN has an additional feature which will allow some of the same function a C programmer would use. This is the POINTER type which is described in the

IBM XL FORTRAN Compiler V3.2 for AIX Language Reference. Use of this will impact the program's portability.

ERRORS

A delete_fn did not return MPI_SUCCESS

Invalid communicator

Invalid keyval keyval is undefined

Predefined keyval cannot modify predefined attributes

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ATTR_GET
MPI_KEYVAL_CREATE

NAME

MPI_Barrier, MPI_BARRIER-Blocks each process in **comm** until all processes have called it.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Barrier(MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_BARRIER(INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

comm is a communicator (handle) (IN)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine blocks until all processes have called it. Processes cannot exit the operation until all group members have entered.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPE_IBARRIER

NAME

MPI_Bcast, MPI_BCAST-Broadcasts a message from **root** to all processes in **comm**.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Bcast(void* buffer,int count,MPI_Datatype datatype,
    int root,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_BCAST(CHOICE BUFFER,INTEGER COUNT,INTEGER DATATYPE,INTEGER ROOT,
      INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

buffer	is the starting address of the buffer (choice) (INOUT)
count	is the number of elements in the buffer (integer) (IN)
datatype	is the datatype of the buffer elements (handle) (IN)
root	is the rank of the root task (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine broadcasts a message from **root** to all processes in **comm**. The contents of **root**'s communication buffer is copied to all processes on return.

The type signature of **count**, **datatype** on any process must be equal to the type signature of **count**, **datatype** at the root. This means the amount of data sent must be equal to the amount of data received, pair wise between each process and the root. Distinct type maps between sender and receiver are allowed.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count **count** < 0

Invalid datatype

Type not committed

Invalid root **root** < 0 or **root** >= groupsize

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

Inconsistent message length

RELATED INFORMATION

MPE_IBCAST

NAME

MPI_Bsend, MPI_BSEND-Blocking buffered mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Bsend(void* buf,int count,MPI_Datatype datatype,
      int dest,int tag,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_BSEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
      INTEGER TAG,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of destination (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a blocking buffered mode send. This is a local operation. It does not depend on the occurrence of a matching receive in order to complete. If a send operation is started and no matching receive is posted, the outgoing message is buffered to allow the send call to complete.

Make sure you have enough buffer space available. An error occurs if the message must be buffered and there is there is insufficient buffer space.

Return from an MPI_BSEND does not guarantee the message was sent. It may remain in the buffer until a matching receive is posted. MPI_BUFFER_DETACH will block until all messages are received.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0

Invalid comm

Insufficient buffer space

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_IBSEND
MPI_SEND
MPI_BUFFER_ATTACH
MPI_BUFFER_DETACH

NAME

MPI_Bsend_init, MPI_BSEND_INIT-Creates a persistent buffered mode send request.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Bsend_init(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_BSEND_INIT(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,
      INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,
      INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements to be sent (integer) (IN)
datatype	is the type of each element (handle) (IN)
dest	is the rank of the destination task (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a persistent communication request for a buffered mode send operation. MPI_START or MPI_STARTALL must be called to activate the send.

NOTES

See MPI_BSEND for additional information.

Because it is the MPI_START which initiates communication, any error related to insufficient buffer space occurs at the MPI_START.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0

Invalid comm

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_START
MPI_IBSEND

NAME

MPI_Buffer_attach, MPI_BUFFER_ATTACH-Provides MPI with a buffer to use for buffering messages sent with MPI_BSEND and MPI_IBSEND.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Buffer_attach(void* buffer,int size);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_BUFFER_ATTACH(CHOICE BUFFER,INTEGER SIZE,INTEGER IERROR)
```

PARAMETERS

buffer	is the initial buffer address (choice) (IN)
size	is the buffer size in bytes (integer) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine provides MPI a buffer in the user's memory which is used for buffering outgoing messages. This buffer is used only by messages sent in buffered mode, and only one buffer is attached to a process at any time.

NOTES

MPI uses part of the buffer space to store information about the buffered messages. The number of bytes required by MPI for each buffered message is given by MPI_BSEND_OVERHEAD.

If a buffer is already attached, it must be detached by MPI_BUFFER_DETACH before a new buffer can be attached.

ERRORS

Invalid size	**size** < 0
Buffer is already attached	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_BUFFER_DETACH
MPI_BSEND
MPI_IBSEND

NAME

MPI_Buffer_detach, MPI_BUFFER_DETACH-Detaches the current buffer.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Buffer_detach(void* buffer,int *size);
```

FORTRAN SYNOPSIS

```
MPI_BUFFER_DETACH(CHOICE BUFFER,INTEGER SIZE,INTEGER IERROR)
```

PARAMETERS

buffer	is the initial buffer address (choice) (OUT)
size	is the buffer size in bytes (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine detaches the current buffer. Blocking occurs until all messages in the active buffer are transmitted. Once this function returns, you can reuse or deallocate the space taken by the buffer.

If there is no active buffer, MPI acts as if a buffer of size 0 is associated with the process.

NOTES

It is important to detach an attached buffer *before* it is deallocated. If this is not done, any buffered message may be lost.

In FORTRAN 77, the **buffer** argument for MPI_BUFFER_DETACH cannot return a useful value because FORTRAN 77 does not support pointers. If a fully portable MPI program written in FORTRAN calls MPI_BUFFER_DETACH, it either passes the name of the original buffer or a throwaway temp as the **buffer** argument.

If a buffer was attached, this implementation of MPI returns the address of the freed buffer in the first word of the **buffer** argument. If the **size** being returned is zero to four bytes, MPI_BUFFER_DETACH will not modify the **buffer** argument. This implementation is harmless for a program that uses either the original buffer or a throwaway temp of at least word size as **buffer**. It also allows the programmer who wants to use an XL FORTRAN POINTER as the **buffer** argument to do so. Using the POINTER type will affect portability.

ERRORS

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_BUFFER_ATTACH
MPI_BSEND
MPI_IBSEND

NAME

MPI_Cancel, MPI_CANCEL-Marks a nonblocking point-to-point operation for cancellation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Cancel(MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CANCEL(INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

request is a communication request (handle) (IN)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine marks a pending nonblocking point-to-point operation for cancellation. The cancel call is local. It returns immediately; it can return even before the communication is actually cancelled. It is necessary to complete a communication marked for cancellation by using a call to MPI_REQUEST_FREE, MPI_WAIT, or MPI_TEST (or any of the derived operations).

You can use MPI_CANCEL to cancel a communication that uses a persistent request in the same way it is used for nonpersistent requests. A successful cancellation cancels the active communication, but not the request itself. After the call to MPI_CANCEL and the subsequent call to MPI_WAIT or MPI_TEST, the request becomes inactive and can be activated for a new communication. It is erroneous to cancel an inactive persistent request.

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

Either the cancellation succeeds or the communication succeeds, but not both. If a send is marked for cancellation, then either the send completes normally, in which case the message sent was received at the destination process, or the send is successfully cancelled, in which case no part of the message was received at the destination. Then, any matching receive has to be satisfied by another send. If a receive is marked for cancellation, then the receive completes normally or the receive is successfully cancelled, in which case no part of the receive buffer is altered. Then, any matching send has to be satisfied by another receive.

If the operation has been cancelled, then information to that effect is returned in the status argument of the operation that completes the communication, and may be retrieved by a call to MPI_TEST_CANCELLED.

Nonblocking collective communication requests cannot be cancelled.

ERRORS

Invalid request

CCL request

Cancel inactive persistent request

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TEST_CANCELLED
MPI_WAIT

NAME

MPI_Cart_coords, MPI_CART_COORDS-Translates process rank in a communicator into cartesian process coordinates.

C SYNOPSIS

```
#include <mpi.h>
MPI_Cart_coords(MPI_Comm comm,int rank,int maxdims,int *coords);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CART_COORDS(INTEGER COMM,INTEGER RANK,INTEGER MAXDIMS,
      INTEGER COORDS(*),INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with cartesian topology (handle) (IN)
rank	is the rank of a process within group **comm** (integer) (IN)
maxdims	is the length of array **coords** in the calling program (integer) (IN)
coords	is an integer array specifying the cartesian coordinates of a process. (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine translates process rank in a communicator into process coordinates.

NOTES

Process coordinates in a cartesian structure begin their numbering at 0. Row-major numbering is always used for the processes in a cartesian structure.

ERRORS

MPI not initialized	
MPI already finalized	
Invalid communicator	
No topology	
Invalid topology	type must be cartesian
Invalid rank	**rank** < 0 or **rank** > = groupsize
Invalid array size	**maxdims** < 0

RELATED INFORMATION

MPI_CART_RANK
MPI_CART_CREATE

NAME

MPI_Cart_create, MPI_CART_CREATE-Creates a communicator containing topology information.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Cart_create(MPI_Comm comm_old,int ndims,int *dims,
    int *periods,int reorder,MPI_Comm *comm_cart);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CART_CREATE(INTEGER COMM_OLD,INTEGER NDIMS,INTEGER DIMS(*),
    INTEGER PERIODS(*),INTEGER REORDER,INTEGER COMM_CART,INTEGER IERROR)
```

PARAMETERS

comm_old	is the input communicator (handle) (IN)
ndims	is the number of cartesian dimensions in grid (integer) (IN)
dims	is an integer array of size **ndims** specifying the number of processes in each dimension (IN)
periods	is a logical array of size **ndims** specifying if the grid is periodic or not in each dimension (IN)
reorder	if true, ranking may be reordered. If false, then rank in **comm_cart** must be the same as in **comm_old**. (logical) (IN)
comm_cart	is a communicator with new cartesian topology (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new communicator containing cartesian topology information defined by **ndims**, **dims**, **periods** and **reorder**. MPI_CART_CREATE returns a handle for this new communicator in **comm_cart**. If there are more tasks in **comm** than required by the grid, some tasks are returned **comm_cart** = MPI_COMM_NULL. **comm_old** must be an intracommunicator.

NOTES

The reorder argument is ignored.

ERRORS

MPI not initialized

MPI already finalized

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid ndims	**ndims** < 0 or **ndims** > groupsize
Invalid dimension	

RELATED INFORMATION

MPI_CART_SUB
MPI_GRAPH_CREATE

NAME

MPI_Cart_get, MPI_CART_GET-Retrieves cartesian topology information from a communicator.

C SYNOPSIS

```
#include <mpi.h>
MPI_Cart_get(MPI_Comm comm,int maxdims,int *dims,int *periods,int *coords);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CART_GET(INTEGER COMM,INTEGER MAXDIMS,INTEGER DIMS(*),
      INTEGER PERIODS(*),INTEGER COORDS(*),INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with cartesian topology (handle) (IN)
maxdims	is the length of **dims, periods,** and **coords** in the calling program (integer) (IN)
dims	is the number of processes for each cartesian dimension (array of integer) (OUT)
periods	is a logical array specifying if each cartesian dimension is periodic or not. (OUT)
coords	is the coordinates of the calling process in the cartesian structure (array of integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine retrieves the cartesian topology information associated with a communicator in **dims, periods** and **coords**.

ERRORS

MPI not initialized	
MPI already finalized	
Invalid communicator	
No topology	
Invalid topology type	must be cartesian
Invalid array size	**maxdims** < 0

RELATED INFORMATION

MPI_CARTDIM_GET
MPI_CART_CREATE

NAME

MPI_Cart_map, MPI_CART_MAP-Computes placement of processes on the physical machine.

C SYNOPSIS

```
#include <mpi.h>
MPI_Cart_map(MPI_Comm comm,int ndims,int *dims,int *periods,
      int *newrank);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CART_MAP(INTEGER COMM,INTEGER NDIMS,INTEGER DIMS(*),
      INTEGER PERIODS(*),INTEGER NEWRANK,INTEGER IERROR)
```

PARAMETERS

comm is the input communicator (handle) (IN)

ndims is the number of dimensions of the cartesian structure (integer) (IN)

dims is an integer array of size **ndims** specifying the number of processes in each coordinate direction (IN)

periods is a logical array of size **ndims** specifying the periodicity in each coordinate direction (IN)

newrank is the reordered rank or MPI_UNDEFINED if the calling process does not belong to the grid (integer) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_CART_MAP allows MPI to compute an optimal placement for the calling process on the physical machine by reordering the processes in **comm**.

NOTES

No reordering is done by this function; it would serve no purpose on an SP. MPI_CART_MAP returns **newrank** as the original rank of the calling process if it belongs to the grid, or MPI_UNDEFINED if it does not.

ERRORS

MPI not initialized

MPI already finalized

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid ndims **ndims** < 1 or **ndims** > groupsize

Invalid dimension	**ndims**[i] <= 0
Invalid grid size	**n**< 0 or **n** > groupsize, where **n** is the product of **dims[i]**

NAME

MPI_Cart_rank, MPI_CART_RANK-Translates process coordinates into a process rank.

C SYNOPSIS

```
#include <mpi.h>
MPI_Cart_rank(MPI_Comm comm,int *coords,int *rank);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CART_RANK(INTEGER COMM,INTEGER COORDS(*),INTEGER RANK,
      INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with cartesian topology (handle) (IN)
coords	is an integer array of size **ndims** specifying the cartesian coordinates of a process (IN)
rank	is an integer specifying the rank of specified process (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine translates cartesian process coordinates into a process rank.

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 ≥ **coords(i)** < **dims(i)** automatically. Out of range coordinates are erroneous for non-periodic dimensions.

NOTES

Process coordinates in a cartesian structure begin their numbering at 0. Row-major numbering is always used for the processes in a cartesian structure.

ERRORS

MPI not initialized	
MPI already finalized	
Invalid communicator	
No topology	
Invalid topology type	must be cartesian
Invalid coordinates	refer to Description above

RELATED INFORMATION

MPI_CART_CREATE
MPI_CART_COORDS

NAME

MPI_Cart_shift, MPI_CART_SHIFT-Returns shifted source and destination ranks for a process.

C SYNOPSIS

```
#include <mpi.h>
MPI_Cart_shift(MPI_Comm comm,int direction,int disp,
      int *rank_source,int *rank_dest);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CART_SHIFT(INTEGER COMM,INTEGER DIRECTION,INTEGER DISP,
      INTEGER RANK_SOURCE,INTEGER RANK_DEST,INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with cartesian topology (handle) (IN)
direction	is the coordinate dimension of shift (integer) (IN)
disp	is the displacement (>0 = upward shift, <0 = downward shift) (integer) (IN)
rank_source	is the rank of the source process (integer) (OUT)
rank_dest	is the rank of the destination process (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine shifts the local rank along a specified coordinate dimension to generate source and destination ranks.

rank_source is obtained by subtracting **disp** from the **n**th coordinate of the local process, where **n** is equal to **direction**. Similarly, **rank_dest** is obtained by adding **disp** to the **n**th coordinate. Coordinate dimensions (**direction**) are numbered starting with **0**.

If the dimension specified by **direction** is non-periodic, off-end shifts result in the value MPI_PROC_NULL being returned for **rank_source** and/or **rank_dest**.

NOTES

In C and FORTRAN, the coordinate is identified by counting from 0. For example, FORTRAN **A(X,Y)** or C **A[x] [y]** both have **x** as direction 0.

ERRORS

MPI not initialized	
MPI already finalized	
Invalid communicator	
Invalid topology type	must be cartesian
No topology	

RELATED INFORMATION

MPI_CART_RANK
MPI_CART_COORDS
MPI_CART_CREATE

NAME

MPI_Cart_sub, MPI_CART_SUB-Partitions a cartesian communicator into lower-dimensional subgroups.

C SYNOPSIS

```
#include <mpi.h>
MPI_Cart_sub(MPI_Comm comm,int *remain_dims,MPI_Comm *newcomm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CART_SUB(INTEGER COMM,LOGICAL REMAIN_DIMS(*),INTEGER NEWCOMM,
      INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with cartesian topology (handle) (IN)
remain_dims	the **i**th entry of **remain_dims** specifies whether the **i**th dimension is kept in the subgrid or is dropped. (logical vector) (IN)
newcomm	is the communicator containing the subgrid that includes the calling process (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

If a cartesian topology was created with MPI_CART_CREATE, you can use the function MPI_CART_SUB:

- to partition the communicator group into subgroups forming lower-dimensional cartesian subgrids, and
- to build a communicator with the associated subgrid cartesian topology for each of those subgroups.

(This function is closely related to MPI_COMM_SPLIT.)

For example, MPI_CART_CREATE (..., **comm**) defined a 2 × 3 × 4 grid. Let **remain_dims** = (true, false, true). Then a call to:

```
MPI_CART_SUB(comm,remain_dims,comm_new),
```

creates three communicators. Each has 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),
```

creates six non-overlapping communicators, each with four processes in a one-dimensional cartesian topology.

ERRORS

MPI not initialized	
MPI already finalized	
Invalid communicator	
Invalid topology	must be cartesian
No topology	

RELATED INFORMATION

MPI_CART_CREATE
MPI_COMM_SPLIT

NAME

MPI_Cartdim_get, MPI_CARTDIM_GET-Retrieves the number of cartesian dimensions from a communicator.

C SYNOPSIS

```
#include <mpi.h>
MPI_Cartdim_get(MPI_Comm comm,int *ndims);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_CARTDIM_GET(INTEGER COMM,INTEGER NDIMS,INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with cartesian topology (handle) (IN)
ndims	is an integer specifying the number of dimensions of the cartesian topology (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine retrieves the number of dimensions in a cartesian topology.

ERRORS

Invalid communicator	
No topology	
Invalid topology type	must be cartesian
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_CART_GET
MPI_CART_CREATE

NAME

MPI_Comm_compare, MPI_COMM_COMPARE-Compares the groups and context of two communicators.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_compare(MPI_Comm comm1,MPI_Comm comm2,int *result);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_COMPARE(INTEGER COMM1,INTEGER COMM2,INTEGER RESULT,INTEGER IERROR)
```

PARAMETERS

comm1 is the first communicator (handle) (IN)

comm2 is the second communicator (handle) (IN)

result is an integer specifying the result. The defined values are: MPI_IDENT, MPI_CONGRUENT, MPI_SIMILAR, and MPI_UNEQUAL. (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine compares the groups and contexts of two communicators. The following is an explanation of each MPI_COMM_COMPARE defined value:

MPI_IDENT comm1 and **comm2** are handles for the identical object

MPI_CONGRUENT the underlying groups are identical in constituents and rank order (both local and remote groups for intercommunications), but are different in context

MPI_SIMILAR the group members of both communicators are the same but are different in rank order (both local and remote groups for intercommunications),

MPI_UNEQUAL if none of the above.

ERRORS

Invalid communicator(s)

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GROUP_COMPARE

NAME

MPI_Comm_create, MPI_COMM_CREATE-Creates a new intracommunicator with a given group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_create(MPI_Comm comm,MPI_Group group,MPI_Comm *newcomm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_CREATE(INTEGER COMM,INTEGER GROUP,INTEGER NEWCOMM,
      INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (IN)
group	is Group which is a subset of the group of **comm** (handle) (IN)
newcomm	is the new communicator (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_COMM_CREATE is a collective function that is invoked by all processes in the group associated with **comm**. This routine creates a new intracommunicator **newcomm** with communication group defined by **group** and a new context. Cached information is not propagated from **comm** to **newcomm**.

For processes that are not in **group**, MPI_COMM_NULL is returned. The call is erroneous if **group** is not a subset of the group associated with **comm**. The call is executed by all processes in **comm** even if they do not belong to the new group.

This call applies only to intracommunicators.

NOTES

MPI_COMM_CREATE provides a way to subset a group of processes for the purpose of separate MIMD computation with separate communication space. You can use **newcomm** in subsequent calls to MPI_COMM_CREATE or other communicator constructors to further subdivide a computation into parallel sub-computations.

ERRORS

Invalid communicator	
Invalid group	**group** is not a subset of the group associated with **comm**
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_COMM_DUP
MPI_COMM_SPLIT

NAME

MPI_Comm_dup, MPI_COMM_DUP-Creates a new communicator that is a duplicate of an existing communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_dup(MPI_Comm comm,MPI_Comm *newcomm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_DUP(INTEGER COMM,INTEGER NEWCOMM,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (IN)
newcomm	is the copy of **comm** (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_COMM_DUP is a collective function that is invoked by the group associated with **comm**. This routine duplicates the existing communicator **comm** with its 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 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 and any copied cached information, but a new context.

This call applies to both intra and inter communicators.

NOTES

Use this operation to produce a duplicate communication space that has the same properties as the original communicator. This includes attributes and topologies.

This call is valid even if there are pending point to point communications involving the communicator **comm**.

ERRORS

A copy_fn did not return MPI_SUCCESS

A delete_fn did not return MPI_SUCCESS

Invalid communicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_KEYVAL_CREATE

NAME

MPI_Comm_free, MPI_COMM_FREE-Marks a communicator for deallocation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_free(MPI_Comm *comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_FREE(INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

comm is the communicator to be freed (handle) (INOUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This collective function marks either an intra or an inter communicator object for deallocation. MPI_COMM_FREE sets the handle to MPI_COMM_NULL. Actual deallocation of the communicator object occurs when active references to it have completed. The delete callback functions for all cached attributes are called in arbitrary order. The delete functions are called immediately and not deferred until deallocation.

ERRORS

A delete_fn did not return MPI_SUCCESS

Invalid communicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_KEYVAL_CREATE

NAME

MPI_Comm_group, MPI_COMM_GROUP-Returns the group handle associated with a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_group(MPI_Comm comm,MPI_Group *group);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_GROUP(INTEGER COMM,INTEGER GROUP,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (IN)
group	is the group corresponding to **comm** (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the group handle associated with a communicator.

NOTES

If **comm** is an intercommunicator, then **group** is set to the local group. To determine the remote group of an intercommunicator, use MPI_COMM_REMOTE_GROUP.

ERRORS

Invalid communicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_COMM_REMOTE_GROUP

NAME

MPI_Comm_rank, MPI_COMM_RANK-Returns the rank of the local process in the group associated with a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_rank(MPI_Comm comm,int *rank);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_RANK(INTEGER COMM,INTEGER RANK,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (IN)
rank	is an integer specifying the rank of the calling process in group of **comm** (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the rank of the local process in the group associated with a communicator.

You can use this routine with MPI_COMM_SIZE to determine the amount of concurrency available for a specific job. MPI_COMM_RANK indicates the rank of the process that calls it in the range from 0...**size**-1, where **size** is the return value of MPI_COMM_SIZE.

This routine is a shortcut to accessing the communicator's group with MPI_COMM_GROUP, computing the rank using MPI_GROUP_RANK and freeing the temporary group by using MPI_GROUP_FREE.

If **comm** is an intercommunicator, **rank** is the rank of the local process in the local group.

ERRORS

Invalid communicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GROUP_RANK

NAME

MPI_Comm_remote_group, MPI_COMM_REMOTE_GROUP-Returns the handle of the remote group of an intercommunicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_remote_group(MPI_Comm comm,MPI_group *group);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_REMOTE_GROUP(INTEGER COMM,MPI_GROUP GROUP,INTEGER IERROR)
```

PARAMETERS

comm	is the intercommunicator (handle) (IN)
group	is the remote group corresponding to **comm**. (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a local operation that returns the handle of the remote group of an intercommunicator.

NOTES

To determine the local group of an intercommunicator, use MPI_COMM_GROUP.

ERRORS

Invalid communicator

Invalid communicator type it must be intercommunicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_COMM_GROUP

NAME

MPI_Comm_remote_size, MPI_COMM_REMOTE_SIZE-Returns the size of the remote group of an intercommunicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_remote_size(MPI_Comm comm,int *size);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_REMOTE_SIZE(INTEGER COMM,INTEGER SIZE,INTEGER IERROR)
```

PARAMETERS

comm	is the intercommunicator (handle) (IN)
size	is an integer specifying the number of processes in the remote group of **comm**. (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a local operation that returns the size of the remote group of an intercommunicator.

NOTES

To determine the size of the local group of an intercommunicator, use MPI_COMM_SIZE.

ERRORS

Invalid communicator

Invalid communicator type it must be intercommunicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_COMM_SIZE

NAME

MPI_Comm_size, MPI_COMM_SIZE-Returns the size of the group associated with a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_size(MPI_Comm comm,int *size);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_SIZE(INTEGER COMM,INTEGER SIZE,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (IN)
size	is an integer specifying the number of processes in the group of **comm** (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the size of the group associated with a communicator. This routine is a shortcut to:

- accessing the communicator's group with MPI_COMM_GROUP,
- computing the size using MPI_GROUP_SIZE, and
- freeing the temporary group using MPI_GROUP_FREE.

If **comm** is an intercommunicator, **size** will be the size of the local group. To determine the size of the remote group of an intercommunicator, use MPI_COMM_REMOTE_SIZE.

You can use this routine with MPI_COMM_RANK to determine the amount of concurrency available for a specific library or program. MPI_COMM_RANK indicates the rank of the process that calls it in the range from 0...**size**-1, where **size** is the return value of MPI_COMM_SIZE. The rank and size information can then be used to partition work across the available processes.

NOTES

This function indicates the number of processes in a communicator. For MPI_COMM_WORLD, it indicates the total number of processes available.

ERRORS

Invalid communicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GROUP_SIZE
MPI_COMM_GROUP
MPI_COMM_RANK
MPI_COMM_REMOTE_SIZE
MPI_GROUP_FREE

NAME

MPI_Comm_split, MPI_COMM_SPLIT-Splits a communicator into multiple communicators based on **color** and **key**.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_split(MPI_Comm comm,int color,int key,MPI_Comm *newcomm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_SPLIT(INTEGER COMM,INTEGER COLOR,INTEGER KEY,
    INTEGER NEWCOMM,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (IN)
color	is an integer specifying control of subset assignment (IN)
key	is an integer specifying control of rank assignment (IN)
newcomm	is the new communicator (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_COMM_SPLIT is a collective function that partitions the group associated with **comm** into disjoint subgroups, one for each value of **color**. Each subgroup contains all processes of the same color. Within each subgroup, the processes are ranked in the order defined by the value of the argument **key**. Ties are broken according to their rank in the old group. A new communicator is created for each subgroup and returned in **newcomm**. If a process supplies the color value MPI_UNDEFINED, **newcomm** returns MPI_COMM_NULL. Even though this is a collective call, each process is allowed to provide different values for **color** and **key**.

This call applies only to intracommunicators.

The value of **color** must be greater than or equal to zero.

ERRORS

Invalid color	**color** < 0
Invalid communicator	
Invalid communicator type	it must be intracommunicator
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_CART_SUB

NAME

MPI_Comm_test_inter, MPI_COMM_TEST_INTER-Returns the type of a communicator (intra or inter).

C SYNOPSIS

```
#include <mpi.h>
int MPI_Comm_test_inter(MPI_Comm comm,int *flag);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_COMM_TEST_INTER(INTEGER COMM,LOGICAL FLAG,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (INOUT)
flag	communicator type (logical)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is used to determine if a communicator is an inter or an intracommunicator.

If **comm** is an intercommunicator, the call returns **true**. If **comm** is an intracommunicator, the call returns **false**.

NOTES

An intercommunicator can be used as an argument to some of the communicator access routines. However, intercommunicators cannot be used as input to some of the constructor routines for intracommunicators, such as MPI_COMM_CREATE.

ERRORS

Invalid communicator

MPI not initialized

MPI already finalized

NAME

MPI_Dims_create, MPI_DIMS_CREATE-Defines a cartesian grid to balance processes.

C SYNOPSIS

```
#include <mpi.h>
MPI_Dims_create(int nnodes,int ndims,int *dims);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_DIMS_CREATE(INTEGER NNODES,INTEGER NDIMS,INTEGER DIMS(*),
      INTEGER IERROR)
```

PARAMETERS

nnodes	is an integer specifying the number of nodes in a grid (IN)
ndims	is an integer specifying the number of cartesian dimensions (IN)
dims	is an integer array of size **ndims** that specifies the number of nodes in each dimension. (INOUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a cartesian grid with a given number of dimensions and a given number of nodes. The dimensions are constrained to be as close to each other as possible.

If **dims[i]** is a positive number when MPI_DIMS_CREATE is called, the routine will not modify the number of nodes in dimension **i**. Only those entries where **dims[i]=0** are modified by the call.

NOTES

MPI_DIMS_CREATE chooses dimensions so that the resulting grid is as close as possible to being an **ndims**– dimensional **cube**.

ERRORS

MPI not initialized	
MPI already finalized	
Invalid ndims	**ndims** < 0
Invalid nnodes	**nnodes**<0
Invalid dimension	**dims[i]** < 0 or **nnodes** is not a multiple of the non-zero entries of **dims**

RELATED INFORMATION

MPI_CART_CREATE

NAME

MPI_Errhandler_create, MPI_ERRHANDLER_CREATE-Registers a user defined error handler.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Errhandler_create(MPI_Handler_function *function,
    MPI_Errhandler *errhandler);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ERRHANDLER_CREATE(EXTERNAL FUNCTION,INTEGER ERRHANDLER,
      INTEGER IERROR)
```

PARAMETERS

function is a user defined error handling procedure (IN)

errhandler is an MPI error handler (handle) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_ERRHANDLER_CREATE registers the user routine **function** for use as an MPI error handler.

You can associate an error handler with a communicator. MPI will use the specified error handling routine for any exception that takes place during a call on this communicator. Different processes can attach different error handlers to the same communicator. MPI calls not related to a specific communicator are considered as attached to the communicator MPI_COMM_WORLD.

ERRORS

NULL function

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ERRHANDLER_SET
MPI_ERRHANDLER_GET
MPI_ERRHANDLER_FREE

NAME

MPI_Errhandler_free, MPI_ERRHANDLER_FREE-Marks an error handler for deallocation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Errhandler_free(MPI_Errhandler *errhandler);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ERRHANDLER_FREE(INTEGER ERRHANDLER,INTEGER IERROR)
```

PARAMETERS

errhandler is an MPI error handler (handle) (IN)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine marks error handler **errhandler** for deallocation and sets **errhandler** to MPI_ERRHANDLER_NULL. Actual deallocation occurs when all communicators associated with the error handler have been deallocated.

ERRORS

Invalid error handler

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ERRHANDLER_CREATE

NAME

MPI_Errhandler_get, MPI_ERRHANDLER_GET- Gets an error handler associated with a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Errhandler_get(MPI_Comm comm,MPI_Errhandler *errhandler);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ERRHANDLER_GET(INTEGER COMM,INTEGER ERRHANDLER,INTEGER IERROR)
```

PARAMETERS

comm is a communicator (handle) (IN)

errhandler is the MPI error handler currently associated with **comm** (handle) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the error handler **errhandler** currently associated with communicator **comm**.

ERRORS

Invalid communicator

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ERRHANDLER_SET
MPI_ERRHANDLER_CREATE

NAME

MPI_Errhandler_set, MPI_ERRHANDLER_SET- Associates a new error handler with a communicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Errhandler_set(MPI_Comm comm,MPI_Errhandler errhandler);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ERRHANDLER_SET(INTEGER COMM, INTEGER ERRHANDLER, INTEGER IERROR)
```

PARAMETERS

comm	is a communicator (handle) (IN)
errhandler	is a new MPI error handler for **comm** (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine associates error handler **errhandler** with communicator **comm**. The association is local.

MPI will use the specified error handling routine for any exception that takes place during a call on this communicator. Different processes can attach different error handlers to the same communicator. MPI calls not related to a specific communicator are considered as attached to the communicator MPI_COMM_WORLD.

NOTES

An error handler that does not end in the MPI job being terminated, creates undefined risks. Some errors are harmless while others are catastrophic. For example, an error detected by one member of a collective operation can result in others members waiting indefinitely for an operation which will never occur.

It is also important to note that the MPI standard does not specify the state the MPI library should be in after an error occurs. MPI does not provide a way for users to determine how much, if any, damage has been done to the MPI state by a particular error.

The default error handler is MPI_ERRORS_ARE_FATAL, which behaves as if it contains a call to MPI_ABORT. MPI_ERRHANDLER_SET allows users to replace MPI_ERRORS_ARE_FATAL with an alternate error handler. The MPI standard provides MPI_ERRORS_RETURN and IBM adds the non-standard MPE_ERRORS_WARN. These are pre-defined handlers that cause the error code to be returned and MPI to continue to run. Error handlers that are written by MPI users may call MPI_ABORT. If they do not abort, they too will cause MPI to deliver an error return code to the caller and continue to run.

Error handlers that let MPI return should only be used if every MPI call checks its return code. Continuing to use MPI after an error involves undefined risks. You may do cleanup after an MPI error is detected, as long as it doesn't use MPI calls. This should normally be followed by a call to MPI_ABORT.

ERRORS

Invalid Communicator

Invalid error handler

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ERRHANDLER_GET
MPI_ERRHANDLER_CREATE

NAME

MPI_Error_class, MPI_ERROR_CLASS-Returns the error class for the corresponding error code.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Error_class(int errorcode,int *errorclass);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ERROR_CLASS(INTEGER ERRORCODE,INTEGER ERRORCLASS,INTEGER IERROR)
```

PARAMETERS

errorcode is the error code returned by an MPI routine (IN)

errorclass is the error class for the **errorcode** (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the error class corresponding to an error code.

The valid error classes are:

Error Classes	Description
MPI_SUCCESS	No error
MPI_ERR_BUFFER	Invalid buffer pointer
MPI_ERR_COUNT	Invalid count argument
MPI_ERR_TYPE	Invalid datatype argument
MPI_ERR_TAG	Invalid tag argument
MPI_ERR_COMM	Invalid communicator
MPI_ERR_RANK	Invalid rank
MPI_ERR_REQUEST	Invalid request (handle)
MPI_ERR_ROOT	Invalid root
MPI_ERR_GROUP	Invalid group
MPI_ERR_OP	Invalid operation
MPI_ERR_TOPOLOGY	Invalid topology
MPI_ERR_DIMS	Invalid dimension argument
MPI_ERR_ARG	Invalid argument
MPI_ERR_UNKNOWN	Unknown error
MPI_ERR_TRUNCATE	Message truncated on receive
MPI_ERR_OTHER	Known error not provided

Error Classes	Description
MPI_ERR_INTERN	Internal MPI error
MPI_ERR_LASTCODE	Last standard error code
MPI_ERR_IN_STATUS	Error code is in status
MPI_PENDING	Pending request

NOTES

For this implementation of MPI refer to the IBM Parallel Environment for AIX: Messages which provides a listing of all the error messages issued as well as the error class to which the message belongs. Be aware that the MPI standard is not explicit enough about error classes to guarantee that every implementation of MPI will use the same error class for every detectable user error.

ERRORS

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ERROR_STRING

NAME

MPI_Error_string, MPI_ERROR_STRING -Returns the error string for a given error code.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Error_string(int errorcode,char *string,
      int *resultlen);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ERROR_STRING(INTEGER ERRORCODE,CHARCTER STRING(*),
     INTEGER RESULTLEN,INTEGER IERROR)
```

PARAMETERS

errorcode	is the error code returned by an MPI routine (IN)
string	is the error message for the **errorcode** (OUT)
resultlen	is the character length of **string** (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the error string for a given error code. The returned **string** is null terminated with the terminating byte not counted in **resultlen**.

Storage for **string** must be at least MPI_MAX_ERROR_STRING characters long. The number of characters actually written is returned in **resultlen**.

ERRORS

Invalid error code	**errorcode** is not defined
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_ERROR_CLASS

NAME

MPI_Finalize, MPI_FINALIZE-Terminates all MPI processing.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Finalize(void);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_FINALIZE(INTEGER IERROR)
```

PARAMETERS

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

Make sure this routine is the last MPI call. Any MPI calls made after MPI_FINALIZE are signalled as errors. You must be sure all pending communications involving a process have completed before the process calls MPI_FINALIZE.

Although MPI_FINALIZE terminates MPI processing, it does not terminate the process. It is possible to continue with non-MPI processing after calling MPI_FINALIZE, but no other MPI calls (including MPI_INIT) can be made.

In a threaded environment both MPI_INIT and MPI_FINALIZE must be called on the same thread. MPI_FINALIZE closes the communication library and terminates the service threads. It does not affect any threads you created, other than returning an error if one subsequently makes an MPI call. If you had registered a SIGIO handler, it is restored as a signal handler; however, the SIGIO signal is blocked when MPI_FINALIZE returns. If you want to catch SIGIO after MPI_FINALIZE has been called, you should unblock it.

NOTES

The MPI 1.1 standard does not specify the state of MPI processes after MPI_FINALIZE, therefore, an assumption that all processes continue may not be portable.

ERRORS

MPI already finalized

MPI not initialized

Develop mode error if:

Buffer still attached

RELATED INFORMATION

MPI_INIT
MPI_ABORT

NAME

MPI_Gather, MPI_GATHER-Collects individual messages from each process in **comm** at the **root** process.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Gather(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,
    MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GATHER(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,
    INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
sendcount	is the number of elements in the send buffer (integer) (IN)
sendtype	is the datatype of the send buffer elements (integer) (IN)
recvbuf	is the address of the receive buffer (choice, significant only at **root**) (OUT)
recvcount	is the number of elements for any single receive (integer, significant only at **root**) (IN)
recvtype	is the datatype of the receive buffer elements (handle, significant only at **root**) (IN)
root	is the rank of the receiving process (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine collects individual messages from each process in **comm** at the **root** process and stores them in rank order.

The type signature of **sendcount**, **sendtype** on process **i** must be equal to the type signature of **recvcount**, **recvtype** at the root. This means 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 allowed.

The following is information regarding MPI_GATHER arguments and processes:

- On the process **root**, all arguments to the function are significant.

- On other processes, only the arguments **sendbuf, sendcount, sendtype, root,** and **comm** are significant.
- The argument **root** must be the same on all processes.

Note that the argument **revcount** at the root indicates the number of items it receives from each process. It is not the total number of items received.

A call where the specification of counts and types causes any location on the root to be written more than once is erroneous.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid root **root** < 0 or **root** >= groupsize

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

Inconsistent message lengths

RELATED INFORMATION

MPE_IGATHER
MPI_SCATTER
MPI_GATHER
MPI_ALLGATHER

NAME

MPI_Gatherv, MPI_GATHERV-Collects individual messages from each process in **comm** at the **root** process. Messages can have different sizes and displacements.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Gatherv(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcounts,int *displs,MPI_Datatype recvtype,
    int root,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GATHERV(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),
    INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf is the starting address of the send buffer (choice) (IN)

sendcount is the number of elements in the send buffer (integer) (IN)

sendtype is the datatype of the send buffer elements (handle) (IN)

recvbuf is the address of the receive buffer (choice, significant only at **root**) (OUT)

recvcounts integer array (of length group size) that contains the number of elements received from each process (significant only at **root**) (IN)

displs integer array (of length group size). Entry **i** specifies the displacement relative to **recvbuf** at which to place the incoming data from process **i** (significant only at **root**) (IN)

recvtype is the datatype of the receive buffer elements (handle, significant only at **root**) (IN)

root is the rank of the receiving process (integer) (IN)

comm is the communicator (handle) (IN)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine collects individual messages from each process in **comm** at the **root** process and stores them in rank order. With **recvcounts** as an array, messages can have varying sizes, and **displs** allows you the flexibility of where the data is placed on the root.

The type signature of **sendcount**, **sendtype** on process **i** must be equal to the type signature of **recvcounts[i]**, **recvtype** at the root. This means 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 allowed.

The following is information regarding MPI_GATHERV arguments and processes:

- On the process **root**, all arguments to the function are significant.
- On other processes, only the arguments. **sendbuf, sendcount, sendtype, root,** and **comm** are significant.
- The argument **root** must be the same on all processes.

A call where the specification of sizes, types and displacements causes any location on the root to be written more than once is erroneous.

NOTES

Displacements are expressed as elements of type **recvtype**, not as bytes.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid root **root** < 0 or **root** >= groupsize

A send and receive have unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

RELATED INFORMATION

MPE_IGATHER
MPI_GATHER

NAME

MPI_Get_count, MPI_GET_COUNT-Returns the number of elements in a message.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Get_count(MPI_Status *status,MPI_Datatype datatype,
    int *count);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GET_COUNT(INTEGER STATUS(MPI_STATUS_SIZE),INTEGER DATATYPE,
    INTEGER COUNT,INTEGER IERROR)
```

PARAMETERS

status is a status object (status) (IN)
Note that in FORTRAN a single status object is an array of integers.

datatype is the datatype of each message element (handle) (IN)

count is the number of elements (integer) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This subroutine returns the number of elements in a message. The **datatype** argument and the argument provided by the call that set the **status** variable should match.

When one of the MPI wait or test calls returns **status** for a non-blocking operation request and the corresponding blocking operation does not provide a **status** argument, the **status** from this wait/test does not contain meaningful source, tag or message size information.

ERRORS

Invalid datatype

Type not committed

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_IRECV
MPI_WAIT
MPI_RECV
MPI_PROBE

NAME

MPI_Get_elements, MPI_GET_ELEMENTS-Returns the number of basic elements in a message.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Get_elements(MPI_Status *status,MPI_Datatype datatype,
    int *count);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GET_ELEMENTS(INTEGER STATUS(MPI_STATUS_SIZE),INTEGER DATATYPE,
    INTEGER COUNT,INTEGER IERROR)
```

PARAMETERS

status is a status of object (status) (IN)
Note that in FORTRAN a single status object is an array of integers.

datatype is the datatype used by the operation (handle) (IN)

count is an integer specifying the number of basic elements (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the number of type map elements in a message. When the number of bytes does not align with the type signature, MPI_GET_ELEMENTS returns MPI_UNDEFINED. For example, given type signature (int, short, int, short) a 10 byte message would return 3 while an 8 byte message would return MPI_UNDEFINED.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argument, the **status** from this wait/test does not contain meaningful source, tag or message size information.

ERRORS

Invalid datatype

Type is not committed

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GET_COUNT

NAME

MPI_Get_processor_name, MPI_GET_PROCESSOR_NAME- Returns the name of the local processor.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Get_processor_name(char *name,int *resultlen);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GET_PROCESSOR_NAME(CHARACTER NAME(*),INTEGER RESULTLEN,
      INTEGER IERROR)
```

PARAMETERS

name is a unique specifier for the actual node (OUT)

resultlen specifies the printable character length of the result returned in **name** (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the name of the local processor at the time of the call. The name is a character string from which it is possible to identify a specific piece of hardware. **name** represents storage that is at least MPI_MAX_PROCESSOR_NAME characters long and MPI_GET_PROCESSOR_NAME can write up to this many characters in **name**.

The actual number of characters written is returned in **resultlen**. The returned **name** is a null terminated C string with the terminating byte not counted in **resultlen**.

ERRORS

MPI not initialized

MPI already finalized

NAME

MPI_Get_version, MPI_GET_VERSION-returns the version of MPI standard supported in this release.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Get_version(int *version,int *subversion);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GET_VERSION(INTEGER VERSION, INTEGER SUBVERSION, NTEGER IERROR)
```

PARAMETERS

version	MPI standard version number (integer) (OUT)
subversion	MPI standard subversion number (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is used to determine the version of the MPI standard supported by the MPI implementation. This function is new for this release of IBM Parallel Environment for AIX(PE) MPI. Previous releases of the PE MPI supported MPI 1.1. MPI_GET_VERSION is part of the MPI 1.2 standard under consideration at the time of this writing.

There are also new symbolic constants, MPI_VERSION and MPI_SUBVERSION, provided in mpi.h and mpif.h that provide similar compile-time information.

MPI_GET_VERSION can be called before MPI_INIT.

NAME

MPI_Graph_create, MPI_GRAPH_CREATE-Creates a new communicator containing graph topology information.

C SYNOPSIS

```
#include <mpi.h>
MPI_Graph_create(MPI_Comm comm_old,int nnodes, int *index,
    int *edges,int reorder,MPI_Comm *comm_graph);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GRAPH_CREATE(INTEGER COMM_OLD,INTEGER NNODES,INTEGER INDEX(*),
    INTEGER EDGES(*),INTEGER REORDER,INTEGER COMM_GRAPH,
    INTEGER IERROR)
```

PARAMETERS

comm_old	is the input communicator (handle) (IN)
nnodes	is an integer specifying the number of nodes in the graph (IN)
index	is an array of integers describing node degrees (IN)
edges	is an array of integers describing graph edges (IN)
reorder	if true, ranking may be reordered (logical) (IN)
comm_graph	is the communicator with the graph topology added (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new communicator containing graph topology information provided by **nnodes, index, edges,** and **reorder**. MPI_GRAPH_CREATE returns the handle for this new communicator in **comm_graph**.

If there are more tasks in **comm_old** then **nnodes**, some tasks are returned **comm_graph** as MPI_COMM_NULL.

NOTES

The reorder argument is currently ignored.

The following is an example showing how to define the arguments **nnodes**, **index**, and **edges**. 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:

Argument	Input
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 ≤ **j** ≤ **index[0]–1** and the list of neighbors of node **i**, **i** > **0**, is stored in **edges[j]**, **index[i–1]** ≤ **j** ≤ **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** ≤ **j** ≤ **index(1)** and the list of neighbors of node **i**, **i** > 0, is stored in **edges(j)**, **index(i)+1** ≤ **j** ≤ **index(i+1)**.

Observe that because node 0 indicates node 1 is a neighbor, that node 1 must indicate that node 0 is its' neighbor. For any edge A→B the edge B→A must also be specified.

ERRORS

MPI not initialized

MPI already finalized

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid nnodes **nnodes**<0 or **nnodes** > groupsize

Invalid node degree (**index[i]**-**index[i-1]**) < 0

Invalid neighbor **edges[i]** < 0 or **edges[i]**>=**nnodes**

Asymmetric graph

RELATED INFORMATION

MPI_CART_CREATE

NAME

MPI_Graph_get, MPI_GRAPH_GET-Retrieves graph topology information from a communicator.

C SYNOPSIS

```
#include <mpi.h>
MPI_Graph_get(MPI_Comm comm,int maxindex,int maxedges,
      int *index,int *edges);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GRAPH_GET(INTEGER COMM,INTEGER MAXINDEX,INTEGER MAXEDGES,
      INTEGER INDEX(*),INTEGER EDGES(*),INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with graph topology (handle) (IN)
maxindex	is an integer specifying the length of **index** in the calling program (IN)
maxedges	is an integer specifying the length of **edges** in the calling program (IN)
index	is an array of integers containing node degrees (OUT)
edges	is an array of integers containing node neighbors (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine retrieves the **index** and **edges** graph topology information associated with a communicator.

ERRORS

MPI not initialized	
MPI already finalized	
Invalid communicator	
No topology	
Invalid topology type	topology type must be graph
Invalid array size	**maxindex** < 0 or **maxedges** < 0

RELATED INFORMATION

MPI_GRAPHDIMS_GET
MPI_GRAPH_CREATE

NAME

MPI_Graph_map, MPI_GRAPH_MAP-Computes placement of processes on the physical machine.

C SYNOPSIS

```
#include <mpi.h>
MPI_Graph_map(MPI_Comm comm,int nnodes,int *index,int *edges,int *newrank);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GRAPH_MAP(INTEGER COMM,INTEGER NNODES,INTEGER INDEX(*),
              INTEGER EDGES(*),INTEGER NEWRANK,INTEGER IERROR)
```

PARAMETERS

comm	is the input communicator (handle) (IN)
nnodes	is the number of graph nodes (integer) (IN)
index	is an integer array specifying node degrees (IN)
edges	is an integer array specifying node adjacency (IN)
newrank	is the reordered rank,or MPI_Undefined if the calling process does not belong to the graph (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_GRAPH_MAP allows MPI to compute an optimal placement for the calling process on the physical machine by reordering the processes in **comm**.

NOTES

MPI_CART_MAP returns **newrank** as the original rank of the calling process if it belongs to the grid or MPI_UNDEFINED if it does not. Currently, no reordering is done by this function.

ERRORS

Invalid communicator	
Invalid communicator type	must be intracommunicator
Invalid nnodes	**nnodes** <0 or **nnodes** > groupsize
Invalid node degree	**index[i]** < 0
Invalid neighbors	**edges[i]** < 0 or **edges[i]** >= **nnodes**
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_GRAPH_CREATE
MPI_CART_MAP

NAME

MPI_Graph_neighbors, MPI_GRAPH_NEIGHBORS-Returns the neighbors of the given process.

C SYNOPSIS

```
#include <mpi.h>
MPI_Graph_neighbors(MPI_Comm comm,int rank,int maxneighbors,int *neighbors);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GRAPH_NEIGHBORS(MPI_COMM COMM,INTEGER RANK,INTEGER MAXNEIGHBORS,
      INTEGER NNEIGHBORS(*),INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with graph topology (handle) (IN)
rank	is the rank of a process within group of **comm** (integer) (IN)
maxneighbors	is the size of array **neighbors** (integer) (IN)
neighbors	is the ranks of processes that are neighbors of the specified process (array of integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine retrieves the adjacency information for a particular task.

ERRORS

Invalid array size	**maxneighbors** < 0
Invalid rank	**rank** < 0 or **rank** > groupsize
MPI not initialized	
MPI already finalized	
Invalid communicator	
No topology	
Invalid topology type	no graph topology associate with communicator

RELATED INFORMATION

MPI_GRAPH_NEIGHBORS_COUNT
MPI_GRAPH_CREATE

NAME

MPI_Graph_neighbors_count, MPI_GRAPH_NEIGHBORS_COUNT- Returns the number of neighbors of the given process.

C SYNOPSIS

```
#include <mpi.h>
MPI_Graph_neighbors_count(MPI_Comm comm,int rank,int *neighbors);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GRAPH_NEIGHBORS_COUNT(INTEGER COMM,INTEGER RANK,INTEGER NEIGHBORS(*),
      INTEGER IERROR)
```

PARAMETERS

comm	is a communicator with graph topology (handle) (IN)
rank	is the rank of a process within **comm** (integer) (IN)
neighbors	is the number of neighbors of the specified process (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the number of neighbors of the given process.

ERRORS

Invalid rank	**rank** < 0 or **rank** > = groupsize
MPI not initialized	
MPI already finalized	
Invalid communicator	
No graph topology associated with communicator	
Invalid topology type	

RELATED INFORMATION

MPI_GRAPH_NEIGHBORS
MPE_GRAPH_CREATE

NAME

MPI_Graphdims_get, MPI_GRAHPDIMS_GET-Retrieves graph topology information from a communicator.

C SYNOPSIS

```
#include <mpi.h>
MPI_Graphdims_get(MPI_Comm comm,int *nnodes,int *nedges);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GRAPHDIMS_GET(INTEGER COMM,INTEGER NNDODES,INTEGER NEDGES,
      INTEGER IERROR)
```

PARAMETERS

comm is a communicator with graph topology (handle) (IN)

nnodes is an integer specifying the number of nodes in the graph. The number of nodes and the number of processes in the group are equal. (OUT)

nedges is an integer specifying the number of edges in the graph. (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine retrieves the number of nodes and the number of edges in the graph topology associated with a communicator.

ERRORS

MPI not initialized

MPI already finalized

Invalid communicator

No topology

Invalid topology type topology type must be graph

RELATED INFORMATION

MPI_GRAPH_GET
MPI_GRAPH_CREATE

NAME

MPI_Group_compare, MPI_GROUP_COMPARE-Compares the contents of two process groups.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_compare(MPI_Group group1,MPI_Group group2,
    int *result);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_COMPARE(INTEGER GROUP1,INTEGER GROUP2,INTEGER RESULT,
    INTEGER IERROR)
```

PARAMETERS

group1	is the first group (handle) (IN)
group2	is the second group (handle) (IN)
result	is the result (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine compares the contents of two process groups and returns one of the following:

MPI_IDENT	both groups have the exact group members and group order
MPI_SIMILAR	group members are the same but group order is different
MPI_UNEQUAL	group size and/or members are different

ERRORS

Invalid group(s)

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_COMM_COMPARE

NAME

MPI_Group_difference, MPI_GROUP_DIFFERENCE-Creates a new group that is the difference of two existing groups.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_difference(MPI_Group group1,MPI_Group group2,
    MPI_Group *newgroup);
```

FORTRAN SYNOPSIS

include 'mpif.h'

```
MPI_GROUP_DIFFERENCE(INTEGER GROUP1,INTEGER GROUP2,
      INTEGER NEWGROUP,INTEGER IERROR)
```

PARAMETERS

group1	is the first group (handle) (IN)
group2	is the second group (handle) (IN)
newgroup	is the difference group (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new group that is the difference of two existing groups. The new group consists of all elements of the first group (**group1**) that are not in the second group (**group2**), and is ordered as in the first group.

ERRORS

Invalid group(s)

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GROUP_UNION
MPI_GROUP_INTERSECTION

NAME

MPI_Group_excl, MPI_GROUP_EXCL-Creates a new group by excluding selected processes of an existing group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_excl(MPI_Group group,int n,int *ranks,
    MPI_Group *newgroup);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_EXCL(INTEGER GROUP,INTEGER N,INTEGER RANKS(*),
      INTEGER NEWGROUP,INTEGER IERROR)
```

PARAMETERS

group	is the group (handle) (IN)
n	is the number of elements in array **ranks** (integer) (IN)
ranks	is the array of integer ranks in **group** not to appear in **newgroup** (IN)
newgroup	is the new group derived from above preserving the order defined by **group** (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine removes selected processes from an existing group to create a new group.

MPI_GROUP_EXCL creates a group of processes **newgroup** obtained by deleting from **group** processes with ranks **ranks[0],... ranks[n-1]**. The ordering of processes in **newgroup** is identical to the ordering in **group**. Each of the **n** elements of **ranks** must be a valid rank in **group** and all elements must be distinct. If **n**= 0, then **newgroup** is identical to **group**.

ERRORS

Invalid group	
Invalid size	n <0 or **n** > groupsize
Invalid rank(s)	**ranks**[i] < 0 or **ranks**[i] > = groupsize
Duplicate rank(s)	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_GROUP_INCL
MPI_GROUP_RANGE_EXCL
MPI_GROUP_RANGE_INCL

NAME

MPI_Group_free, MPI_GROUP_FREE-Marks a group for deallocation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_free(MPI_Group *group);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_FREE(INTEGER GROUP,INTEGER IERROR)
```

PARAMETERS

group is the group (handle) (INOUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_GROUP_FREE sets the handle **group** to MPI_GROUP_NULL and marks the group object for deallocation. Actual deallocation occurs only after all operations involving **group** are completed. Any active operation using **group** completes normally but no new calls with meaningful references to the freed group are possible.

ERRORS

Invalid group

MPI not initialized

MPI already finalized

NAME

MPI_Group_incl, MPI_GROUP_INCL-Creates a new group consisting of selected processes from an existing group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_incl(MPI_Group group,int n,int *ranks,
    MPI_Group *newgroup);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_INCL(INTEGER GROUP,INTEGER N,INTEGER RANKS(*),
      INTEGER NEWGROUP,INTEGER IERROR)
```

PARAMETERS

group	is the group (handle) (IN)
n	is the number of elements in array **ranks** and the size of **newgroup**(integer) (IN)
ranks	is the ranks of processes in **group** to appear in **newgroup** (array of integers) (IN)
newgroup	is the new group derived from above in the order defined by **ranks** (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new group consisting of selected processes from an existing group.

MPI_GROUP_INCL creates a group **newgroup** consisting of **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. If **n** = 0, then **newgroup** is MPI_GROUP_EMPTY. This function can be used to reorder the elements of a group.

ERRORS

Invalid group	
Invalid size	n <0 or **n** > groupsize
Invalid rank(s)	**ranks**[i] < 0 or **ranks**[i] >= groupsize
Duplicate rank(s)	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_GROUP_EXCL
MPI_GROUP_RANGE_INCL
MPI_GROUP_RANGE_EXCL

NAME

MPI_Group_intersection, MPI_GROUP_INTERSECTION-Creates a new group that is the intersection of two existing groups.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_intersection(MPI_Group group1,MPI_Group group2,
    MPI_Group *newgroup);
```

FORTRAN SYNOPSIS

include 'mpif.h'

```
MPI_GROUP_INTERSECTION(INTEGER GROUP1,INTEGER GROUP2,
      INTEGER NEWGROUP,INTEGER IERROR)
```

PARAMETERS

group1	is the first group (handle) (IN)
group2	is the second group (handle) (IN)
newgroup	is the intersection group (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new group that is the intersection of two existing groups. The new group consists of all elements of the first group (**group1**) that are also part of the second group (**group2**), and is ordered as in the first group.

ERRORS

Invalid group(s)

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GROUP_UNION
MPI_GROUP_DIFFERENCE

NAME

MPI_Group_range_excl, MPI_GROUP_RANGE_EXCL-Creates a new group by removing selected ranges of processes from an existing group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_range_excl(MPI_Group group,int n,
    int ranges[][3],MPI_Group *newgroup);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_RANGE_EXCL(INTEGER GROUP,INTEGER N,INTEGER RANGES(3,*),
    INTEGER NEWGROUP,INTEGER IERROR)
```

PARAMETERS

group	is the group (handle) (IN)
n	is the number of triplets in array ranges (integer) (IN)
ranges	is an array of integer triplets of the form (first rank, last rank, stride) specifying the ranks in **group** of processes that are to be excluded from the output group **newgroup**. (IN)
newgroup	is the new group derived from above that preserves the order in **group** (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new group by removing selected ranges of processes from an existing group. Each computed rank must be a valid rank in **group** and all computed ranks must be distinct.

The function of this routine is equivalent to expanding the array **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**.

ERRORS

Invalid group	
Invalid size	**n** < 0 or **n** > groupsize
Invalid rank(s)	a computed rank < 0 or >= groupsize
Duplicate rank(s)	
Invalid stride(s)	**stride[i]** = **0**
Too many ranks	Number of ranks > groupsize

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GROUP_RANGE_INCL
MPI_GROUP_EXCL
MPI_GROUP_INCL

NAME

MPI_Group_range_incl, MPI_GROUP_RANGE_INCL-Creates a new group consisting of selected ranges of processes from an existing group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_range_incl(MPI_Group group,int n,
    int ranges[][3],MPI_Group *newgroup);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_RANGE_INCL(INTEGER GROUP,INTEGER N,INTEGER RANGES(3,*),
    INTEGER NEWGROUP,INTEGER IERROR)
```

PARAMETERS

group	is the group (handle) (IN)
n	is the number of triplets in array **ranges** (integer) (IN)
ranges	is an array of integer triplets of the form (first rank, last rank, stride) indicating ranks in **group** of processes to be included in **newgroup** (IN)
newgroup	is the new group derived from above in the order defined by **ranges** (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new group consisting of selected ranges of processes from an existing group. The function of this routine is 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**.

ERRORS

Invalid group	
Invalid size	**n** <0 or **n** > groupsize
Invalid rank(s)	a computed rank < 0 or >= groupsize
Duplicate rank(s)	
Invalid stride(s)	**stride[i]** = 0
Too many ranks	**nranks** > groupsize
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_GROUP_RANGE_EXCL
MPI_GROUP_INCL
MPI_GROUP_EXCL

NAME

MPI_Group_rank, MPI_GROUP_RANK-Returns the rank of the local process with respect to **group**.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_rank(MPI_Group group,int *rank);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_RANK(INTEGER GROUP,INTEGER RANK,INTEGER IERROR)
```

PARAMETERS

group	is the group (handle) (IN)
rank	is an integer that specifies the rank of the calling process in group or MPI_UNDEFINED if the process is not a member. (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the rank of the local process with respect to **group**. This local operation does not require any interprocess communication.

ERRORS

Invalid group

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_COMM_RANKS

NAME

MPI_Group_size, MPI_GROUP_SIZE-Returns the number of processes in a group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_size(MPI_Group group,int *size);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_SIZE(INTEGER GROUP,INTEGER SIZE,INTEGER IERROR)
```

PARAMETERS

group	is the group (handle) (IN)
size	is the number of processes in the group (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the number of processes in a group. This is a local operation and does not require any interprocess communication.

ERRORS

Invalid group

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_COMM_SIZE

NAME

MPI_Group_translate_ranks, MPI_GROUP_TRANSLATE_RANKS-Converts process ranks of one group into ranks of another group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_translate_ranks(MPI_Group group1,int n,
    int *ranks1,MPI_Group group2,int *ranks2);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_GROUP_TRANSLATE_RANKS(INTEGER GROUP1, INTEGER N,
    INTEGER RANKS1(*),INTEGER GROUP2,INTEGER RANKS2(*),INTEGER IERROR)
```

PARAMETERS

group1	is group1 (handle) (IN)
n	is an integer that specifies the number of ranks in **ranks1** and **ranks2** arrays (IN)
ranks1	is an array of zero or more valid ranks in **group1** (IN)
group2	is group2 (handle) (IN)
ranks2	is an array of corresponding ranks in **group2**. If the process of **ranks1(i)** is not a member of **group2**, **ranks2(i)** returns MPI_UNDEFINED. (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This subroutine converts process ranks of one group into ranks of another group. For example, if you know the ranks of processes in one group, you can use this function to find the ranks of processes in another group.

ERRORS

Invalid group(s)	
Invalid rank count	**n** < 0
Invalid rank	**ranks1[i]** < 0 or **ranks1[i]** > = size of **group1**
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_COMM_COMPARE

NAME

MPI_Group_union, MPI_GROUP_UNION-Creates a new group that is the union of two existing groups.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Group_union(MPI_Group group1,MPI_Group group2,
    MPI_Group *newgroup);
```

FORTRAN SYNOPSIS

include 'mpif.h'

```
MPI_GROUP_UNION(INTEGER GROUP1,INTEGER GROUP2,INTEGER NEWGROUP,
      INTEGER IERROR)
```

PARAMETERS

group1	is the first group (handle) (IN)
group2	is the second group (handle) (IN)
newgroup	is the union group (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a new group that is the union of two existing groups. The new group consists of the elements of the first group (**group1**) followed by all the elements of the second group (**group2**) not in the first group.

ERRORS

Invalid group(s)

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_GROUP_INTERSECTION
MPI_GROUP_DIFFERENCE

NAME

MPI_Ibsend, MPI_IBSEND-Nonblocking buffered mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Ibsend(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_IBSEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
    INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of the destination task in **comm** (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_IBSEND starts a buffered mode, nonblocking send. The send buffer may not be modified until the request has been completed by MPI_WAIT, MPI_TEST, or one of the other MPI wait or test functions.

NOTES

See MPI_BSEND for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Invalid destination	
Type not committed	
	dest < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	

MPI not initialized

MPI already finalized

Develop mode error if:

Illegal buffer update

RELATED INFORMATION

MPI_BSEND
MPI_BSEND_INIT
MPI_WAIT
MPI_BUFFER_ATTACH

NAME

MPI_Init, MPI_INIT-Initializes MPI.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Init(int *argc,char ***argv);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_INIT(INTEGER IERROR)
```

PARAMETERS

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine initializes MPI. All MPI programs must call this routine before any other MPI routine (with the exception of MPI_INITIALIZED). More than one call to MPI_INIT by any task is erroneous.

NOTES

argc and **argv** are the arguments passed to **main**. The IBM MPI implementation of the MPI Standard does not examine or modify these arguments when passed to MPI_INIT.

In a threaded environment, MPI_INIT needs to be called once per process and not once per thread. You don't need to call it on the main thread but both MPI_INIT and MPI_FINALIZE must be called on the same thread.

MPI_INIT opens a local socket and binds it to a port, sends that information to POE, receives a list of destination addresses and ports, opens a socket to send to each one, verifies that communication can be established, and distributes MPI internal state to each task.

In the signal-handling library, this work is done in the initialization stub added by POE, so that the library is open when your main program is called. MPI_INIT sets a flag saying that you called it.

In the threaded library, the work of MPI_INIT is done when the function is called. It is not open when your main program starts. This may affect the numbering of file descriptors, the use of the environment strings, and the treatment of stdin (the MP_HOLD_STDIN variable). If an existing non-threaded program is recompiled using the threaded library, the code prior to calling MPI_INIT should be examined with these thoughts in mind.

Also for the threaded library, if you had registered a function as an AIX signal handler for the SIGIO signal at the time that MPI_INIT was called, that function will be added to the interrupt service thread and be processed as a thread function rather than as a signal handler. You'll need to set the environment variable

MP_CSS_INTERRUPT=YES to get arriving packets to invoke the interrupt service thread.

ERRORS

MPI already initialized

MPI already finalized

RELATED INFORMATION

MPI_INITIALIZED
MPI_FINALIZE

NAME

MPI_Initialized, MPI_INITIALIZED-Determines if MPI is initialized.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Initialized(int *flag);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_INITIALIZED(INTEGER FLAG,INTEGER IERROR)
```

PARAMETERS

flag is true if MPI_INIT was called; otherwise is false.

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine determines if MPI is initialized. This and MPI_GET_VERSION are the only MPI calls that can be made before MPI_INIT is called.

NOTES

Because it is erroneous to call MPI_INIT more than once per process, use MPI_INITIALIZED if there is doubt as to the state of MPI.

RELATED INFORMATION

MPI_INIT

NAME

MPI_Intercomm_create, MPI_INTERCOMM_CREATE-Creates an intercommunicator from two intracommunicators.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Intercomm_create(MPI_Comm local_comm,int local_leader,
    MPI_Comm peer_comm,int remote_leader,int tag,MPI_Comm *newintercom);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_INTERCOMM_CREATE(INTEGER LOCAL_COMM,INTEGER LOCAL_LEADER,
    INTEGER PEER_COMM,INTEGER REMOTE_LEADER,INTEGER TAG,
    INTEGER NEWINTERCOM,INTEGER IERROR)
```

PARAMETERS

local_comm	is the local intracommunicator (handle) (IN)
local_leader	is an integer specifying the rank of local group leader in **local_comm** (IN)
peer_comm	is the "peer" intracommunicator (significant only at the **local_leader**) (handle) (IN)
remote_leader	is the rank of remote group leader in **peer_comm** (significant only at the **local_leader**) (integer) (IN)
tag	"safe" tag (integer) (IN)
newintercom	is the new intercommunicator (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates an intercommunicator from two intracommunicators and is collective over the union of the local and the remote groups. Processes should provide identical **local_comm** and **local_leader** arguments within each group. Wildcards are not permitted for **remote_leader**, **local_leader**, and **tag**.

MPI_INTERCOMM_CREATE uses point-to-point communication with communicator **peer_comm** and tag **tag** between the leaders. Make sure that there are no pending communications on **peer_comm** that could interfere with this communication. It is recommended that you use a dedicated peer communicator, such as a duplicate of MPI_COMM_WORLD, to avoid trouble with peer communicators.

ERRORS

Invalid communicator(s)	
Invalid communicator type(s)	must be intracommunicator(s)
Invalid rank(s)	**rank** < 0 or **rank** > = groupsize
Invalid tag	**tag** < 0
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_COMM_DUP
MPI_INTERCOMM_MERGE

NAME

MPI_Intercomm_merge, MPI_INTERCOMM_MERGE-Creates an intracommunicator by merging the local and the remote groups of an intercommunicator.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Intercomm_merge(MPI_Comm intercomm,int high,
    MPI_Comm *newintracom);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_INTERCOMM_MERGE(INTEGER INTERCOMM,INTEGER HIGH,
      INTEGER NEWINTRACOMM,INTEGER IERROR)
```

PARAMETERS

intercomm	is the intercommunicator (handle) (IN)
high	(logical) (IN)
newintracomm	is the new intracommunicator (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates an intracommunicator from the union of two groups associated with **intercomm**. Processes should provide the same **high** value within each of the two groups. If processes in one group provide the value **high** = **false** and processes in the other group provide 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.

ERRORS

Invalid communicator

Invalid communicator type must be intercommunicator

Inconsistent high within group

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_INTERCOMM_CREATE

NAME

MPI_Iprobe, MPI_IPROBE-Checks if a message matching **source**, **tag**, and **comm** has arrived.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Iprobe(int source,int tag,MPI_Comm comm,int *flag,
      MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_IPROBE(INTEGER SOURCE,INTEGER TAG,INTEGER COMM,INTEGER FLAG,
      INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

source	is a source rank or MPI_ANY_SOURCE (integer) (IN)
tag	is a tag value or MPI_ANY_TAG (integer) (IN)
comm	is a communicator (handle) (IN)
flag	(logical) (OUT)
status	is a status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine allows you to check for incoming messages without actually receiving them.

MPI_IPROBE(**source, tag, comm, flag, status**) returns **flag** = **true** when there is a message that can be received 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 values that would have been returned by MPI_RECV(). Otherwise, the call returns **flag** = **false** and leaves **status** undefined.

When MPI_IPROBE returns **flag** = **true**, the content of the status object can be accessed to find the source, tag and length of the probed message.

A subsequent receive executed with the same **comm**, and the source and tag returned in **status** by MPI_IPROBE receives the message that was matched by the probe, if no other intervening receive occurs after the initial probe.

source can be MPI_ANY_SOURCE and **tag** can be MPI_ANY_TAG. This allows you to probe messages from any source and/or with any tag, but you must provide a specific communicator with **comm**.

When a message is not received immediately after it is probed, the same message can be probed for several times before it is received.

NOTES

In a threaded environment, MPI_PROBE or MPI_IPROBE followed by MPI_RECV, based on the information from the probe, may not be a thread-safe operation. You must ensure that no other thread received the detected message.

An MPI_IPROBE cannot prevent a message from being cancelled successfully by the sender, making it unavailable for the MPI_RECV. Structure your program so this will not occur.

ERRORS

Invalid source	**source** < 0 or **source** > = groupsize
Invalid tag	**tag** < 0
Invalid communicator	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_PROBE
MPI_RECV

NAME

MPI_Irecv, MPI_IRECV-Nonblocking receive.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Irecv(void* buf,int count,MPI_Datatype datatype,
    int source,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_IRECV(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER SOURCE,
      INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the receive buffer (choice) (OUT)
count	is the number of elements in the receive buffer (integer) (IN)
datatype	is the datatype of each receive buffer element (handle) (IN)
source	is the rank of source or MPI_ANY_SOURCE (integer) (IN)
tag	is the message tag or MPI_ANY_TAG (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine starts a nonblocking receive and returns a handle to a request object. You can later use the **request** to query the status of the communication or wait for it to complete.

A nonblocking receive call means the system may start writing data into the receive buffer. Once the nonblocking receive operation is called, do not access any part of the receive buffer until the receive is complete.

NOTES

The message received must be less than or equal to the length of the receive buffer. If all incoming messages do not fit without truncation, an overflow error occurs. If a message arrives that is shorter than the receive buffer, then only those locations corresponding to the actual message are changed. If an overflow occurs, it is flagged at the MPI_WAIT or MPI_TEST. See MPI_RECV for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid source	**source** < 0 or **source** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_RECV
MPI_RECV_INIT
MPI_WAIT

NAME

MPI_Irsend, MPI_IRSEND-Nonblocking ready mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Irsend(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_IRSEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
      INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of the destination task in **comm** (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_IRSEND starts a ready mode, nonblocking send. The send buffer may not be modified until the request has been completed by MPI_WAIT, MPI_TEST, or one of the other MPI wait or test functions.

NOTES

See MPI_RSEND for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
No receive posted	error flagged at destination

MPI not initialized

MPI already finalized

Develop mode error if:

Illegal buffer update

RELATED INFORMATION

MPI_RSEND
MPI_RSEND_INIT
MPI_WAIT

NAME

MPI_Isend, MPI_ISEND-Nonblocking standard mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Isend(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm,MPI_request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ISEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
    INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of the destination task in **comm** (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine starts a nonblocking standard mode send. The send buffer may not be modified until the request has been completed by MPI_WAIT, MPI_TEST, or one of the other MPI wait or test functions.

NOTES

See MPI_SEND for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	

MPI already finalized

Develop mode error if:

Illegal buffer update

RELATED INFORMATION

MPI_SEND
MPI_SEND_INIT
MPI_WAIT

NAME

MPI_Issend, MPI_ISSEND-Nonblocking synchronous mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Issend(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_ISSEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
    INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of the destination task in **comm** (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_ISSEND starts a synchronous mode, nonblocking send. The send buffer may not be modified until the request has been completed by MPI_WAIT, MPI_TEST, or one of the other MPI wait or test functions.

NOTES

See MPI_SSEND for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	

MPI already finalized

Develop mode error if:

Illegal buffer update

RELATED INFORMATION

MPI_SSEND
MPI_SSEND_INIT
MPI_WAIT

NAME

MPI_Keyval_create, MPI_KEYVAL_CREATE-Generates a new attribute key.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Keyval_create(MPI_Copy_function *copy_fn,
    MPI_Delete_function *delete_fn,int *keyval,
    void* extra_state);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_KEYVAL_CREATE(EXTERNAL COPY_FN,EXTERNAL DELETE_FN,
      INTEGER KEYVAL,INTEGER EXTRA_STATE,INTEGER IERROR)
```

PARAMETERS

copy_fn	is the copy callback function for keyval (IN)
delete_fn	is the delete callback function for keyval (IN)
keyval	is an integer specifying the key value for future access (OUT)
extra_state	is the extra state for callback functions (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine generates a new attribute key. Keys are locally unique in a process, opaque to the user, and are explicitly stored in integers. Once allocated, **keyval** can be used to associate attributes and access them on any locally defined communicator. **copy_fn** is invoked when a communicator is duplicated by MPI_COMM_DUP. It should be of type MPI_COPY_FUNCTION, which is defined as follows:

In C:

```
typedef int MPI_Copy_function (MPI_Comm oldcomm,int keyval,
        void *extra_state,void *attribute_val_in,
        void *attribute_val_out,int *flag);
```

In FORTRAN:

```
PROCEDURE COPY_FUNCTION(INTEGER OLDCOMM,INTEGER KEYVAL,
        INTEGER EXTRA_STATE,INTEGER ATTRIBUTE_VAL_IN,
        INTEGER ATTRIBUTE_VAL_OUT,LOGICAL FLAG,INTEGER IERROR)
```

You can use the predefined functions MPI_NULL_COPY_FN and MPI_DUP_FN to never copy or to always copy, respectively.

delete_fn is invoked when a communicator is deleted by MPI_COMM_FREE or when a call is made to MPI_ATTR_DELETE. A call to MPI_ATTR_PUT that over-

lays a previously put attribute also causes **delete_fn** to be called. It should be defined as follows:

In C:

```
typedef int MPI_Delete_function (MPI_Comm comm,int keyval,
        void *attribute_val, void *extra_state);
```

In FORTRAN:

```
PROCEDURE DELETE_FUNCTION(INTEGER COMM,INTEGER KEYVAL,
        INTEGER ATTRIBUTE_VAL,INTEGER EXTRA_STATE,
        INTEGER IERROR)
```

You can use the predefined function MPI_NULL_DELETE_FN if no special handling of attribute deletions is required.

In FORTRAN, the value of **extra_state** is recorded by MPI_KEYVAL_CREATE and the callback functions should not attempt to modify this value.

The MPI standard requires that when **copy_fn** or **delete_fn** gives a return code other than MPI_SUCCESS, the MPI routine in which this occurs must fail. The standard does not suggest that the **copy_fn** or **delete_fn** return code be used as the MPI routine's return value. The standard does require that an MPI return code be in the range between MPI_SUCCESS and MPI_ERR_LASTCODE. It places no range limits on **copy_fn** or **delete_fn** return codes. For this reason, we provide a specific error code for a **copy_fn** failure and another for a **delete_fn** failure. These error codes can be found in error class MPI_ERR_OTHER. The **copy_fn** or the **delete_fn** return code is not preserved.

ERRORS

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ATTR_PUT
MPI_ATTR_DELETE
MPI_COMM_DUP
MPI_COMM_FREE

NAME

MPI_Keyval_free, MPI_KEYVAL_FREE-Marks an attribute key for deallocation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Keyval_free(int *keyval);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_KEYVAL_FREE(INTEGER KEYVAL,INTEGER IERROR)
```

PARAMETERS

keyval attribute key (integer) (INOUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine sets **keyval** to MPI_KEYVAL_INVALID and marks the attribute key for deallocation. You can free an attribute key that is in use because the actual deallocation occurs only when all active references to it are complete. These references, however, need to be explicitly freed. Use calls to MPI_ATTR_DELETE to free one attribute instance. To free all attribute instances associated with a communicator, use MPI_COMM_FREE.

ERRORS

Invalid attribute key attribute key is undefined

Predefined attribute key attribute key is predefined

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_ATTR_DELETE

MPI_COMM_FREE

NAME

MPI_Op_create, MPI_OP_CREATE-Binds a user-defined reduction operation to an **op** handle.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Op_create(MPI_User_function *function,int commute,
      MPI_Op *op);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_OP_CREATE(EXTERNAL FUNCTION,INTEGER COMMUTE,INTEGER OP,
      INTEGER IERROR)
```

PARAMETERS

function	is the user-defined reduction function (function) (IN)
commute	is **true** if commutative; otherwise it's **false** (IN)
op	is the reduction operation (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine binds a user-defined reduction operation to an **op** handle which you can then use in MPI_REDUCE, MPI_ALLREDUCE, MPI_REDUCE_SCATTER and MPI_SCAN and their nonblocking equivalents.

The user-defined operation is assumed to be associative. If **commute** = **true**, then the operation must be both commutative and associative. If **commute** = **false**, then the order of the operation is fixed. The order is defined in ascending, process rank order and begins with process zero.

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

The following is the ANSI-C prototype for the function:

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

The following is the FORTRAN declaration for the function:

```
FUNCTION USER_FUNCTION(INVEC(*), INOUTVEC(*), LEN, TYPE)
<type> INVEC(LEN), INOUTVEC(LEN)
 INTEGER LEN, TYPE
```

NOTES

See Appendix D, “Reduction Operations” on page 277 for information about reduction functions.

ERRORS

Null function

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_OP_FREE
MPI_REDUCE
MPI_ALLREDUCE
MPI_REDUCE_SCATTER
MPI_SCAN

NAME

MPI_Op_free, MPI_OP_FREE-Marks a user-defined reduction operation for deallocation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Op_free(MPI_Op *op);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_OP_FREE(INTEGER OP,INTEGER IERROR)
```

PARAMETERS

op is the reduction operation (handle) (IN)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This function marks a reduction operation for deallocation, and set **op** to MPI_OP_NULL. Actual deallocation occurs when the operation's reference count is zero.

ERRORS

Invalid operation

Predefined operation

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_OP_CREATE

NAME

MPI_Pack, MPI_PACK-Packs the message in the specified send buffer into the specified buffer space.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Pack(void* inbuf,int incount,MPI_Datatype datatype,
    void *outbuf,int outsize,int *position,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_PACK(CHOICE INBUF,INTEGER INCOUNT,INTEGER DATATYPE,
    CHOICE OUTBUF,INTEGER OUTSIZE,INTEGER POSITION,INTEGER COMM
    INTEGER IERROR)
```

PARAMETERS

inbuf	is the input buffer start (choice) (IN)
incount	is an integer specifying the number of input data items (IN)
datatype	is the datatype of each input data item (handle) (IN)
outbuf	is the output buffer start (choice) (OUT)
outsize	is an integer specifying the output buffer size in bytes (OUT)
position	is the current position in the output buffer counted in bytes (integer) (INOUT)
comm	is the communicator for sending the packed message (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine packs the message specified by **inbuf**, **incount**, and **datatype** into the buffer space specified by **outbuf** and **outsize**. The input buffer is any communication buffer allowed in MPI_SEND. The output buffer is any contiguous storage space containing **outsize** bytes and starting at the address **outbuf**.

The input value of **position** is the beginning offset in the output buffer that will be used for packing. The output value of **position** is the offset in the output buffer following the locations occupied by the packed message. **comm** is the communicator that will be used for sending the packed message.

ERRORS

Invalid incount	**incount** < 0
Invalid datatype	
Type not committed	
Invalid communicator	
Outbuf too small	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_UNPACK
MPI_PACK_SIZE

NAME

MPI_Pack_size, MPI_PACK_SIZE-Returns the number of bytes required to hold the data.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Pack_size(int incount,MPI_Datatype datatype,
    MPI_Comm comm, int *size);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_PACK_SIZE(INTEGER INCOUNT,INTEGER DATATYPE,INTEGER COMM,INTEGER SIZE,
      INTEGER IERROR)
```

PARAMETERS

incount	is an integer specifying the count argument to a packing call (IN)
datatype	is the datatype argument to a packing call (handle) (IN)
comm	is the communicator to a packing call (handle) (IN)
size	size of packed message in bytes (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the number of bytes required to pack **incount** replications of the datatype. You can use MPI_PACK_SIZE to determine the size required for a packing buffer or to track space needed for buffered sends.

ERRORS

Invalid datatype

Type is not committed

MPI not initialized

MPI already finalized

Invalid communicator

Invalid incount **incount** < 0

RELATED INFORMATION

MPI_PACK

NAME

MPI_Pcontrol, MPI_PCONTROL-Provides profiler control.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Pcontrol(const int level, ...);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_PCONTROL(INTEGER LEVEL, ...)
```

PARAMETERS

level	is the profiling level (IN) The proper values for **level** and the meanings of those values are determined by the profiler being used.
...	0 or more parameters
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_PCONTROL is a placeholder to allow applications to run with or without an independent profiling package without modification. MPI implementations do not use this routine and do not have any control of the implementation of the profiling code.

Calls to this routine allow a profiling package to be controlled from MPI programs. The nature of control and the arguments required are determined by the profiling package. The MPI library routine by this name returns to the caller without any action.

NOTES

For each additional call level introduced by the profiling code, the global variable VT_instaddr_depth needs to be incremented so the Visualization Tool Tracing Subsystem(VT) can record where the application called the MPI message passing library routine. The VT_instaddr_depth variable is defined in /usr/lpp/ppe.vt/include/VT_mpi.h.

ERRORS

MPI does not report any errors for MPI_PCONTROL.

NAME

MPI_Probe, MPI_PROBE-Waits until a message matching **source**, **tag**, and **comm** arrives.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Probe(int source,int tag,MPI_Comm comm,MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_PROBE(INTEGER SOURCE,INTEGER TAG,INTEGER COMM,
    INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

source	is a source rank or MPI_ANY_SOURCE (integer) (IN)
tag	is a source tag or MPI_ANY_TAG (integer) (IN)
comm	is a communicator (handle) (IN)
status	is a status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_PROBE behaves like MPI_IPROBE. It allows you to check for an incoming message without actually receiving it. MPI_PROBE is different in that it is a blocking call that returns only after a matching message has been found.

NOTES

In a threaded environment, MPI_PROBE or MPI_IPROBE followed by MPI_RECV, based on the information from the probe, may not be a thread-safe operation. You must ensure that no other thread received the detected message.

An MPI_IPROBE cannot prevent a message from being cancelled successfully by the sender, making it unavailable for the MPI_RECV. Structure your program so this will not occur.

ERRORS

Invalid source	**source** < 0 or **source** > = groupsize
Invalid tag	**tag** < 0
Invalid communicator	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_IPROBE
MPI_RECV

NAME

MPI_Recv, MPI_RECV-Blocking receive.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Recv(void* buf,int count,MPI_Datatype datatype,
    int source,int tag,MPI_Comm comm,MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_RECV(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER SOURCE,
      INTEGER TAG,INTEGER COMM,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the receive buffer (choice) (OUT)
count	is the number of elements to be received (integer) (IN)
datatype	is the datatype of each receive buffer element (handle) (IN)
source	is the rank of the source task in **comm** or MPI_ANY_SOURCE (integer) (IN)
tag	is the message tag or MPI_ANY_TAG (integer) (IN)
comm	is the communicator (handle) (IN)
status	is the status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_RECV is a blocking receive. The receive buffer is storage containing room for **count** consecutive elements of the type specified by **datatype**, starting at address **buf**.

The message received must be less than or equal to the length of the receive buffer. If all incoming messages do not fit without truncation, an overflow error occurs. If a message arrives that is shorter than the receive buffer, then only those locations corresponding to the actual message are changed.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid source	**source** < 0 or **source** > = groupsize

Invalid tag	**tag** < 0
Invalid comm	
Truncation occurred	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_IRECV
MPI_SENDRECV
MPI_SEND

NAME

MPI_Recv_init, MPI_RECV_INIT-Creates a persistent receive request.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Recv_init(void* buf,int count,MPI_Datatype datatype,
    int source,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_RECV_INIT(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,
      INTEGER SOURCE,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the receive buffer (choice) (OUT)
count	is the number of elements to be received (integer) (IN)
datatype	is the type of each element (handle) (IN)
source	is the rank of source or MPI_ANY_SOURCE (integer) (IN)
tag	is the tag or MPI_ANY_TAG (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a persistent communication request for a receive operation. The argument **buf** is marked as OUT because the user gives permission to write to the receive buffer by passing the argument to MPI_RECV_INIT.

A persistent communication request is inactive after it is created. No active communication is attached to the request.

A send or receive communication using a persistent request is initiated by the function MPI_START.

NOTES

See MPI_RECV for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid source	**source** < 0 or **source** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_START
MPI_IRECV

NAME

MPI_Reduce, MPI_REDUCE-Applies a reduction operation to the vector **sendbuf** over the set of processes specified by **comm** and places the result in **recvbuf** on **root**.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Reduce(void* sendbuf,void* recvbuf,int count,
    MPI_Datatype datatype,MPI_Op op,int root,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_REDUCE(CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,
    INTEGER DATATYPE,INTEGER OP,INTEGER ROOT,INTEGER COMM,
    INTEGER IERROR)
```

PARAMETERS

sendbuf	is the address of the send buffer (choice) (IN)
recvbuf	is the address of the receive buffer (choice, significant only at root) (OUT)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of elements of the send buffer (handle) (IN)
op	is the reduction operation (handle) (IN)
root	is the rank of the root process (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine applies a reduction operation to the vector **sendbuf** over the set of processes specified by **comm** and places the result in **recvbuf** on **root**.

Both the input and output buffers have the same number of elements with the same type. The arguments **sendbuf**, **count**, and **datatype** define the send or input buffer and **recvbuf**, **count** and **datatype** define the output buffer. MPI_REDUCE is called by all group members using the same arguments for **count**, **datatype**, **op**, and **root**. If a sequence of elements is provided to a process, then the reduce operation is executed element-wise on each entry of the sequence. Here's an 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**)).

Users may define their own operations or use the predefined operations provided by MPI. User defined operations can be overloaded to operate on several

datatypes, either basic or derived. A list of the MPI predefined operations is in this manual. Refer to Appendix D, “Reduction Operations” on page 277.

The argument **datatype** of MPI_REDUCE must be compatible with **op**. For a list of predefined operations refer to Appendix I, “Predefined Datatypes” on page 349.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

NOTES

See Appendix D, “Reduction Operations” on page 277.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid op	
Invalid root	**root** < 0 or **root** > = groupsize
Invalid communicator	
Invalid communicator type	must be intracommunicator
Unequal message lengths	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent op

Inconsistent datatype

Inconsistent root

Inconsistent message length

RELATED INFORMATION

MPE_IREDUCE
MPI_ALLREDUCE
MPI_REDUCE_SCATTER
MPI_SCAN
MPI_OP_CREATE

NAME

MPI_Reduce_scatter, MPI_REDUCE_SCATTER-Applies a reduction operation to the vector **sendbuf** over the set of processes specified by **comm** and scatters the result according to the values in **recvcounts**.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Reduce_scatter(void* sendbuf,void* recvbuf,int *recvcounts,
    MPI_Datatype datatype,MPI_Op op,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_REDUCE_SCATTER(CHOICE SENDBUF,CHOICE RECVBUF,
    INTEGER RECVCOUNTS(*),INTEGER DATATYPE,INTEGER OP,
    INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
recvbuf	is the starting address of the receive buffer (choice) (OUT)
recvcounts	integer array specifying the number of elements in result distributed to each process. Must be identical on all calling processes. (IN)
datatype	is the datatype of elements in the input buffer (handle) (IN)
op	is the reduction operation (handle) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_REDUCE_SCATTER first performs an element-wise reduction on vector of **count** = Σ **i recvcounts[i]** elements in the send buffer defined by **sendbuf**, **count** and **datatype**. Next, the resulting vector is split into **n** disjoint segments, where **n** is the number of members in the group. Segment **i** contains **recvcounts[i]** elements. The **i**th segment is sent to process **i** and stored in the receive buffer defined by **recvbuf**, **recvcounts[i]** and **datatype**.

NOTES

MPI_REDUCE_SCATTER is functionally equivalent to MPI_REDUCE with **count** equal to the sum of **recvcounts[i]** followed by MPI_SCATTERV with **sendcounts** equal to **recvcounts**. When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid recvcounts	**recvcounts[i]** < 0
Invalid datatype	
Type not committed	
Invalid op	
Invalid communicator	
Invalid communicator type	must be intracommunicator
Unequal message lengths	
MPI not initialized	
MPI already finalized	

Develop mode error if:

Inconsistent op

Inconsistent datatype

RELATED INFORMATION

MPE_IREDUCE_SCATTER
MPI_REDUCE
MPI_OP_CREATE

NAME

MPI_Request_free, MPI_REQUEST_FREE-Marks a request for deallocation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Request_free(int MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_REQUEST_FREE(INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

request is a communication request (handle) (INOUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine marks a request object for deallocation and sets **request** to MPI_REQUEST_NULL. An ongoing communication associated with the request is allowed to complete before deallocation occurs.

NOTES

This function marks a communication request as **free**. Actual deallocation occurs when the **request** is complete. Active receive requests and collective communication requests cannot be freed.

ERRORS

Invalid request

Attempt to free receive request

Attempt to free CCL request

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_WAIT

NAME

MPI_Rsend, MPI_RSEND-Blocking ready mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Rsend(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_RSEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
      INTEGER TAG,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of destination (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a blocking ready mode send. It can be started only when a matching receive is posted. If a matching receive is not posted, the operation is erroneous and its outcome is undefined.

The completion of MPI_RSEND indicates that the send buffer can be reused.

NOTES

A ready send for which no receive exists produces an asynchronous error at the destination. The error is not detected at the MPI_RSEND and it returns MPI_SUCCESS.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	

No receive posted	error flagged at destination
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_IRSEND
MPI_SEND

NAME

MPI_Rsend_init, MPI_RSEND_INIT-Creates a persistent ready mode send request.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Rsend_init(void* buf,int count,MPI_Datatype datatype,
      int dest,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_RSEND_INIT(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,
      INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements to be sent (integer) (IN)
datatype	is the type of each element (handle) (IN)
dest	is the rank of the destination task (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_RSEND_INIT creates a persistent communication object for a ready mode send operation. MPI_START or MPI_STARTALL is used to activate the send.

NOTES

See MPI_RSEND for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	

MPI already finalized

RELATED INFORMATION

MPI_START
MPI_IRSEND

NAME

MPI_Scan, MPI_SCAN -Performs a parallel prefix reduction on data distributed across a group.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Scan(void* sendbuf,void* recvbuf,int count,
    MPI_Datatype datatype,MPI_Op op,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SCAN(CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,
    INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the starting address of the send buffer (choice) (IN)
recvbuf	is the starting address of the receive buffer (choice) (OUT)
count	is the number of elements in **sendbuf** (integer) (IN)
datatype	is the datatype of elements in **sendbuf** (handle) (IN)
op	is the reduction operation (handle) (IN)
comm	is the communicator (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

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, ..., **i** (inclusive). The type of operations supported, their semantics, and the restrictions on send and receive buffers are the same as for MPI_REDUCE.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid op	
Invalid communicator	

Invalid communicator type must be intracommunicator

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent op

Inconsistent datatype

Inconsistent message length

RELATED INFORMATION

MPE_ISCAN
MPI_REDUCE
MPI_OP_CREATE

NAME

MPI_Scatter, MPI_SCATTER-Distributes individual messages from **root** to each process in **comm**.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Scatter(void* sendbuf,int sendcount,MPI_Datatype sendtype,
    void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,
    MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SCATTER(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
    CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,
    INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the address of the send buffer (choice, significant only at **root**) (IN)
sendcount	is the number of elements to be sent to each process (integer, significant only at **root**) (IN)
sendtype	is the datatype of the send buffer elements (handle, significant only at **root**) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcount	is the number of elements in the receive buffer (integer) (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
root	is the rank of the sending process (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_SCATTER distributes individual messages from **root** to each process in **comm**. This routine is the inverse operation to MPI_GATHER.

The type signature associated with **sendcount**, **sendtype** at the root must be equal to the type signature associated with **recvcount**, **recvtype** at all processes. (Type maps can be different.) This means 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 allowed.

The following is information regarding MPI_SCATTER arguments and processes:

- On the process **root**, all arguments to the function are significant.

- On other processes, only the arguments **recvbuf, recvcount, recvtype, root,** and **comm** are significant.
- The argument **root** must be the same on all processes.

A call where the specification of counts and types causes any location on the root to be read more than once is erroneous.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, "Programming Considerations for User Applications in POE" on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid root (**root** < 0 or **root** >= groupsize)

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

Inconsistent message length

RELATED INFORMATION

MPE_ISCATTER
MPI_SCATTER
MPI_GATHER

NAME

MPI_Scatterv, MPI_SCATTERV-Distributes individual messages from **root** to each process in **comm**. Messages can have different sizes and displacements.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Scatterv(void* sendbuf,int *sendcounts,
    int *displs,MPI_Datatype sendtype,void* recvbuf,
    int recvcount,MPI_Datatype recvtype,int root,
    MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SCATTERV(CHOICE SENDBUF,INTEGER SENDCOUNTS(*),INTEGER DISPLS(*),
    INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,
    INTEGER ROOT,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

sendbuf	is the address of the send buffer (choice, significant only at **root**) (IN)
sendcounts	integer array (of length group size) that contains the number of elements to send to each process (significant only at **root**) (IN)
displs	integer array (of length group size). Entry **i** specifies the displacement relative to **sendbuf** from which to send the outgoing data to process **i** (significant only at **root**) (IN)
sendtype	is the datatype of the send buffer elements (handle, significant only at **root**) (IN)
recvbuf	is the address of the receive buffer (choice) (OUT)
recvcount	is the number of elements in the receive buffer (integer) (IN)
recvtype	is the datatype of the receive buffer elements (handle) (IN)
root	is the rank of the sending process (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine distributes individual messages from **root** to each process in **comm**. Messages can have different sizes and displacements.

With **sendcounts** as an array, messages can have varying sizes of data that can be sent to each process. **displs** allows you the flexibility of where the data can be taken from on the **root**.

The type signature of **sendcount[i]**, **sendtype** at the **root** must be equal to the type signature of **recvcount**, **recvtype** at process **i**. (The type maps can be different.) This means 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 allowed.

The following is information regarding MPI_SCATTERV arguments and processes:

- On the process **root**, all arguments to the function are significant.
- On other processes, only the arguments **recvbuf, recvcount, recvtype, root,** and **comm** are significant.
- The argument **root** must be the same on all processes.

A call where the specification of sizes, types and displacements causes any location on the root to be read more than once is erroneous.

When you use this routine in a threaded application, make sure all collective operations on a particular communicator occur in the same order at each task. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid communicator

Invalid communicator type must be intracommunicator

Invalid count(s) **count** < 0

Invalid datatype(s)

Type not committed

Invalid root **root** < 0 or **root** >= groupsize

Unequal message lengths

MPI not initialized

MPI already finalized

Develop mode error if:

Inconsistent root

RELATED INFORMATION

MPI_SCATTER
MPI_GATHER

NAME

MPI_Send, MPI_SEND-Blocking standard mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Send(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
      INTEGER TAG,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (non-negative integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of the destination task in **comm**(integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a blocking standard mode send. MPI_SEND causes **count** elements of type **datatype** to be sent from **buf** to the task specified by **dest**. **dest** is a task rank which can be any value from 0 to n–1, where n is the number of tasks in **comm**.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_ISEND
MPI_BSEND
MPI_SSEND
MPI_RSEND
MPI_SENDRECV

NAME

MPI_Send_init, MPI_SEND_INIT-Creates a persistent standard mode send request.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Send_init(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SEND_INIT(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
      INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements to be sent (integer) (IN)
datatype	is the type of each element (handle) (IN)
dest	is the rank of the destination task (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a persistent communication request for a standard mode send operation, and binds to it all arguments of a send operation. MPI_START or MPI_STARTALL is used to activate the send.

NOTES

See MPI_SEND for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_START
MPI_ISEND

NAME

MPI_Sendrecv, MPI_SENDRECV-A blocking send and receive operation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Sendrecv(void* sendbuf,int sendcount,MPI_Datatype sendtype,
     int dest,int sendtag,void *recvbuf,int recvcount,MPI_Datatype recvtype,
     int source,int recvtag,MPI_Comm comm,MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SENDRECV(CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,
      INTEGER DEST,INTEGER SENDTAG,CHOICE RECVBUF,INTEGER RECVCOUNT,
      INTEGER RECVTYPE,INTEGER SOURCE,INTEGER RECVTAG,INTEGER COMM,
      INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

sendbuf	is the initial address of the send buffer (choice) (IN)
sendcount	is the number of elements to be sent (integer) (IN)
sendtype	is the type of elements in the send buffer (handle) (IN)
dest	is the rank of the destination task (integer) (IN)
sendtag	is the send tag (integer) (IN)
recvbuf	is the initial address of the receive buffer (choice) (OUT)
recvcount	is the number of elements to be received (integer) (IN)
recvtype	is the type of elements in the receive buffer (handle) (IN)
source	is the rank of the source task or MPI_ANY_SOURCE (integer) (IN)
recvtag	is the receive tag or MPI_ANY_TAG (integer) (IN)
comm	is the communicator (handle) (IN)
status	is the status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a blocking send and receive operation. Send and receive use the same communicator but can use different tags. The send and the receive buffers must be disjoint and can have different lengths and datatypes.

ERRORS

Invalid count(s)	**count** < 0
Invalid datatype(s)	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid source	**source** < 0 or **source** > = groupsize
Invalid communicator	
Invalid tag(s)	**tag** < 0
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_SENDRECV_REPLACE
MPI_SEND
MPI_RECV

NAME

MPI_Sendrecv_replace, MPI_SENDRECV_REPLACE- Blocking send and receive operation using a common buffer.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Sendrecv_replace(void* buf,int count,MPI_Datatype datatype,
     int dest,int sendtag,int source,int recvtag,MPI_Comm comm,
     MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SENDRECV_REPLACE(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,
      INTEGER DEST,INTEGER SENDTAG,INTEGER SOURCE,INTEGER RECVTAG,
      INTEGER COMM,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send and receive buffer (choice) (INOUT)
count	is the number of elements to be sent and received (integer) (IN)
datatype	is the type of elements in the send and receive buffer (handle) (IN)
dest	is the rank of the destination task (integer) (IN)
sendtag	is the send message tag (integer) (IN)
source	is the rank of the source task or MPI_ANY_SOURCE (integer) (IN)
recvtag	is the receive message tag or MPI_ANY_TAGE (integer) (IN)
comm	is the communicator (handle) (IN)
status	is the status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a blocking send and receive operation using a common buffer. Send and receive use the same buffer so the message sent is replaced with the message received.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid source	**source** < 0 or **source** > = groupsize
Invalid communicator	
Invalid tag(s)	**tag** < 0
Out of memory	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_SENDRECV

NAME

MPI_Ssend, MPI_SSEND-Blocking synchronous mode send.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Ssend(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SSEND(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,
    INTEGER TAG,INTEGER COMM,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements in the send buffer (integer) (IN)
datatype	is the datatype of each send buffer element (handle) (IN)
dest	is the rank of the destination task (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine is a blocking synchronous mode send. This a non-local operation. It can be started whether or not a matching receive was posted. However, the send will complete only when a matching receive is posted and the receive operation has started to receive the message sent by MPI_SSEND.

The completion of MPI_SSEND indicates that the send buffer is freed and also that the receiver has started executing the matching receive. If both sends and receives are blocking operations, the synchronous mode provides synchronous communication.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	

MPI already finalized

RELATED INFORMATION

MPI_ISSEND
MPI_SEND

NAME

MPI_Ssend_init, MPI_SSEND_INIT-Creates a persistent synchronous mode send request.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Ssend_init(void* buf,int count,MPI_Datatype datatype,
    int dest,int tag,MPI_Comm comm,MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_SSEND_INIT(CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST

      INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

buf	is the initial address of the send buffer (choice) (IN)
count	is the number of elements to be sent (integer) (IN)
datatype	is the type of each element (handle) (IN)
dest	is the rank of the destination task (integer) (IN)
tag	is the message tag (integer) (IN)
comm	is the communicator (handle) (IN)
request	is the communication request (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine creates a persistent communication object for a synchronous mode send operation. MPI_START or MPI_STARTALL can be used to activate the send.

NOTES

See MPI_SSEND for additional information.

ERRORS

Invalid count	**count** < 0
Invalid datatype	
Type not committed	
Invalid destination	**dest** < 0 or **dest** > = groupsize
Invalid tag	**tag** < 0
Invalid comm	
MPI not initialized	

MPI already finalized

RELATED INFORMATION

MPI_START
MPI_ISSEND

NAME

MPI_Start, MPI_START-Activates a persistent request operation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Start(MPI_Request *request);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_START(INTEGER REQUEST,INTEGER IERROR)
```

PARAMETERS

request is a communication request (handle) (INOUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_START activates a persistent request operation. **request** is a handle returned by MPI_RECV_INIT, MPI_RSEND_INIT, MPI_SSEND_INIT, MPI_BSEND_INIT or MPI_SEND_INIT. Once the call is made, do not access the communication buffer until the operation completes.

If the request is for a send with ready mode, then a matching receive must be posted before the call is made. If the request is for a buffered send, adequate buffer space must be available.

ERRORS

Invalid request

Request not persistent

Request already active

Insufficient buffer space only if buffered send

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_STARTALL
MPI_SEND_INIT
MPI_BSEND_INIT
MPI_RSEND_INIT
MPI_SSEND_INIT
MPI_RECV_INIT

NAME

MPI_Startall, MPI_STARTALL-Activates a collection of persistent request operations.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Startall(int count,MPI_request *array_of_requests);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_STARTALL(INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER IERROR)
```

PARAMETERS

count	is the list length (integer) (IN)
array_of_requests	is the array of requests (array of handle) (INOUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_STARTALL starts all communications associated with request operations in **array_of_requests**.

A communication started with MPI_STARTALL is completed by a call to one of the MPI wait or test operations. The request becomes inactive after successful completion but is not deallocated and can be reactivated by an MPI_STARTALL. If a request is for a send with ready mode, then a matching receive must be posted before the call. If a request is for a buffered send, adequate buffer space must be available.

ERRORS

Invalid count	
Invalid request array	
Request(s) invalid	
Request(s) not persistent	
Request(s) active	
Insufficient buffer space	only if a buffered send
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_START

NAME

MPI_Test, MPI_TEST-Checks to see if a nonblocking send or receive operation has completed.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Test(MPI_Request *request,int *flag,MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TEST(INTEGER REQUEST,INTEGER FLAG,INTEGER STATUS(MPI_STATUS_SIZE),
    INTEGER IERROR)
```

PARAMETERS

request	is the communication request (handle) (INOUT)
flag	*true* if operation completed (logical) (OUT)
status	status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_TEST returns **flag** = **true** if the operation identified by **request** is complete. The status object is set to contain information on the completed operation. If the communication object associated was created by a nonblocking send or receive call, then the object is deallocated and **request** is set to MPI_REQUEST_NULL. Otherwise, **flag** = **false** and the status object is undefined. MPI_TEST is a local operation. The status object can be queried for information about the operation. (See MPI_WAIT.)

You can call MPI_TEST with a null or inactive **request** argument. The operation returns **flag** = **true** and empty status.

The error field of MPI_Status is never modified. The success or failure is indicated by the return code only.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argument, the **status** from this wait/test does not contain meaningful source, tag or message size information.

When you use this routine in a threaded application, make sure the request is tested on only one thread. The request does not have to be tested on the thread that created the request. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid request handle

Truncation occurred

MPI not initialized

MPI already finalized

Develop mode error if:

Illegal buffer update

RELATED INFORMATION

MPI_TESTALL
MPI_TESTSOME
MPI_TESTANY
MPI_WAIT

NAME

MPI_Test_cancelled, MPI_TEST_CANCELLED-Tests whether a nonblocking operation was cancelled.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Test_cancelled(MPI_Status status,int *flag);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TEST_CANCELLED(INTEGER STATUS(MPI_STATUS_SIZE),INTEGER FLAG,
    INTEGER IERROR)
```

PARAMETERS

status is a status object (status) (IN)
Note that in FORTRAN a single status object is an array of integers.

flag **true** if the operation was cancelled (logical) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_TEST_CANCELLED returns **flag** = **true** if the communication associated with the status object was cancelled successfully. In this case, all other fields of **status** (such as count or tag) are undefined. Otherwise, **flag** = **false** is returned. If a receive operation might be cancelled, you should call MPI_TEST_CANCELLED first to check if the operation was cancelled, before checking on the other fields of the return status.

ERRORS

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_CANCEL

NAME

MPI_Testall, MPI_TESTALL-Tests a collection of nonblocking operations for completion.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Testall(int count,MPI_Request *array_of_requests,
    int *flag,MPI_Status *array_of_statuses);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TESTALL(INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER FLAG,
      INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*),INTEGER IERROR)
```

PARAMETERS

count	is the number of requests to test (integer) (IN)
array_of_requests	is an array of requests of length **count** (array of handles) (INOUT)
flag	(logical) (OUT)
array_of_statuses	is an array of status of length **count** objects (array of status) (OUT) Note that in FORTRAN a status object is itself an array.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine tests a collection of nonblocking operations for completion. **flag** = **true** is returned if all communications associated with active handles in the array completed, or when no handle in the list is active.

Each status entry of an active handle request is set to the status of the corresponding communication. A request allocated by a nonblocking communication call is deallocated and the handle is set to MPI_REQUEST_NULL.

Each status entry of a null or inactive handle is set to **empty**. If one or more requests have not completed, **flag** = **false** is returned. No request is modified and the values of the status entries are undefined.

The error fields are never modified unless the function gives a return code of MPI_ERR_IN_STATUS. In which case, the error field of every MPI_Status is modified to reflect the result of the corresponding request.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argument, the **status** from this wait/test does not contain meaningful source, tag or message size information.

When you use this routine in a threaded application, make sure the request is tested on only one thread. The request does not have to be tested on the thread that created it. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count	**count** <0
Invalid request array	
Invalid request(s)	
Truncation occurred	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_TEST
MPI_WAITALL

NAME

MPI_Testany, MPI_TESTANY-Tests for the completion of any specified request.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Testany(int count,MPI_Request *array_of_requests,
    int *index,int *flag,MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TESTANY(INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER INDEX,
     INTEGER FLAG,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

count	is the list length (integer) (IN)
array_of_requests	is the array of request (array of handles) (INOUT)
index	is the index of the operation that completed, or MPI_UNDEFINED is no operation completed (OUT)
flag	**true** if one of the operations is complete (logical) (OUT)
status	status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

If one of the operations has completed, MPI_TESTANY returns **flag** = **true** and returns in **index** the index of this request in the array, and returns in **status** the status of the operation. If the request was allocated by a nonblocking communication call, then the request is deallocated and the handle is set to MPI_REQUEST_NULL.

If none of the operations has completed, it returns **flag** = **false** and returns a value of MPI_UNDEFINED in **index**, and **status** is undefined. The array can contain null or inactive handles. When the array contains no active handles, then the call returns immediately with **flag** = **true**, **index** = MPI_UNDEFINED, and empty **status**.

MPI_TESTANY(**count, array_of_requests, index, flag, status**) has the same effect as the execution of MPI_TEST(**array_of_requests[i], flag, status**), for **i** = 0, 1, ..., **count-1**, in some arbitrary order, until one call returns **flag** = **true**, or all fail.

The error fields are never modified unless the function gives a return code of MPI_ERR_IN_STATUS. In which case, the error field of every MPI_Status is modified to reflect the result of the corresponding request.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argument, the **status** from this wait/test does not contain meaningful source, tag or message size information.

When you use this routine in a threaded application, make sure the request is tested on only one thread. The request does not have to be tested on the thread that created it. See Appendix G, "Programming Considerations for User Applications in POE" on page 329 for more information on programming with MPI in a threaded environment.

NOTES

The array is indexed from zero in C and from one in FORTRAN.

ERRORS

Invalid count	**count** <0
Invalid request array	
Invalid request(s)	
Truncation occurred	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_TEST
MPI_WAITANY

NAME

MPI_Testsome, MPI_TESTSOME-Tests a collection of nonblocking operations for completion.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Testsome(int incount,MPI_Request *array_of_requests,
    int *outcount,int *array_of_indices,
    MPI_Status *array_of_statuses);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TESTSOME(INTEGER INCOUNT,INTEGER ARRAY_OF_REQUESTS(*),
      INTEGER OUTCOUNT,INTEGER ARRAY_OF_INDICES(*),
      INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*),INTEGER IERROR)
```

PARAMETERS

incount	is the length of **array_of_requests** (integer) (IN)
array_of_requests	is the array of requests (array of handles) (INOUT)
outcount	is the number of completed requests (integer) (OUT)
array_of_indices	is the array of indices of operations that completed (array of integers) (OUT)
array_of_statuses	is the array of status objects for operations that completed (array of status) (OUT) Note that in FORTRAN a status object is itself an array.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine tests a collection of nonblocking operations for completion. MPI_TESTSOME behaves like MPI_WAITSOME except that MPI_TESTSOME is a local operation and returns immediately. **outcount** = 0 is returned when no operation has completed.

When a request for a receive repeatedly appears in a list of requests passed to MPI_TESTSOME and a matching send is posted, then the receive eventually succeeds unless the send is satisfied by another receive. This fairness requirement also applies to send requests.

The error fields are never modified unless the function gives a return code of MPI_ERR_IN_STATUS. In which case, the error field of every MPI_Status is modified to reflect the result of the corresponding request.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argu-

ment, the **status** from this wait/test does not contain meaningful source, tag or message size information.

When you use this routine in a threaded application, make sure the request is tested on only one thread. The request does not have to be tested on the thread that created it. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count	**count** < 0
Invalid request array	
Invalid request(s)	
Truncation occurred	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_TEST
MPI_TESTSOME

NAME

MPI_Topo_test, MPI_TOPO_TEST-Returns the type of virtual topology associated with a communicator.

C SYNOPSIS

```
#include <mpi.h>
MPI_Topo_test(MPI_Comm comm,int *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TOPO_TEST(INTEGER COMM,INTEGER STATUS,INTEGER IERROR)
```

PARAMETERS

comm	is the communicator (handle) (IN)
status	is the topology type of communicator **comm** (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the type of virtual topology associated with a communicator. The output of **status** will be as follows:

MPI_GRAPH	graph topology
MPI_CART	cartesian topology

MPI_UNDEFINED no topology

ERRORS

MPI not initialized

MPI already finalized

Invalid communicator

RELATED INFORMATION

MPI_CART_CREATE
MPI_GRAPH_CREATE

NAME

MPI_Type_commit, MPI_TYPE_COMMIT-Makes a datatype ready for use in communications.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_commit(MPI_Datatype *datatype);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_COMMIT(INTEGER DATATYPE,INTEGER IERROR)
```

PARAMETERS

datatype is the datatype that is to be committed (handle) (INOUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

A datatype object must be committed before you can use it in communication. You can use an uncommitted datatype as an argument in datatype constructors.

This routine makes a datatype ready for use in communications. The datatype is the formal description of a communication buffer. It is not the content of the buffer.

Once the datatype is committed it can be repeatedly reused to communicate the changing contents of a buffer or buffers with different starting addresses.

NOTES

Basic datatypes are precommitted.

ERRORS

Invalid datatype

Type is already committed

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_CONTIGUOUS
MPI_TYPE_VECTOR
MPI_TYPE_HVECTOR
MPI_TYPE_INDEXED
MPI_TYPE_HINDEXED
MPI_TYPE_STRUCT
MPI_TYPE_FREE

NAME

MPI_Type_contiguous, MPI_TYPE_CONTIGUOUS-Returns a new datatype that represents the concatenation of **count** instances of **oldtype**.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_contiguous(int count,MPI_Datatype oldtype,
    MPI_Datatype *newtype);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_CONTIGUOUS(INTEGER COUNT,INTEGER OLDTYPE,INTEGER NEWTYPE,
      INTEGER IERROR)
```

PARAMETERS

count	is the replication **count** (non-negative integer) (IN)
oldtype	is the old datatype (handle) (IN)
newtype	is the new datatype (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns a new datatype that represents the concatenation of **count** instances of **oldtype**. MPI_TYPE_CONTIGUOUS allows replication of a datatype into contiguous locations.

NOTES

newtype must be committed using MPI_TYPE_COMMIT before being used for communication.

ERRORS

Invalid count	**count** < 0
Undefined oldtype	
Oldtype is MPI_LB, MPI_UB, or MPI_PACKED	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_TYPE_VECTOR
MPI_TYPE_COMMIT

NAME

MPI_Type_extent, MPI_TYPE_EXTENT-Returns the extent of any defined datatype.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_extent(MPI_Datatype datatype,int *extent);
```

FORTRAN SYNOPSIS

```
MPI_TYPE_EXTENT(INTEGER DATATYPE,INTEGER EXTENT,INTEGER IERROR)
```

PARAMETERS

datatype is the datatype (handle) (IN)

extent is the datatype extent (integer) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the extent of a datatype. The extent of a datatype is the span from the first byte to the last byte occupied by entries in this datatype and rounded up to satisfy alignment requirements.

NOTES

Rounding for alignment is not done when MPI_UB is used to define the datatype. Types defined with MPI_LB, MP_UB or with any type that itself contains MPI_LB or MPI_UB may return an extent which is not directly related to the layout of data in memory. Refer to MPI_Type_struct for more information on MPI_LB and MPI_UB.

ERRORS

Invalid datatype

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_SIZE
MPI_TYPE_STRUCT

NAME

MPI_Type_free, MPI_TYPE_FREE-Marks a datatype for deallocation.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_free(MPI_Datatype *datatype);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_FREE(INTEGER DATATYPE,INTEGER IERROR)
```

PARAMETERS

datatype is the datatype to be freed (handle) (INOUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine marks the datatype object associated with **datatype** for deallocation. It sets **datatype** to MPI_DATATYPE_NULL. All communication currently using this datatype completes normally. Derived datatypes defined from the freed datatype are not affected.

ERRORS

Invalid datatype

Predefined datatype

Type is already free

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_COMMIT

NAME

MPI_Type_hindexed, MPI_TYPE_HINDEXED-Returns a new datatype that represents **count** blocks. Each block is defined by an entry in **array_of_blocklengths** and **array_of_displacements**. Displacements are expressed in bytes.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_hindexed(int count,int *array_of_blocklengths,
    MPI_Aint *array_of_displacements,MPI_Datatype oldtype,
    MPI_Datatype *newtype);
```

FORTRAN SYNOPSIS

```
MPI_TYPE_HINDEXED(INTEGER COUNT,INTEGER ARRAY_OF_BLOCKLENGTHS(*),
    INTEGER ARRAY_OF DISPLACEMENTS(*),INTEGER OLDTYPE,INTEGER NEWTYPE,
    INTEGER IERROR)
```

PARAMETERS

count	is the number of blocks and the number of entries in **array_of_displacements** and **array_of_blocklengths** (non-negative integer) (IN)
array_of_blocklengths	is the number of instances of **oldtype** for each block (array of non-negative integers) (IN)
array_of_displacements	is a byte displacement for each block (array of integer) (IN)
oldtype	is the old datatype (handle) (IN)
newtype	is the new datatype (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns a new datatype that represents **count** blocks. Each is defined by an entry in **array_of_blocklengths** and **array_of_displacements**. Displacements are expressed in bytes rather than in multiples of the **oldtype** extent as in MPI_TYPE_INDEXED.

NOTES

newtype must be committed using MPI_TYPE_COMMIT before being used for communication.

ERRORS

Invalid count	**count** < 0
Invalid blocklength	**blocklength** [i] < 0
Undefined oldtype	
Oldtype is MPI_LB, MPI_UB or MPI_PACKED	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_TYPE_INDEXED
MPI_TYPE_COMMIT

NAME

MPI_Type_hvector, MPI_TYPE_HVECTOR-Returns a new datatype that represents equally spaced blocks. The spacing between the start of each block is given in bytes.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_hvector(int count,int blocklength,MPI_Aint stride,
    MPI_Datatype oldtype,MPI_Datatype *newtype);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_HVECTOR(INTEGER COUNT,INTEGER BLOCKLENGTH,INTEGER STRIDE,
    INTEGER OLDTYPE,INTEGER NEWTYPE,INTEGER IERROR)
```

PARAMETERS

count	is the number of blocks (non-negative integer) (IN)
blocklength	is the number of **oldtype** instances in each block (non-negative integer) (IN)
stride	is an integer specifying the number of bytes between start of each block. (IN)
oldtype	is the old datatype (handle) (IN)
newtype	is the new datatype (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns a new datatype that represents **count** equally spaced blocks. Each block is a concatenation of **blocklength** instances of **oldtype**. The origins of the blocks are spaced **stride** units apart where the counting unit is one byte.

NOTES

newtype must be committed using MPI_TYPE_COMMIT before being used for communication.

ERRORS

Invalid count	**count** < 0
Invalid blocklength	**blocklength** < 0
Undefined oldtype	
Oldtype is MPI_LB, MPI_UB or MPI_PACKED	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_TYPE_VECTOR
MPI_TYPE_COMMIT

NAME

MPI_Type_indexed, MPI_TYPE_INDEXED-Returns a new datatype that represents **count** blocks. Each block is defined by an entry in **array_of_blocklengths** and **array_of_displacements**. Displacements are expressed in units of extent(**oldtype**).

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_indexed(int count,int *array_of_blocklengths,
    int *array_of_displacements,MPI_Datatype oldtype,
    MPI_datatype *newtype);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_INDEXED(INTEGER COUNT,INTEGER ARRAY_OF_BLOCKLENGTHS(*),
    INTEGER ARRAY_OF DISPLACEMENTS(*),INTEGER OLDTYPE,INTEGER NEWTYPE,
    INTEGER IERROR)
```

PARAMETERS

count	is the number of blocks and the number of entries in **array_of_displacements** and **array_of_blocklengths** (non-negative integer) (IN)
array_of_blocklengths	is the number of instances of **oldtype** in each block (array of non-negative integers) (IN)
array_of_displacements	is the displacement of each block in units of extent(**oldtype**) (array of integer) (IN)
oldtype	is the old datatype (handle) (IN)
newtype	is the new datatype (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns a new datatype that represents **count** blocks. Each is defined by an entry in **array_of_blocklengths** and **array_of_displacements**. Displacements are expressed in units of extent(**oldtype**).

NOTES

newtype must be committed using MPI_TYPE_COMMIT before being used for communication.

ERRORS

Invalid count	**count** < 0
Invalid count	**blocklength** [i] < 0
Undefined oldtype	
Oldtype is MPI_LB, MPI_UB or MPI_PACKED	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_TYPE_HINDEXED
MPI_TYPE_COMMIT

NAME

MPI_Type_lb, MPI_TYPE_LB-Returns the lower bound of a datatype.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_lb(MPI_Datatype datatype,int *displacement);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_LB(INTEGER DATATYPE,INTEGER DISPLACEMENT,INTEGER IERROR)
```

PARAMETERS

datatype is the datatype (handle) (IN)

displacement is the displacement of lower bound from the origin in bytes (integer) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the lower bound of a specific datatype.

Normally the lower bound is the offset of the lowest address byte in the datatype. Datatype constructors with explicit MPI_LB and vector constructors with negative stride can produce lb < 0. Lower bound cannot be greater than upper bound. For a type with MPI_LB in its ancestry, the value returned by MPI_TYPE_LB may not be related to the displacement of the lowest address byte. Refer to MPI_TYPE_STRUCT for more information on MPI_LB and MPI_UB.

ERRORS

Invalid datatype

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_UB
MPI_TYPE_STRUCT

NAME

MPI_Type_size, MPI_TYPE_SIZE-Returns the number of bytes represented by any defined datatype.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_size(MPI_Datatype datatype,int *size);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_SIZE(INTEGER DATATYPE,INTEGER SIZE,INTEGER IERROR)
```

PARAMETERS

datatype	is the datatype (handle) (IN)
size	is the datatype size (integer) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the total number of bytes in the type signature associated with **datatype**. Entries with multiple occurrences in the datatype are counted.

ERRORS

Invalid datatype

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_EXTENT

NAME

MPI_Type_struct, MPI_TYPE_STRUCT-Returns a new datatype that represents **count** blocks. Each is defined by an entry in **array_of_blocklengths**, **array_of_displacements** and **array_of_types**. Displacements are expressed in bytes.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_struct(int count,int *array_of_blocklengths,
    MPI_Aint *array_of_displacements,MPI_Datatype *array_of_types,
    MPI_datatype *newtype);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_STRUCT(INTEGER COUNT,INTEGER ARRAY_OF_BLOCKLENGTHS(*),
    INTEGER ARRAY_OF DISPLACEMENTS(*),INTEGER ARRAY_OF_TYPES(*),
    INTEGER NEWTYPE,INTEGER IERROR)
```

PARAMETERS

count	is an integer specifying the number of blocks. It is also the number of entries in arrays **array_of_types**, **array_of_displacements** and **array_of_blocklengths**. (IN)
array_of_blocklengths	is the number of elements in each block (array of integer). That is, **array_of_blocklengths(i)** specifies the number of instances of type **array_of_types(i)**in block(i). (IN)
array_of_displacements	is the byte displacement of each block (array of integer) (IN)
array_of_types	is the datatype comprising each block. That is, block(i) is made of a concatenation of type **array_of_types(i)**. (array of handles to datatype objects) (IN)
newtype	is the new datatype (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns a new datatype that represents **count** blocks. Each is defined by an entry in **array_of_blocklengths**, **array_of_displacements** and **array_of_types**. Displacements are expressed in bytes.

MPI_TYPE_STRUCT is the most general type constructor. It allows each block to consist of replications of different datatypes. This is the only constructor which allows MPI pseudo types MPI_LB and MPI_UB. Without these pseudo types, the

extent of a datatype is the range from the first byte to the last byte rounded up as needed to meet boundary requirements. For example, if a type is made of an integer followed by 2 characters, it will still have an extent of 8 because it is padded to meet the boundary constraints of an int. This is intended to match the behavior of a compiler defining an array of such structures.

Because there may be cases in which this default behavior is not correct, MPI provides a means to set explicit upper and lower bounds which may not be directly related to the lowest and highest displacement datatype. When the pseudo type MPI_UB is used, the upper bound will be the value specified as the displacement of the MPI_UB block. No rounding for alignment is done. MPI_LB can be used to set an explicit lower bound but its use does not suppress rounding. When MPI_UB is not used, the upper bound of the datatype is adjusted to make the extent a multiple of the type's most boundary constrained component.

The marker placed by a MPI_LB or MPI_UB is 'sticky'. For example, assume type A is defined with a MPI_UB at 100. Type B is defined with a type A at 0 and a MPI_UB at 50. In effect, type B has received a MPI_UB at 50 and an inherited MPI_UB at 100. Because the inherited MPI_UB is higher, it is kept in the type B definition and the MPI_UB explicitly placed at 50 is discarded.

NOTES

newtype must be committed using MPI_TYPE_COMMIT before being used for communication.

ERRORS

Invalid count **count** < 0

Invalid blocklength **blocklength**[i] < 0

Undefined oldtype in array_of_types

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_HINDEXED
MPI_TYPE_COMMIT

NAME

MPI_Type_ub, MPI_TYPE_UB-Returns the upper bound of a datatype.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_ub(MPI_Datatype datatype,int *displacement);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_UB(INTEGER DATATYPE,INTEGER DISPLACEMENT,INTEGER IERROR)
```

PARAMETERS

datatype is the datatype (handle) (IN)

displacement is the displacement of upper bound from origin in bytes (integer) (OUT)

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine returns the upper bound of a specific datatype.

The upper bound is the displacement you use in locating the origin byte of the next instance of **datatype** for operations which use count and datatype. In the normal case, ub represents the displacement of the highest address byte of the datatype + e (where e >= 0 and results in (ub – lb) being a multiple of the boundary requirement for the most boundary constrained type in the datatype). If MPI_UB is used in a type constructor, no alignment adjustment is done so ub is exactly as you set it.

For a type with MPI_UB in its ancestry, the value returned by MPI_TYPE_UB may not be related to the displacement of the highest address byte (with rounding). Refer to MPI_TYPE_STRUCT for more informatin on MPI_LB and MPI_UB.

ERRORS

Invalid datatype

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_TYPE_LB

NAME

MPI_Type_vector, MPI_TYPE_VECTOR-Returns a new datatype that represents equally spaced blocks. The spacing between the start of each block is given in units of extent (**oldtype**).

C SYNOPSIS

```
#include <mpi.h>
int MPI_Type_vector(int count,int blocklength,int stride,
    MPI_Datatype oldtype,MPI_Datatype *newtype);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_TYPE_VECTOR(INTEGER COUNT,INTEGER BLOCKLENGTH,
    INTEGER STRIDE,INTEGER OLDTYPE,INTEGER NEWTYPE,INTEGER IERROR)
```

PARAMETERS

count	is the number of blocks (non-negative integer) (IN)
blocklength	is the number of **oldtype** instances in each block (non-negative integer) (IN)
stride	is the number of units between the start of each block (integer) (IN)
oldtype	is the old datatype (handle) (IN)
newtype	is the new datatype (handle) (OUT)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This function returns a new datatype that represents **count** equally spaced blocks. Each block is a a concatenation of **blocklength** instances of **oldtype**. The origins of the blocks are spaced **stride** units apart where the counting unit is extent(**oldtype**). That is, from one origin to the next in bytes = **stride** * extent (**oldtype**).

NOTES

newtype must be committed using MPI_TYPE_COMMIT before being used for communication.

ERRORS

Invalid count	**count** < 0
Invalid blocklength	**blocklength** < 0
Undefined oldtype	
Oldtype is MPI_LB, MPI_UB or MPI_PACKED	
MPI not initialized	

MPI already finalized

RELATED INFORMATION

MPI_TYPE_HVECTOR
MPI_TYPE_COMMIT

NAME

MPI_Unpack, MPI_UNPACK-Unpacks the message into the specified receive buffer from the specified packed buffer.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Unpack(void* inbuf,int insize,int *position,
    void *outbuf,int outcount,MPI_Datatype datatype,
    MPI_Comm comm);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_UNPACK(CHOICE INBUF,INTEGER INSIZE,INTEGER POSITION,
     CHOICE OUTBUF,INTEGER OUTCOUNT,INTEGER DATATYPE,INTEGER COMM,
     INTEGER IERROR)
```

PARAMETERS

inbuf	is the input buffer start (choice) (IN)
insize	is an integer specifying the size of input buffer in bytes (IN)
position	is an integer specifying the current packed buffer offset in bytes (INOUT)
outbuf	is the output buffer start (choice) (OUT)
outcount	is an integer specifying the number of instances of **datatype** to be unpacked (IN)
datatype	is the datatype of each output data item (handle) (IN)
comm	is the communicator for the packed message (handle) (IN)
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine unpacks the message specified by **outbuf**, **outcount**, and **datatype** from the buffer space specified by **inbuf** and **insize**. The output buffer is any receive buffer allowed in MPI_RECV. The input buffer is any contiguous storage space containing **insize** bytes and starting at address **inbuf**.

The input value of **position** is the beginning offset in the input buffer for the data to be unpacked. The output value of **position** is the offset in the input buffer following the data already unpacked. That is, the starting point for another call to MPI_UNPACK. **comm** is the communicator that was used to receive the packed message.

NOTES

In MPI_UNPACK the **outcount** argument specifies the actual number of items to be unpacked. The size of the corresponding message is the increment in **position**.

ERRORS

Invalid outcount	**outcount** < 0
Invalid datatype	
Type is not committed	
Invalid communicator	
Inbuf too small	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_PACK

NAME

MPI_Wait, MPI_WAIT-Waits for a nonblocking operation to complete.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Wait(MPI_Request *request,MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_WAIT(INTEGER REQUEST,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

request is the request to wait for (handle) (INOUT)

status is the status object (status) (OUT)
Note that in FORTRAN a single status object is an array of integers.

IERROR is the FORTRAN return code. It is always the last argument.

DESCRIPTION

MPI_WAIT returns after the operation identified by **request** completes. If the communication object associated with **request** was created by a nonblocking send or receive call, then the object is deallocated by the call to MPI_WAIT and **request** is set to MPI_REQUEST_NULL. MPI_WAIT is a non-local operation.

You can call MPI_WAIT with a null or inactive **request** argument. The operation returns immediately. The **status** argument returns **tag** = MPI_ANY_TAG, **source** = MPI_ANY_SOURCE. The **status** argument is also internally configured so that calls to MPI_GET_COUNT and MPI_GET_ELEMENTS return **count** = 0. (This is called an **empty** status.)

Information on the completed operation is found in **status**. You can query the status object for a send or receive operation with a call to MPI_TEST_CANCELLED. For receive operations, you can also retrieve information from **status** with MPI_GET_COUNT and MPI_GET_ELEMENTS. If wildcards were used by the receive for either the source or tag, the actual source and tag can be retrieved by:

In C:

```
source = status.MPI_SOURCE
tag = status.MPI_TAG
```

In FORTRAN:

```
source = status(MPI_SOURCE)
tag = status(MPI_TAG)
```

The error field of MPI_Status is never modified. The success or failure is indicated by the return code only.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argument, the **status** from this wait/test does not contain meaningful source, tag or message size information.

When you use this routine in a threaded application, make sure that the wait for a given request is done on only one thread. The wait does not have to be done on the thread that created the request. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid request handle

Truncation occurred

MPI not initialized

MPI already finalized

Develop mode error if:

Illegal buffer update

RELATED INFORMATION

MPI_WAITALL
MPI_WAITSOME
MPI_WAITANY
MPI_TEST

NAME

MPI_Waitall, MPI_WAITALL-Waits for a collection of nonblocking operations to complete.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Waitall(int count,MPI_Request *array_of_requests,
    MPI_Status *array_of_statuses);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_WAITALL(INTEGER COUNT,INTEGER ARRAY_OF_ REQUESTS(*),
      INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*),INTEGER IERROR)
```

PARAMETERS

count	is the lists length (integer) (IN)
array_of_requests	is an array of requests of length **count** (array of handles) (INOUT)
array_of_statuses	is an array of status objects of length **count** (array of status) (OUT) Note that in FORTRAN a status object is itself an array.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine blocks until all communication operations associated with active handles in the list complete, and returns the status of each operation. **array_of_requests** and **array_of statuses** contain **count** entries.

The **i**th entry in **array_of_statuses** is set to the return status of the **i**th operation. Requests created by nonblocking communication operations are deallocated and the corresponding handles in the array are set to MPI_REQUEST_NULL. If **array_of_requests** contains null or inactive handles, MPI_WAITALL sets the status of each one to **empty**.

MPI_WAITALL(**count, array_of_requests, array_of_statuses**) has the same effect as the execution of MPI_WAIT(**array_of_requests[i], array_of_statuses[i]**) for **i** = 0, 1, ..., **count-1**, in some arbitrary order. MPI_WAITALL with an array of length one is equivalent to MPI_WAIT.

The error fields are never modified unless the function gives a return code of MPI_ERR_IN_STATUS. In which case, the error field of every MPI_Status is modified to reflect the result of the corresponding request.

When you use this routine in a threaded application, make sure that the wait for a given request is done on only one thread. The wait does not have to be done on the thread that created it. See Appendix G, “Programming Considerations for User

Applications in POE" on page 329 for more information on programming with MPI in a threaded environment.

ERRORS

Invalid count	**count** <0
Invalid request array	
Invalid request(s)	
Truncation occurred	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_WAIT
MPI_TESTALL

NAME

MPI_Waitany, MPI_WAITANY-Waits for any specified request to complete.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Waitany(int count,MPI_Request *array_of_requests,
    int *index,MPI_Status *status);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_WAITANY(INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER INDEX,
      INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR)
```

PARAMETERS

count	is the list length (integer) (IN)
array_of_requests	is the array of requests (array of handles) (INOUT)
index	is the index of the handle for the operation that completed (integer) (OUT)
status	status object (status) (OUT) Note that in FORTRAN a single status object is an array of integers.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine blocks until one of the operations associated with the active requests in the array has completed. If more then one operation can complete, one is arbitrarily chosen. MPI_WAITANY returns in **index** the index of that request in the array, and in **status** the status of the completed communication. When the request is allocated by a nonblocking communication operation, it is deallocated and the request handle is set to MPI_REQUEST_NULL.

The **array_of_requests** list can contain null or inactive handles. When the list has a length of zero or all entries are null or inactive, the call returns immediately with **index** = MPI_UNDEFINED, and an empty status.

MPI_WAITANY(**count, array_of_requests, index, status**) has the same effect as the execution of MPI_WAIT(**array_of_requests[i], status**), where **i** is the value returned by **index**. MPI_WAITANY with an array containing one active entry is equivalent to MPI_WAIT.

The error fields are never modified unless the function gives a return code of MPI_ERR_IN_STATUS. In which case, the error field of every MPI_Status is modified to reflect the result of the corresponding request.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argu-

ment, the **status** from this wait/test does not contain meaningful source, tag or message size information.

When you use this routine in a threaded application, make sure that the wait for a given request is done on only one thread. The wait does not have to be done on the thread that created it. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

NOTES

In C, the array is indexed from zero and in FORTRAN from one.

ERRORS

Invalid count	**count** < 0
Invalid requests array	
Invalid request(s)	
Truncation occurred	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_WAIT
MPI_TESTANY

NAME

MPI_Waitsome, MPI_WAITSOME-Waits for at least one of a list of nonblocking operations to complete.

C SYNOPSIS

```
#include <mpi.h>
int MPI_Waitsome(int incount,MPI_Request *array_of_requests,
    int *outcount,int *array_of_indices,MPI_Status *array_of_statuses);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
MPI_WAITSOME(INTEGER INCOUNT,INTEGER ARRAY_OF_REQUESTS,INTEGER OUTCOUNT,
      INTEGER ARRAY_OF_INDICES(*),INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*),
      INTEGER IERROR)
```

PARAMETERS

incount	is the length of **array_of_requests**, **array_of_indices**, and **array_of_statuses** (integer) (IN)
array_of_requests	is an array of requests (array of handles) (INOUT)
outcount	is the number of completed requests (integer) (OUT)
array_of_indices	is the array of indices of operations that completed (array of integers) (OUT)
array_of_statuses	is the array of status objects for operations that completed (array of status) (OUT) Note that in FORTRAN a status object is itself an array.
IERROR	is the FORTRAN return code. It is always the last argument.

DESCRIPTION

This routine waits for at least one of a list of nonblocking operations associated with active handles in the list to complete. The number of completed requests from the list of **array_of_requests** is returned in **outcount**. Returns in the first **outcount** locations of the array **array_of_indices** the indices of these operations.

The status for the completed operations is returned in the first **outcount** locations of the array **array_of_statuses**. When a completed request is allocated by a non-blocking communication call, then it is deallocated and the associated handle is set to MPI_REQUEST_NULL.

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

When a request for a receive repeatedly appears in a list of requests passed to MPI_WAITSOME and a matching send was posted, then the receive eventually succeeds unless the send is satisfied by another receive. This fairness requirement also applies to send requests.

The error fields are never modified unless the function gives a return code of MPI_ERR_IN_STATUS. In which case, the error field of every MPI_Status is modified to reflect the result of the corresponding request.

When one of the MPI wait or test calls returns **status** for a nonblocking operation request and the corresponding blocking operation does not provide a **status** argument, the **status** from this wait/test does not contain meaningful source, tag or message size information.

When you use this routine in a threaded application, make sure that the wait for a given request is done on only one thread. The wait does not have to be done on the thread that created it. See Appendix G, “Programming Considerations for User Applications in POE” on page 329 for more information on programming with MPI in a threaded environment.

NOTES

In C, the index within the array **array_of_requests**, is indexed from zero and from one in FORTRAN.

ERRORS

Invalid count	**count** <0
Invalid request(s)	
Invalid index array	
Truncation occurred	
MPI not initialized	
MPI already finalized	

RELATED INFORMATION

MPI_WAIT
MPI_TESTSOME

NAME

MPI_Wtick, MPI_WTICK-Returns the resolution of MPI_Wtime in seconds.

C SYNOPSIS

```
#include <mpi.h>
double MPI_Wtick(void);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
DOUBLE PRECISION MPI_WTICK()
```

PARAMETERS

None.

DESCRIPTION

This routine returns the resolution of MPI_WTIME in seconds, the time in seconds between successive clock ticks.

ERRORS

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_WTIME

NAME

MPI_Wtime, MPI_WTIME-Returns the current value of **time** as a floating point value.

C SYNOPSIS

```
#include <mpi.h>
double MPI_Wtime(void);
```

FORTRAN SYNOPSIS

```
include 'mpif.h'
DOUBLE PRECISION MPI_WTIME()
```

PARAMETERS

None.

DESCRIPTION

This routine returns the current value of **time** as a double precision floating point number of seconds. This value represents elapsed time since some point in the past. This time in the past will not change during the life of the process. You are responsible for converting the number of seconds into other units if you prefer.

NOTES

You can use the attribute key MPI_WTIME_IS_GLOBAL to determine if the values returned by MPI_WTIME on different nodes are synchronized. See MPI_ATTR_GET.

ERRORS

MPI not initialized

MPI already finalized

RELATED INFORMATION

MPI_WTICK
MPI_ATTR_GET

Appendix A. MPI Subroutine Bindings Quick Reference

The following tables summarize the C and FORTRAN binding information for all the subroutines listed in this book.

Note: FORTRAN refers to FORTRAN 77 bindings which are officially supported for MPI. However, FORTRAN 77 bindings can be used by Fortran 90. Fortran 90 and High Performance Fortran (HPF) offer array section and assumed shape arrays as parameters on calls. These are not safe with MPI.

Bindings for Nonblocking Collective Communication

The following table lists the C and FORTRAN bindings for nonblocking collective communication routines.

C/FORTRAN Subroutine	C/FORTRAN Binding
MPE_Ibarrier	int MPE_Ibarrier(*MPI_Comm comm,MPI_Request *request*);
MPE_IBARRIER	MPE_IBARRIER(*INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Ibcast	int MPE_Ibcast(*void* buffer,int count,MPI_Datatype datatype,int root,MPI_Comm comm,MPI_Request *request*);
MPE_IBCAST	MPE_IBCAST(*CHOICE BUFFER,INTEGER COUNT,INTEGER DATATYPE,INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Igather	int MPE_Igather(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype,int root, MPI_Comm comm,MPI_Request *request*);
MPE_IGATHER	MPE_IGATHER(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Igatherv	int MPE_Igatherv(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int *recvcounts,int *displs,MPI_Datatype recvtype,int root,MPI_Comm comm,MPI_Request *request*);
MPE_IGATHERV	MPE_IGATHERV(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE, CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Iscatter	int MPE_Iscatter(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,MPI_Comm comm,MPI_Request *request*);
MPE_ISCATTER	MPE_ISCATTER(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Iscatterv	int MPE_Iscatterv(*void* sendbuf,int *sendcounts,int *displs,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,MPI_Comm comm,MPI_Request *request*);

© Copyright IBM Corp. 1997

C/FORTRAN Subroutine	C/FORTRAN Binding
MPE_ISCATTERV	MPE_ISCATTERV(*CHOICE SENDBUF,INTEGER SENDCOUNTS(*),INTEGER DISPLS(*),INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Iallgather	int MPE_Iallgather(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype, MPI_Comm comm,MPI_Request *request*);
MPE_IALLGATHER	MPE_IALLGATHER(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE, CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Iallgatherv	int MPE_Iallgatherv(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int *recvcounts,int *displs,MPI_Datatype recvtype,MPI_Comm comm,MPI_Request *request*);
MPE_IALLGATHERV	MPE_IALLGATHERV(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE, CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),INTEGER RECVTYPE,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Ialltoall	int MPE_Ialltoall(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype,MPI_Comm comm,MPI_Request *request*);
MPE_IALLTOALL	MPE_IALLTOALL(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Ialltoallv	int MPE_Ialltoallv(*void* sendbuf,int *sendcounts,int *sdispls,MPI_Datatype sendtype,void* recvbuf,int *recvcounts,int *rdispls,MPI_Datatype recvtype,MPI_Comm comm,MPI_Request *request*);
MPE_IALLTOALLV	MPE_IALLTOALV(*CHOICE SENDBUF,INTEGER SENDCOUNTS(*),INTEGER SDISPLS(*),INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER RDISPLS(*),INTEGER RECVTYPE,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Ireduce	int MPE_Ireduce(*void* sendbuf,void* recvbuf,int count,MPI_Datatype datatype,MPI_Op op,int root,MPI_Comm comm,MPI_Request *request*);
MPE_IREDUCE	MPE_IREDUCE(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER OP,INTEGER ROOT,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Iallreduce	int MPE_Iallreduce(*void* sendbuf,void* recvbuf,int count,MPI_Datatype datatype,MPI_Op op,MPI_Comm comm,MPI_Request *request*);
MPE_IALLREDUCE	MPE_IALLREDUCE(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Ireduce_scatter	int MPE_Ireduce_scatter(*void* sendbuf,void* recvbuf,int *recvcounts,MPI_Datatype datatype,MPI_Op op,MPI_Comm comm,MPI_Request *request*);
MPE_IREDUCE_SCATTER	MPE_IREDUCE_SCATTER(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPE_Iscan	int MPE_Iscan(*void* sendbuf,void* recvbuf,int count,MPI_Datatype datatype,MPI_Op op,MPI_Comm comm,MPI_Request *request*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPE_ISCAN	MPE_ISCAN(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)

Bindings for Point-to-Point

Communication and Derived Datatypes

The following table lists the C and FORTRAN bindings for point-to-point communication and derived datatype routines.

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Send	int MPI_Send(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm*);
MPI_SEND	MPI_SEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM, INTEGER IERROR*)
MPI_Recv	int MPI_Recv(*void* buf,int count,MPI_Datatype datatype,int source,int tag, MPI_Comm comm, MPI_Status *status*);
MPI_RECV	MPI_RECV(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER SOURCE, INTEGER TAG,INTEGER COMM,INTEGER STATUS(MPI_STATUS_SIZE),,INTEGER IERROR*)
MPI_Get_count	int MPI_Get_count(*MPI_Status *status, MPI_Datatype datatype, int *count*);
MPI_GET_COUNT	MPI_GET_COUNT(*INTEGER STATUS(MPI_STATUS_SIZE),,INTEGER DATATYPE,INTEGER COUNT, INTEGER IERROR*)
MPI_Bsend	int MPI_Bsend(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm*);
MPI_BSEND	MPI_BSEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST, INTEGER TAG,INTEGER COMM,INTEGER IERROR*)
MPI_Ssend	int MPI_Ssend(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm*);
MPI_SSEND	MPI_SSEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER IERROR*)
MPI_Rsend	int MPI_Rsend(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm*);
MPI_RSEND	MPI_RSEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER IERROR*)
MPI_Buffer_attach	int MPI_Buffer_attach(*void* buffer,int size*);
MPI_BUFFER_ATTACH	MPI_BUFFER_ATTACH(*CHOICE BUFFER,INTEGER SIZE,INTEGER IERROR*)
MPI_Buffer_detach	int MPI_Buffer_detach(*void* buffer,int* size*);
MPI_BUFFER_DETACH	MPI_BUFFER_DETACH(*CHOICE BUFFER,INTEGER SIZE,INTEGER IERROR*)

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Isend	int MPI_Isend(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_ISEND	MPI_ISEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Ibsend	int MPI_Ibsend(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_IBSEND	MPI_IBSEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Issend	int MPI_Issend(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_ISSEND	MPI_ISSEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Irsend	int MPI_Irsend(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_IRSEND	MPI_IRSEND(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Irecv	int MPI_Irecv(*void* buf,int count,MPI_Datatype datatype,int source,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_IRECV	MPI_IRECV(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER SOURCE,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Wait	int MPI_Wait(*MPI_Request *request,MPI_Status *status*);
MPI_WAIT	MPI_WAIT(*INTEGER REQUEST,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR*)
MPI_Test	int MPI_Test(*MPI_Request *request,int *flag,MPI_Status *status*);
MPI_TEST	MPI_TEST(*INTEGER REQUEST,INTEGER FLAG,INTEGER STATUS(MPI_STATUS_SIZE), INTEGER IERROR*)
MPI_Request_free	int MPI_Request_free(*MPI_Request *request*);
MPI_REQUEST_FREE	MPI_REQUEST_FREE(*INTEGER REQUEST,INTEGER IERROR*)
MPI_Waitany	int MPI_Waitany(*int count,MPI_Request *array_of_requests,int *index,MPI_Status *status*);
MPI_WAITANY	MPI_WAITANY(*INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER INDEX, INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR*)
MPI_Testany	int MPI_Testany(*int count, MPI_Request *array_of_requests, int *index, int *flag,MPI_Status *status*);
MPI_TESTANY	MPI_TESTANY(*INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER INDEX,INTEGER FLAG,INTEGER STATUS(MPI_STATUS_SIZE), INTEGER IERROR*)
MPI_Waitall	int MPI_Waitall(*int count,MPI_Request *array_of_requests,MPI_Status *array_of_statuses*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_WAITALL	MPI_WAITALL(*INTEGER COUNT,INTEGER ARRAY_OF_ REQUESTS(*),INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*), INTEGER IERROR*)
MPI_Testall	int MPI_Testall(*int count,MPI_Request *array_of_requests,int *flag,MPI_Status *array_of_statuses*);
MPI_TESTALL	MPI_TESTALL(*INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER FLAG, INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE,*),INTEGER IERROR*)
MPI_Waitsome	int MPI_Waitsome(*int incount,MPI_Request *array_of_requests,int *outcount,int *array_of_indices,MPI_Status *array_of_statuses*);
MPI_WAITSOME	MPI_WAITSOME(*INTEGER INCOUNT,INTEGER ARRAY_OF_REQUESTS,INTEGER OUTCOUNT,INTEGER ARRAY_OF_INDICES(*),INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE),*),INTEGER IERROR*)
MPI_Testsome	int MPI_Testsome(*int incount,MPI_Request *array_of_requests,int *outcount,int *array_of_indices,MPI_Status *array_of_statuses*);
MPI_TESTSOME	MPI_TESTSOME(*INTEGER INCOUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER OUTCOUNT,INTEGER ARRAY_OF_INDICES(*),INTEGER ARRAY_OF_STATUSES(MPI_STATUS_SIZE),*),INTEGER IERROR*)
MPI_Iprobe	int MPI_Iprobe(*int source,int tag,MPI_Comm comm,int *flag,MPI_Status *status*);
MPI_IPROBE	MPI_IPROBE(*INTEGER SOURCE,INTEGER TAG,INTEGER COMM,INTEGER FLAG,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR*)
MPI_Probe	int MPI_Probe(*int source,int tag,MPI_Comm comm,MPI_Status *status*);
MPI_PROBE	MPI_PROBE(*INTEGER SOURCE,INTEGER TAG,INTEGER COMM,INTEGER STATUS(MPI_STATUS_SIZE), INTEGER IERROR*)
MPI_Cancel	int MPI_Cancel(*MPI_Request *request*);
MPI_CANCEL	MPI_CANCEL(*INTEGER REQUEST,INTEGER IERROR*)
MPI_Test_cancelled	int MPI_Test_cancelled(*MPI_Status *status,int *flag*);
MPI_TEST_CANCELLED	MPI_TEST_CANCELLED(*INTEGER STATUS(MPI_STATUS_SIZE),INTEGER FLAG,INTEGER IERROR*)
MPI_Send_init	int MPI_Send_init(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_SEND_INIT	MPI_SEND_INIT(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Bsend_init	int MPI_Bsend_init(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_BSEND_INIT	MPI_SEND_INIT(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Ssend_init	int MPI_Ssend_init(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_SSEND_INIT	MPI_SSEND_INIT(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,IERROR*)
MPI_Rsend_init	int MPI_Rsend_init(*void* buf,int count,MPI_Datatype datatype,int dest,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_RSEND_INIT	MPI_RSEND_INIT(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Recv_init	int MPI_Recv_init(*void* buf,int count,MPI_Datatype datatype,int source,int tag,MPI_Comm comm,MPI_Request *request*);
MPI_RECV_INIT	MPI_RECV_INIT(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER SOURCE,INTEGER TAG,INTEGER COMM,INTEGER REQUEST,INTEGER IERROR*)
MPI_Start	int MPI_Start(*MPI_Request *request*);
MPI_START	MPI_START(*INTEGER REQUEST,INTEGER IERROR*)
MPI_Startall	int MPI_Startall(*int count,MPI_Request *array_of_requests*);
MPI_STARTALL	MPI_STARTALL(*INTEGER COUNT,INTEGER ARRAY_OF_REQUESTS(*),INTEGER IERROR*)
MPI_Sendrecv	int MPI_Sendrecv(*void *sendbuf,int sendcount,MPI_Datatype sendtype,int dest,int sendtag,void *recvbuf,int recvcount, MPI_Datatype recvtype,int source,int recvtag,MPI_Comm comm,MPI_Status *status*);
MPI_SENDRECV	MPI_SENDRECV(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,INTEGER DEST,INTEGER SENDTAG,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER SOURCE,INTEGER RECVTAG,INTEGER COMM,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR*)
MPI_Sendrecv_replace	int MPI_Sendrecv_replace(*void* buf,int count,MPI_Datatype datatype,int dest,int sendtag,int source,int recvtag,MPI_Comm comm,MPI_Status *status*);
MPI_SENDRECV_REPLACE	MPI_SENDRECV_REPLACE(*CHOICE BUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER DEST,INTEGER SENDTAG,INTEGER SOURCE,INTEGER RECVTAG,INTEGER COMM,INTEGER STATUS(MPI_STATUS_SIZE),INTEGER IERROR*)
MPI_Type_contiguous	int MPI_Type_contiguous(*int count,MPI_Datatype oldtype,MPI_Datatype *newtype*);
MPI_TYPE_CONTIGUOUS	MPI_TYPE_CONTIGUOUS(*INTEGER COUNT,INTEGER OLDTYPE,INTEGER NEWTYPE,INTEGER IERROR*)
MPI_Type_vector	int MPI_Type_vector(*int count,int blocklength,int stride,MPI_Datatype oldtype,MPI_Datatype *newtype*);
MPI_TYPE_VECTOR	MPI_TYPE_VECTOR(*INTEGER COUNT,INTEGER BLOCKLENGTH,INTEGER STRIDE,INTEGER OLDTYPE,INTEGER NEWTYPE,INTEGER IERROR*)
MPI_Type_hvector	int MPI_Type_hvector(*int count,int blocklength,MPI_Aint stride,MPI_Datatype oldtype,MPI_Datatype *newtype*);
MPI_TYPE_HVECTOR	MPI_TYPE_HVECTOR(*INTEGER COUNT,INTEGER BLOCKLENGTH,INTEGER STRIDE,INTEGER OLDTYPE,INTEGER NEWTYPE,INTEGER IERROR*)

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Type_indexed	int MPI_Type_indexed(*int count,int *array_of_blocklengths,int *array_of_displacements,MPI_Datatype oldtype, MPI_Datatype *newtype*);
MPI_TYPE_INDEXED	MPI_TYPE_INDEXED(*INTEGER COUNT,INTEGER ARRAY_OF_BLOCKLENGTHS(*),INTEGER ARRAY_OF DISPLACEMENTS(*),INTEGER OLDTYPE,INTEGER NEWTYPE,INTEGER IERROR*)
MPI_Type_hindexed	int MPI_Type_hindexed(*int count,int *array_of_blocklengths,MPI_Aint *array_of_displacements,MPI_Datatype oldtype, MPI_Datatype *newtype*);
MPI_TYPE_HINDEXED	MPI_TYPE_HINDEXED(*INTEGER COUNT,INTEGER ARRAY_OF_BLOCKLENGTHS(*),INTEGER ARRAY_OF DISPLACEMENTS(*),INTEGER OLDTYPE,INTEGER NEWTYPE,INTEGER IERROR*)
MPI_Type_struct	int MPI_Type_struct(*int count,int *array_of_blocklengths, MPI_Aint *array_of_displacements,MPI_Datatype *array_of_types, MPI_Datatype *newtype*);
MPI_TYPE_STRUCT	MPI_TYPE_STRUCT(*INTEGER COUNT,INTEGER ARRAY_OF_BLOCKLENGTHS(*),INTEGER ARRAY_OF DISPLACEMENTS(*),INTEGER ARRAY_OF_TYPES(*),INTEGER NEWTYPE,INTEGER IERROR*)
MPI_Address	int MPI_Address(*void* location,MPI_Aint *address*);
MPI_ADDRESS	MPI_ADDRESS(*CHOICE LOCATION,INTEGER ADDRESS,INTEGER IERROR*)
MPI_Type_extent	int MPI_Type_extent(*MPI_Datatype datatype,int *extent*);
MPI_TYPE_EXTENT	MPI_TYPE_EXTENT(*INTEGER DATATYPE,INTEGER EXTENT,INTEGER IERROR*)
MPI_Type_size	int MPI_Type_size(*MPI_Datatype datatype,int *size*);
MPI_TYPE_SIZE	MPI_TYPE_SIZE(*INTEGER DATATYPE,INTEGER SIZE,INTEGER IERROR*)
MPI_Type_lb	int MPI_Type_lb(*MPI_Datatype datatype,int* displacement*);
MPI_TYPE_LB	MPI_TYPE_LB(*INTEGER DATATYPE,INTEGER DISPLACEMENT,INTEGER IERROR*)
MPI_Type_ub	int MPI_Type_ub(*MPI_Datatype datatype,int* displacement*);
MPI_TYPE_UB	MPI_TYPE_UB(*INTEGER DATATYPE,INTEGER DISPLACEMENT,INTEGER IERROR*)
MPI_Type_commit	int MPI_Type_commit(*MPI_Datatype *datatype*);
MPI_TYPE_COMMIT	MPI_TYPE_COMMIT(*INTEGER DATATYPE,INTEGER IERROR*)
MPI_Type_free	int MPI_Type_free(*MPI_Datatype *datatype*);
MPI_TYPE_FREE	MPI_TYPE_FREE(*INTEGER DATATYPE,INTEGER IERROR*)
MPI_Get_elements	int MPI_Get_elements(*MPI_Status *status,MPI_Datatype datatype,int *count*);
MPI_GET_ELEMENTS	MPI_GET_ELEMENTS(*INTEGER STATUS(MPI_STATUS_SIZE),INTEGER DATATYPE,INTEGER COUNT,INTEGER IERROR*)
MPI_Pack	int MPI_Pack(*void* inbuf,int incount,MPI_Datatype datatype,void *outbuf,int outsize,int *position,MPI_Comm comm*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_PACK	MPI_PACK(*CHOICE INBUF,INTEGER INCOUNT,INTEGER DATATYPE,CHOICE OUTBUF,INTEGER OUTSIZE,INTEGER POSITION,INTEGER COMM,INTEGER IERROR*)
MPI_Unpack	int MPI_Unpack(*void* inbuf,int insize,int *position,void *outbuf,int outcount,MPI_Datatype datatype,MPI_Comm comm*);
MPI_UNPACK	MPI_UNPACK(*CHOICE INBUF,INTEGER INSIZE,INTEGER POSITION,CHOICE OUTBUF,INTEGER OUTCOUNT,INTEGER DATATYPE,INTEGER COMM, INTEGER IERRROR*)
MPI_Pack_size	int MPI_Pack_size(*int incount,MPI_Datatype datatype,MPI_Comm comm,int *size*);
MPI_PACK_SIZE	MPI_PACK_SIZE(*INTEGER INCOUNT,INTEGER DATATYPE,INTEGER COMM,INTEGER SIZE,INTEGER IERROR*)

Bindings for Collective Communication

The following table lists the C and FORTRAN bindings for collective communication routines.

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Barrier	int MPI_Barrier(*MPI_Comm comm*);
MPI_BARRIER	MPI_BARRIER(*INTEGER COMM,INTEGER IERROR*)
MPI_Bcast	int MPI_Bcast(*void* buffer,int count,MPI_Datatype datatype,int root,MPI_Comm comm*);
MPI_BCAST	MPI_BCAST(*CHOICE BUFFER,INTEGER COUNT,INTEGER DATATYPE,INTEGER ROOT,INTEGER COMM,INTEGER IERROR*)
MPI_Gather	int MPI_Gather(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,MPI_Comm comm*);
MPI_GATHER	MPI_GATHER(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER IERROR*)
MPI_Gatherv	int MPI_Gatherv(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int *recvcounts,int *displs,MPI_Datatype recvtype,int root,MPI_Comm comm*);
MPI_GATHERV	MPI_GATHERV(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER IERROR*)
MPI_Scatter	int MPI_Scatter(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype,int root MPI_Comm comm*);
MPI_SCATTER	MPI_SCATTER(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER IERROR*)
MPI_Scatterv	int MPI_Scatterv(*void* sendbuf,int *sendcounts,int *displs,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype,int root,MPI_Comm comm*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_SCATTERV	MPI_SCATTERV(*CHOICE SENDBUF,INTEGER SENDCOUNTS(*),INTEGER DISPLS(*),INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER ROOT,INTEGER COMM,INTEGER IERROR*)
MPI_Allgather	int MPI_Allgather(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype, MPI_Comm comm*);
MPI_ALLGATHER	MPI_ALLGATHER(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,INTEGER IERROR*)
MPI_Allgatherv	int MPI_Allgatherv(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int *recvcounts,int *displs, MPI_Datatype recvtype,MPI_Comm comm*);
MPI_ALLGATHERV	MPI_ALLGATHERV(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DISPLS(*),INTEGER RECVTYPE,INTEGER COMM,INTEGER IERROR*)
MPI_Alltoall	int MPI_Alltoall(*void* sendbuf,int sendcount,MPI_Datatype sendtype,void* recvbuf,int recvcount,MPI_Datatype recvtype, MPI_Comm comm*);
MPI_ALLTOALL	MPI_ALLTOALL(*CHOICE SENDBUF,INTEGER SENDCOUNT,INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNT,INTEGER RECVTYPE,INTEGER COMM,INTEGER IERROR*)
MPI_Alltoallv	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	MPI_ALLTOALLV(*CHOICE SENDBUF,INTEGER SENDCOUNTS(*),INTEGER SDISPLS(*),INTEGER SENDTYPE,CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER RDISPLS(*),INTEGER RECVTYPE,INTEGER COMM,INTEGER IERROR*)
MPI_Reduce	int MPI_Reduce(*void* sendbuf,void* recvbuf,int count,MPI_Datatype datatype,MPI_Op op,int root,MPI_Comm comm*);
MPI_REDUCE	MPI_REDUCE(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER OP,INTEGER ROOT,INTEGER COMM,INTEGER IERROR*)
MPI_Op_create	int MPI_Op_create(*MPI_User_function *function, int commute, MPI_Op *op*);
MPI_OP_CREATE	MPI_OP_CREATE(*EXTERNAL FUNCTION,INTEGER COMMUTE,INTEGER OP,INTEGER IERROR*)
MPI_Op_free	int MPI_Op_free(*MPI_Op *op*);
MPI_OP_FREE	MPI_OP_FREE(*INTEGER OP,INTEGER IERROR*)
MPI_Allreduce	int MPI_Allreduce(*void* sendbuf,void* recvbuf,int count,MPI_Datatype datatype,MPI_Op op,MPI_Comm comm*);
MPI_ALLREDUCE	MPI_ALLREDUCE(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER IERROR*)
MPI_Reduce_scatter	int MPI_Reduce_scatter(*void* sendbuf,void* recvbuf,int *recvcounts,MPI_Datatype datatype,MPI_Op op,MPI_Comm comm*);
MPI_REDUCE_SCATTER	MPI_REDUCE_SCATTER(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER RECVCOUNTS(*),INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER IERROR*)
MPI_Scan	int MPI_Scan(*void* sendbuf,void* recvbuf,int count,MPI_Datatype datatype,MPI_Op op,MPI_Comm comm*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_SCAN	MPI_SCAN(*CHOICE SENDBUF,CHOICE RECVBUF,INTEGER COUNT,INTEGER DATATYPE,INTEGER OP,INTEGER COMM,INTEGER IERROR*)

Bindings for Groups and Communicators

The following table lists the C and FORTRAN bindings for group and communicator routines.

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Group_size	int MPI_Group_size(*MPI_Group group,int *size*);
MPI_GROUP_SIZE	MPI_GROUP_SIZE(*INTEGER GROUP,INTEGER SIZE,INTEGER IERROR*)
MPI_Group_rank	int MPI_Group_rank(*MPI_Group group,int *rank*);
MPI_GROUP_RANK	MPI_GROUP_RANK(*INTEGER GROUP,INTEGER RANK,INTEGER IERROR*)
MPI_Group_translate_ranks	int MPI_Group_translate_ranks (*MPI_Group group1,int n,int *ranks1,MPI_Group group2,int *ranks2*);
MPI_GROUP_TRANSLATE_RANKS	MPI_GROUP_TRANSLATE_RANKS(*INTEGER GROUP1, INTEGER N,INTEGER RANKS1(*),INTEGER GROUP2,INTEGER RANKS2(*),INTEGER IERROR*)
MPI_Group_compare	int MPI_Group_compare(*MPI_Group group1,MPI_Group group2,int *result*);
MPI_GROUP_COMPARE	MPI_GROUP_COMPARE(*INTEGER GROUP1,INTEGER GROUP2,INTEGER RESULT,INTEGER IERROR*)
MPI_Comm_group	int MPI_Comm_group(*MPI_Comm comm,MPI_Group *group*);
MPI_COMM_GROUP	MPI_COMM_GROUP(*INTEGER COMM,INTEGER GROUP,INTEGER IERROR*)
MPI_Group_union	int MPI_Group_union(*MPI_Group group1,MPI_Group group2,MPI_Group *newgroup*);
MPI_GROUP_UNION	MPI_GROUP_UNION(*INTEGER GROUP1,INTEGER GROUP2,INTEGER NEWGROUP,INTEGER IERROR*)
MPI_Group_intersection	int MPI_Group_intersection(*MPI_Group group1,MPI_Group group2,MPI_Group *newgroup*);
MPI_GROUP_INTERSECTION	MPI_GROUP_INTERSECTION(*INTEGER GROUP1,INTEGER GROUP2,INTEGER NEWGROUP,INTEGER IERROR*)
MPI_Group_difference	int MPI_Group_difference(*MPI_Group group1,MPI_Group group2,MPI_Group *newgroup*);
MPI_GROUP_DIFFERENCE	MPI_GROUP_DIFFERENCE(*INTEGER GROUP1,INTEGER GROUP2,INTEGER NEWGROUP,INTEGER IERROR*)
MPI_Group_incl	int MPI_Group_incl(*MPI_Group group,int n,int *ranks,MPI_Group *newgroup*);
MPI_GROUP_INCL	MPI_GROUP_INCL(*INTEGER GROUP,INTEGER N,INTEGER RANKS(*),INTEGER NEWGROUP,INTEGER IERROR*)
MPI_Group_excl	int MPI_Group_excl(*MPI_Group group,int n,int *ranks,MPI_Group *newgroup*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_GROUP_EXCL	MPI_GROUP_EXCL(*INTEGER GROUP,INTEGER N,INTEGER RANKS(*),INTEGER NEWGROUP,INTEGER IERROR*)
MPI_Group_range_incl	int MPI_Group_range_incl(*MPI_Group group,int n,int ranges[][3],MPI_Group *newgroup*);
MPI_GROUP_RANGE_INCL	MPI_GROUP_RANGE_INCL(*INTEGER GROUP,INTEGER N,INTEGER RANGES(3,*),INTEGER NEWGROUP,INTEGER IERROR*)
MPI_Group_range_excl	int MPI_Group_range_excl(*MPI_Group group,int n,int ranges [][3],MPI_Group *newgroup*);
MPI_GROUP_RANGE_EXCL	MPI_GROUP_RANGE_EXCL(*INTEGER GROUP,INTEGER N,INTEGER RANGES(3,*),INTEGER NEWGROUP,INTEGER IERROR*)
MPI_Group_free	int MPI_Group_free(*MPI_Group *group*);
MPI_GROUP_FREE	MPI_GROUP_FREE(*INTEGER GROUP,INTEGER IERROR*)
MPI_Comm_size	int MPI_Comm_size(*MPI_Comm comm,int *size*);
MPI_COMM_SIZE	MPI_COMM_SIZE(*INTEGER COMM,INTEGER SIZE,INTEGER IERROR*)
MPI_Comm_rank	int MPI_Comm_rank(*MPI_Comm comm,int *rank*);
MPI_COMM_RANK	MPI_COMM_RANK(*INTEGER COMM,INTEGER RANK,INTEGER IERROR*)
MPI_Comm_compare	int MPI_Comm_compare(*MPI_Comm comm1,MPI_Comm comm2,int *result*);
MPI_COMM_COMPARE	MPI_COMM_COMPARE(*INTEGER COMM1,INTEGER COMM2,INTEGER RESULT,INTEGER IERROR*)
MPI_Comm_dup	int MPI_Comm_dup(*MPI_Comm comm,MPI_Comm *newcomm*);
MPI_COMM_DUP	MPI_COMM_DUP(*INTEGER COMM,INTEGER NEWCOMM,INTEGER IERROR*)
MPI_Comm_create	int MPI_Comm_create(*MPI_Comm comm,MPI_Group group,MPI_Comm *newcomm*);
MPI_COMM_CREATE	MPI_COMM_CREATE(*INTEGER COMM,INTEGER GROUP,INTEGER NEWCOMM,INTEGER IERROR*)
MPI_Comm_split	int MPI_Comm_split(*MPI_Comm comm,int color,int key,MPI_Comm *newcomm*);
MPI_COMM_SPLIT	MPI_COMM_SPLIT(*INTEGER COMM,INTEGER COLOR,INTEGER KEY,INTEGER NEWCOMM,INTEGER IERROR*)
MPI_Comm_free	int MPI_Comm_free(*MPI_Comm *comm*);
MPI_COMM_FREE	MPI_COMM_FREE(*INTEGER COMM,INTEGER IERROR*)
MPI_Comm_test_inter	int MPI_Comm_test_inter(*MPI_Comm comm,int *flag*);
MPI_COMM_TEST_INTER	MPI_COMM_TEST_INTER(*INTEGER COMM,LOGICAL FLAG,INTEGER IERROR*)
MPI_Comm_remote_size	int MPI_Comm_remote_size(*MPI_Comm comm,int *size*);
MPI_COMM_REMOTE_SIZE	MPI_COMM_REMOTE_SIZE(*INTEGER COMM,INTEGER SIZE,INTEGER IERROR*)
MPI_Comm_remote_group	int MPI_Comm_remote_group(*MPI_Comm comm,MPI_Group *group*);
MPI_COMM_REMOTE_GROUP	MPI_COMM_REMOTE_GROUP(*INTEGER COMM,INTEGER GROUP,INTEGER IERROR*)

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Intercomm_create	int MPI_Intercomm_create(*MPI_Comm local_comm,int local_leader,MPI_Comm peer_comm,int remote_leader,int tag, MPI_Comm *newintercomm*);
MPI_INTERCOMM_CREATE	MPI_INTERCOMM_CREATE(*INTEGER LOCAL_COMM,INTEGER LOCAL_LEADER,INTEGER PEER_COMM,INTEGER REMOTE_LEADER,INTEGER TAG,INTEGER NEWINTERCOM,INTEGER IERROR*)
MPI_Intercomm_merge	int MPI_Intercomm_merge(*MPI_Comm intercomm,int high,MPI_Comm *newintracomm*);
MPI_INTERCOMM_MERGE	MPI_INTERCOMM_MERGE(*INTEGER INTERCOMM,INTEGER HIGH,INTEGER NEWINTRACOMM,INTEGER IERROR*)
MPI_Keyval_create	int MPI_Keyval_create(*MPI_Copy_function *copy_fn,MPI_Delete_function *delete_fn,int *keyval, void* extra_state*);
MPI_KEYVAL_CREATE	MPI_KEYVAL_CREATE(*EXTERNAL COPY_FN,EXTERNAL DELETE_FN,INTEGER KEYVAL,INTEGER EXTRA_STATE,INTEGER IERROR*)
MPI_Keyval_free	int MPI_Keyval_free(*int *keyval*);
MPI_KEYVAL_FREE	MPI_KEYVAL_FREE(*INTEGER KEYVAL,INTEGER IERROR*)
MPI_Attr_put	int MPI_Attr_put(*MPI_Comm comm,int keyval,void* attribute_val*);
MPI_ATTR_PUT	MPI_ATTR_PUT(*INTEGER COMM,INTEGER KEYVAL,INTEGER ATTRIBUTE_VAL,INTEGER IERROR*)
MPI_Attr_get	int MPI_Attr_get(*MPI_Comm comm,int keyval,void *attribute_val,int *flag*);
MPI_ATTR_GET	MPI_ATTR_GET(*INTEGER COMM,INTEGER KEYVAL,INTEGER ATTRIBUTE_VAL, LOGICAL FLAG,INTEGER IERROR*)
MPI_Attr_delete	int MPI_Attr_delete(*MPI_Comm comm,int keyval*);
MPI_ATTR_DELETE	MPI_ATTR_DELETE(*INTEGER COMM,INTEGER KEYVAL,INTEGER IERROR*)

Bindings for Topologies

The following table lists the C and FORTRAN bindings for topology routines.

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Cart_create	int MPI_Cart_create(*MPI_Comm comm_old,int ndims,int *dims,int *periods,int reorder,MPI_Comm *comm_cart*);
MPI_CART_CREATE	MPI_CART_CREATE(*INTEGER COMM_OLD,INTEGER NDIMS,INTEGER DIMS(*), INTEGER PERIODS(*),INTEGER REORDER,INTEGER COMM_CART,INTEGER IERROR*)
MPI_Dims_create	int MPI_Dims_create(*int nnodes,int ndims,int *dims*);
MPI_DIMS_CREATE	MPI_DIMS_CREATE(*INTEGER NNODES,INTEGER NDIMS,INTEGER DIMS(*), INTEGER IERROR*)
MPI_Graph_create	int MPI_Graph_create(*MPI_Comm comm_old,int nnodes,int *index,int *edges,int reorder,MPI_Comm *comm_graph*);
MPI_GRAPH_CREATE	MPI_GRAPH_CREATE(*INTEGER COMM_OLD,INTEGER NNODES,INTEGER INDEX(*), INTEGER EDGES(*),INTEGER REORDER,INTEGER COMM_GRAPH,INTEGER IERROR*)

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Topo_test	int MPI_Topo_test(*MPI_Comm comm,int *status*);
MPI_TOPO_TEST	MPI_TOPO_TEST(*INTEGER COMM,INTEGER STATUS,INTEGER IERROR*)
MPI_Graphdims_get	int MPI_Graphdims_get(*MPI_Comm comm,int *nnodes,int *nedges*);
MPI_GRAPHDIMS_GET	MPI_GRAPHDIMS_GET(*INTEGER COMM,INTEGER NNDODES,INTEGER NEDGES, INTEGER IERROR*)
MPI_Graph_get	int MPI_Graph_get(*MPI_Comm comm,int maxindex,int maxedges,int *index, int *edges*);
MPI_GRAPH_GET	MPI_GRAPH_GET(*INTEGER COMM,INTEGER MAXINDEX,INTEGER MAXEDGES,INTEGER INDEX(*),INTEGER EDGES(*),INTEGER IERROR*)
MPI_Cartdim_get	int MPI_Cartdim_get(*MPI_Comm comm, int *ndims*);
MPI_CARTDIM_GET	MPI_CARTDIM_GET(*INTEGER COMM,INTEGER NDIMS,INTEGER IERROR*)
MPI_Cart_get	int MPI_Cart_get(*MPI_Comm comm,int maxdims,int *dims,int *periods,int *coords*);
MPI_CART_GET	MPI_CART_GET(*INTEGER COMM,INTEGER MAXDIMS,INTEGER DIMS(*),INTEGER PERIODS(*),INTEGER COORDS(*),INTEGER IERROR*)
MPI_Cart_rank	int MPI_Cart_rank(*MPI_Comm comm,int *coords,int *rank*);
MPI_CART_RANK	MPI_CART_RANK(*INTEGER COMM,INTEGER COORDS(*),INTEGER RANK,INTEGER IERROR*)
MPI_Cart_coords	int MPI_Cart_coords(*MPI_Comm comm,int rank,int maxdims,int *coords*);
MPI_CART_COORDS	MPI_CART_COORDS(*INTEGER COMM,INTEGER RANK,INTEGER MAXDIMS,INTEGER COORDS(*),INTEGER IERROR*)
MPI_Graph_neighbors_count	int MPI_Graph_neighbors_count(*MPI_Comm comm,int rank,int *nneighbors*);
MPI_GRAPH_NEIGHBORS_COUNT	MPI_GRAPH_NEIGHBORS_COUNT(*INTEGER COMM,INTEGER RANK,INTEGER NEIGHBORS, INTEGER IERROR*)
MPI_Graph_neighbors	int MPI_Graph_neighbors(*MPI_Comm comm,int rank,int maxneighbors,int *neighbors*);
MPI_GRAPH_NEIGHBORS	MPI_GRAPH_NEIGHBORS(*MPI_COMM COMM,INTEGER RANK,INTEGER MAXNEIGHBORS,INTEGER NNEIGHBORS(*),INTEGER IERROR*)
MPI_Cart_shift	int MPI_Cart_shift(*MPI_Comm comm,int direction,int disp,int *rank_source,int *rank_dest*);
MPI_CART_SHIFT	MPI_CART_SHIFT(*INTEGER COMM,INTEGER DIRECTION,INTEGER DISP, INTEGER RANK_SOURCE,INTEGER RANK_DEST,INTEGER IERROR*)
MPI_Cart_sub	int MPI_Cart_sub(*MPI_Comm comm,int *remain_dims,MPI_Comm *newcomm*);
MPI_CART_SUB	MPI_CART_SUB(*INTEGER COMM,INTEGER REMAIN_DIMS,INTEGER NEWCOMM, INTEGER IERROR*)
MPI_Cart_map	int MPI_Cart_map(*MPI_Comm comm,int ndims,int *dims,int *periods,int *newrank*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_CART_MAP	MPI_CART_MAP(*INTEGER COMM,INTEGER NDIMS,INTEGER DIMS(*),INTEGER PERIODS(*),INTEGER NEWRANK,INTEGER IERROR*)
MPI_Graph_map	int MPI_Graph_map(*MPI_Comm comm,int nnodes,int *index,int *edges,int *newrank*);
MPI_GRAPH_MAP	MPI_GRAPH_MAP(*INTEGER COMM,INTEGER NNODES,INTEGER INDEX(*),INTEGER EDGES(*),INTEGER NEWRANK,INTEGER IERROR*)

Bindings for Environmental Management

The following table lists the C and FORTRAN bindings for environmental management routines.

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Get_version	int MPI_Get_version(*int *version,int *subversion*);
MPI_GET_VERSION	MPI_GET_VERSION(*INTEGER VERSION,INTEGER SUBVERSION,INTEGER IERROR*)
MPI_Get_processor_name	int MPI_Get_processor_name(*char *name,int *resultlen*);
MPI_GET_PROCESSOR_NAME	MPI_GET_PROCESSOR_NAME(*CHARACTER NAME(*),INTEGER RESULTLEN,INTEGER IERROR*)
MPI_Errhandler_create	int MPI_Errhandler_create(*MPI_Handler_function *function, MPI_Errhandler *errhandler*);
MPI_ERRHANDLER_CREATE	MPI_ERRHANDLER_CREATE(*EXTERNAL FUNCTION,INTEGER ERRHANDLER, INTEGER IERROR*)
MPI_Errhandler_set	int MPI_Errhandler_set(*MPI_Comm comm,MPI_Errhandler errhandler*);
MPI_ERRHANDLER_SET	MPI_ERRHANDLER_SET(*INTEGER COMM,INTEGER ERRHANDLER,INTEGER IERROR*)
MPI_Errhandler_get	int MPI_Errhandler_get(*MPI_Comm comm,MPI_Errhandler *errhandler*);
MPI_ERRHANDLER_GET	MPI_ERRHANDLER_GET(*INTEGER COMM,INTEGER ERRHANDLER,INTEGER IERROR*)
MPI_Errhandler_free	int MPI_Errhandler_free(*MPI_Errhandler *errhandler*);
MPI_ERRHANDLER_FREE	MPI_ERRHANDLER_FREE(*INTEGER ERRHANDLER,INTEGER IERROR*)
MPI_Error_string	int MPI_Error_string(*int errorcode, char *string, int *resultlen*);
MPI_ERROR_STRING	MPI_ERROR_STRING(*INTEGER ERRORCODE,CHARACTER STRING(*),INTEGER RESULTLEN,INTEGER IERROR*)
MPI_Error_class	int MPI_Error_class(*int errorcode, int *errorclass*);
MPI_ERROR_CLASS	MPI_ERROR_CLASS(*INTEGER ERRORCODE,INTEGER ERRORCLASS,INTEGER IERROR*)
MPI_Wtime	double MPI_Wtime(*void*);
MPI_WTIME	DOUBLE PRECISION MPI_WTIME()
MPI_Wtick	double MPI_Wtick(*void*);

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_WTICK	DOUBLE PRECISION MPI_WTICK()
MPI_Init	int MPI_Init(*int *argc, char ***argv*);
MPI_INIT	MPI_INIT(*INTEGER IERROR*)
MPI_Finalize	int MPI_Finalize(*void*);
MPI_FINALIZE	MPI_FINALIZE(*INTEGER IERROR*)
MPI_Initialized	int MPI_Initialized(*int *flag*);
MPI_INITIALIZED	MPI_INITIALIZED(*INTEGER FLAG,INTEGER IERROR*)
MPI_Abort	int MPI_Abort(*MPI_Comm comm, int errorcode*);
MPI_ABORT	MPI_ABORT(*INTEGER COMM,INTEGER ERRORCODE,INTEGER IERROR*)

Bindings for Profiling

The following table lists the C and FORTRAN bindings for profiling.

C/FORTRAN Subroutine	C/FORTRAN Binding
MPI_Pcontrol	int MPI_Pcontrol(*const int level,...*);
MPI_PCONTROL	MPI_PCONTROL(*INTEGER LEVEL,...*)

Appendix B. Profiling Message Passing

AIX Profiling

If you use **prof**, **gprof** or **xprofiler** and the appropriate compiler flag (-p or -pg), you can profile your program.

The message passing library is not enabled for **prof** or **gprof**, profiling counts. You can obtain profiling information by using the name-shifted MPI functions provided.

MPI Nameshift Profiling

To use nameshift profiling routines that are written to the C bindings with an MPI program written in C, or the FORTRAN bindings with an MPI program written in FORTRAN, do the following:

1. Create .o files for your profiling routines.
2. Use one of the following commands to list both the MPI program **.o** files and the profiling **.o** files as inputs:

   - **mpcc**
   - **mpxlf**
   - **mpcc_r**
   - **mpxlf_r**
   - **mpCC**
   - **mpCC_r**

   See IBM Parallel Environment for AIX: Operation and Use, Volume 1 for more information on these commands.
3. Run the resulting executable normally.

To use nameshift profiling routines which are written to the C bindings with an MPI program written in FORTRAN, follow these steps:

- If you are both the creator and user of the profiling library, follow all the steps (1 through 17).
- If you are the creator of the profiling library, follow steps 1 through 6. You also need to provide the user with the file created in step 2.
- If you are the user of the profiling library, follow steps 7 through 17. For step 14, use the file generated by the creator of the profiling library in step 2.

Based on the above, follow the appropriate steps:

1. Create a source file containing profiling versions of all the MPI routines you want to profile. As an example, create a source file called myprof.c containing the following code:

© Copyright IBM Corp. 1997

```
#include <stdio.h>
#include "mpi.h"
int MPI_Init(int *argc, char ***argv) {
  int rc;

  printf("hello from profiling layer MPI_Init...\n");
  rc = PMPI_Init(argc, argv);
  printf("goodbye from profiling layer MPI_Init...\n");
  return(rc);
}
```

2. Create an export file containing all of the symbols your profiling library will export. Begin this file with the name your profiling library will have and the name of the **.o** that will have the object code of your profiling routines. As an example, create a file called myprof.exp containing the following statements:

   ```
   #!libmyprof.a(newmyprof.o)
   MPI_Init
   ```

3. Create a file called **mpicore.imp**. This file will import all of the PMPI symbols that your profiling library needs. Begin this file with the statement #!libmpi.a(mpicore.o). The following is an example of mpicore.imp:

   ```
   #!libmpi.a(mpicore.o)
   PMPI_Init
   ```

4. Compile the source file containing your profiling MPI routines. For example:

   ```
   cc -c myprof.c -I/usr/lpp/ppe.poe/include
   ```

   The -I defines the location of mpi.h.

5. Create your profiling MPI library. Use the file created in step 2 as the export file and the file created in step 3 as the import file. Include any other libraries your profiling code needs, such as libc. For example:

   ```
   ld -o newmyprof.o myprof.o -bM:SRE -H512 -T512 -bnoentry
   -bI:mpicore.imp -bE:myprof.exp -lc
   ```

6. Archive the object module created in step 5 into a library. The library name should be the same as that listed in the first statement of the export file created in step 2. For example:

   ```
   ar rv libmyprof.a newmyprof.o
   ```

7. Use the following command to extract mpifort.o from libmpi.a:

   ```
   ar -xv /usr/lpp/ppe.poe/lib/libmpi.a mpifort.o
   ```

8. Use the following command to create a non-shared version of mpifort.o:

   ```
   ld -o mpifort.tmp mpifort.o -r -bnso -bnoentry
   ```

9. Use the following command to extract mpicore.o from libmpi.a:

   ```
   ar -xv /usr/lpp/ppe.poe/lib/libmpi.a mpicore.o
   ```

10. Use the following command to create an export list from the extracted mpicore.o:

    ```
    /usr/bin/dump -nvp mpicore.o  | /usr/bin/grep "¬\[" | cut -f2-
    | cut -c26- | grep -y "¬exp" | cut -c35- | sort | uniq > mpicore.exp
    ```

11. Delete all of the symbols selected for profiling in step 2 from mpicore.exp. Then create a new line at the top of the file. We'll call this the *new line 1* To the *new line 1*, add #!libmpi.a(mpicore.o). Continuing with our example:

MPI_Init would now be deleted from mpicore.exp. and #!libmpi.a(mpicore.o) would now comprise line 1 of mpicore.exp.

12. Create a file called **vt.exp** with the following statements:

    ```
    #!libvtd.a(dynamic.o)
    VT_instaddr_depth
    ```

13. Use the following command to create an export list from the extracted mpifort.o:

    ```
    /usr/bin/dump -nvp mpifort.o  | /usr/bin/grep "¬\[" | cut -f2-
    | cut -c26- | grep -y "¬exp" | cut -c35- | sort | uniq > mpifort.exp
    ```

    insert **#!libpmpi.a**(newmpifort.o) as the first line of the new mpifort.exp file

14. Create a new version of mpifort.o from the non-shared version you created in step 8. It will import the symbols representing your profiling functions from your profiling library using the file created in step 2. It will import the remaining MPI symbols from mpicore.o using the file created in step 11. One additional symbol must be imported using the file created in step 12. The new mpifort.o will export symbols using the file created in step 13.

    ```
    ld -o newmpifort.o mpifort.tmp -bI:mpicore.exp -bI:myprof.exp -bI:vt.exp
    -bE:mpifort.exp -bM:SRE -H512 -T512 -bnoentry
    ```

15. Use the following command to create a library containing a FORTRAN object which will reference your profiling library:

    ```
    ar rv libpmpi.a newmpifort.o
    ```

16. Create a program that uses an MPI function you've profiled. An example would be a file called hwinit.f that contains the following statements:

    ```
    c -------------------------------------
          program  hwinit
          include 'mpif.h'
          integer forterr
    c
          call MPI_INIT(forterr)
    c
    c  Write comments to screen.
    c
          write(6,*)'Hello from task '
    c
          call MPI_FINALIZE(forterr)
    c
          stop
          end
    c
    ```

17. Compile your program linking in the library created in step 15. For example:

    ```
    mpxlf -o hwinit hwinit.f -lpmpi -L.
    ```

Sample CPU MPI Time Program

The following is a sample MPI program that uses the name-shifted MPI interface to separate the amount of user and system CPU time used by MPI.

CPU MPI Time Example

```
#include "mpi.h"
#include <sys/types.h>
#include <time.h>
#include <sys/times.h>

#define ARRAY_SIZE 1000000
#define VALUE 123

struct tms mpitms;
double mpi_elapsed;

void main()
{
   int in[ARRAY_SIZE],out[ARRAY_SIZE],tasks,me,src,dest;
   int i;
   MPI_Status status[2];
   MPI_Request msgid[2];

   for (i=0;i<ARRAY_SIZE;i++)out[i]=VALUE;

   MPI_Init(&argc,&argv);
   MPI_Comm_size(MPI_COMM_WORLD,&tasks);
   MPI_Comm_rank(MPI_COMM_WORLD,&me);

   mpi_elapsed = MPI_Wtime();

   dest = (me==tasks-1) ? 0 : me+1;
   MPI_Isend(out,ARRAY_SIZE,MPI_INT,dest,5,MPI_COMM_WORLD,&msgid[0]);
   src = (me==0) ? tasks-1 : me-1;
   MPI_Irecv(in,ARRAY_SIZE,MPI_INT,src,5,MPI_COMM_WORLD,&msgid[1]);
   MPI_Waitall(2,msgid,status);

   for (i=0; i< ARRAY_SIZE; i++) {
   if(in[i] != VALUE )
      printf("ERROR on node %d, in = %d\n",me,in[i]);
          break;
          }

   MPI_Barrier(MPI_COMM_WORLD);

   mpi_elapsed = MPI_Wtime() - mpi_elapsed;

   printf("MPI CPU times: user %f, system %f, total %f sec\n",
                ((float)mpitms.tms_utime)/CLK_TCK,
                ((float)mpitms.tms_stime)/CLK_TCK,
                (float)(mpitms.tms_utime+mpitms.tms_stime)/CLK_TCK);
   printf("MPI Elapsed time: %f sec\n", mpi_elapsed);

   MPI_Finalize();
}

/*************************************/
/* Replacement functions for profiling */
/*************************************/

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

```
        struct tms beforetms, aftertms;
        int rc;
        times(&beforetms);

    rc = PMPI_Isend(buf,count,datatype,dest,tag,comm,request);

        times(&aftertms);
        mpitms.tms_utime += (aftertms.tms_utime - beforetms.tms_utime);
        mpitms.tms_stime += (aftertms.tms_stime - beforetms.tms_stime);
        return (rc);
        }


int MPI_Waitall(int count, MPI_Request *array_of_requests,
              MPI_Status *array_of_statuses)
{
        struct tms beforetms, aftertms;
        int rc;
        times(&beforetms);

    rc = PMPI_Waitall(count,array_of_requests,array_of_statuses);

        times(&aftertms);
        mpitms.tms_utime += (aftertms.tms_utime - beforetms.tms_utime);
        mpitms.tms_stime += (aftertms.tms_stime - beforetms.tms_stime);
        return (rc);
        }



int MPI_Irecv(void* buf, int count, MPI_Datatype datatype,
       int source, int tag, MPI_Comm comm, MPI_Request *request)
{
        struct tms beforetms, aftertms;
        int rc;
        times(&beforetms);

    rc = PMPI_Irecv(buf,count,datatype,source,tag,comm,request);

        times(&aftertms);
        mpitms.tms_utime += (aftertms.tms_utime - beforetms.tms_utime);
        mpitms.tms_stime += (aftertms.tms_stime - beforetms.tms_stime);
        return (rc);
        }


int MPI_Barrier(MPI_Comm comm )
{
        struct tms beforetms, aftertms;
        int rc;
        times(&beforetms);

    rc = PMPI_Barrier(comm);

        times(&aftertms);
        mpitms.tms_utime += (aftertms.tms_utime - beforetms.tms_utime);
        mpitms.tms_stime += (aftertms.tms_stime - beforetms.tms_stime);
        return (rc);
```

```
}
```

Appendix C. MPI Size Limits

MPI Tunables and Limits

The following is a list of MPI size limits. This list includes system limits on the size of various MPI elements and the relevant environment variable or tunable parameter.

- Number of tasks: MP_PROCS
- Maximum number of tasks: 128
- Maximum message size for Point-to-Point communication: No specific limit
- Default receive buffer size: (MP_BUFFER_MEM)

  When using Internet Protocol (IP), 2,800,000 bytes

  When using User Space (US), 64MB
- Default eager limit: See table below.
- Maximum eager limit: 64K bytes
- To ensure that at least 32 messages can be outstanding between any 2 tasks, MPI_EAGER_LIMIT is adjusted according to the following table (and when MP__EAGER_LIMIT and MP_BUFFER_MEM have not been set by the user):

Number of Tasks	MP_EAGER_LIMIT
1 - 16	4096
17 - 32	2048
33 - 64	1024
65 - 128	512

- The maximum number of outstanding unmatched send requests (smaller than MP_EAGER_LIMIT) per node for any single destination node is given by the formula:

  ```
  (0.75*MP_BUFFER_MEM)/(MP_PROCS*(max(MP_EAGER_LIMIT,64)))
  ```
- Maximum aggregate unsent data, per task: No specific limit
- Maximum number of communicators: approximately 2000
- Maximum number of data types: Depends on MP_BUFFER_MEM
- Maximum data type depth: Default is 5 (MP_MAX_TYPEDEPTH)
- Maximum number of distinct tags: All non-negative integers less than 2**32-1

© Copyright IBM Corp. 1997

Appendix D. Reduction Operations

Predefined Reduction Operations

The following is a list of the predefined operations for use with MPI_REDUCE, MPI_ALLREDUCE, MPI_REDUCE_SCATTER and MPI_SCAN. To invoke a predefined operation, place any of the following in **op**.

Reduction Operation	Description
MPI_MAX	maximum
MPI_MIN	minimum
MPI_SUM	sum
MPI_PROD	product
MPI_LAND	logical AND
MPI_BAND	bitwise AND
MPI_LOR	logical OR
MPI_BOR	bitwise OR
MPI_LXOR	logical XOR
MPI_BXOR	bitwise XOR
MPI_MAXLOC	max value and location
MPI_MINLOC	min value and location

Reduction Operations - Valid Datatype Arguments Operations

The reduction operations have the following basic **datatype** arguments.

Type	Valid Datatype Arguments
C integer	MPI_INT, MPI_LONG, MPI_SHORT, MPI_UNSIGNED_SHORT, MPI_UNSIGNED, MPI_UNSIGNED_LONG
FORTRAN integer	MPI_INTEGER
Floating point	MPI_FLOAT, MPI_DOUBLE, MPI_REAL, MPI_DOUBLE_PRECISION, MPI_LONG_DOUBLE
Logical	MPI_LOGICAL
Complex	MPI_COMPLEX
Byte	MPI_BYTE
C Pair	MPI_FLOAT_INT, MPI_DOUBLE_INT, MPI_LONG_INT, MPI_2INT, MPI_SHORT_INT, MPI_LONG_DOUBLE_INT
FORTRAN Pair	MPI_2REAL, MPI_2DOUBLE_PRECISION, MPI_2INTEGER

© Copyright IBM Corp. 1997

op Option - Valid Datatypes

The following are the valid datatypes for each **op** option.

Type	Valid Datatypes For op Option
C integer	MPI_MAX, MPI_MIN, MPI_SUM, MPI_PROD, MPI_LAND, MPI_LOR, MPI_LXOR, MPI_BAND, MPI_BOR, MPI_BXOR
FORTRAN integer	MPI_MAX, MPI_MIN, MPI_SUM, MPI_PROD, MPI_BAND, MPI_BOR, MPI_BXOR
Floating point	MPI_MAX, MPI_MIN, MPI_SUM, MPI_PROD
Logical	MPI_LAND, MPI_LOR, MPI_LXOR
Complex	MPI_SUM, MPI_PROD
Byte	MPI_BAND, MPI_BOR, MPI_BXOR
C Pair	MPI_MINLOC, MPI_MAXLOC
FORTRAN Pair	MPI_MINLOC, MPI_MAXLOC

Examples

Examples of user-defined reduction functions for integer vector addition.

C Example

```
void int_sum (int *in, int *inout,
     int *len, MPI_Datatype *type);

{
   int i
   for (i=0; i<*len; i++)  {
      inout[i] + = in[i];
   }
}
```

FORTRAN Example

```
      SUBROUTINE INT_SUM(IN,INOUT,LEN,TYPE)

      INTEGER IN(*),INOUT(*),LEN,TYPE,I

      DO I = 1,LEN
         INOUT(I) = IN(I) + INOUT(I)
      ENDDO
      END
```

User-supplied reduction operations have four arguments:

- The first argument, **in**, is an array or scalar variable. The length, in elements, is specified by the third argument, **len**.

  This argument is an input array to be reduced.

- The second argument, **inout**, is an array or scalar variable. The length, in elements, is specified by the third argument, **len**.

  This argument is an input array to be reduced and the result of the reduction will be placed here.

- The third argument, **len** is the number of elements in **in** and **inout** to be reduced.

- The fourth argument **type** is the datatype of the elements to be reduced.

Users may code their own reduction operations, with the restriction that the operations must be associative. Also, C programmers should note that the values of **len** and **type** will be passed as pointers. No communication calls are allowed in user-defined reduction operations. See "Limitations In Setting The Thread Stacksize" on page 341 in Appendix G, "Programming Considerations for User Applications in POE" on page 329 for thread stacksize considerations when using the threaded MPI library.

Appendix E. Parallel Utility Functions

This chapter contains the man pages for the Parallel Utility functions. These user-callable, thread-safe functions exploit features of the IBM Parallel Environment for AIX. Included are functions for:

- updating the Program Marker Array lights
- controlling distribution of STDIN and STDOUT
- synchronizing parallel tasks without using the message passing library
- improving control of interrupt driven programs.

There is a C and FORTRAN version for most of the functions.

The Parallel Utility functions are:

MP_DISABLEINTR and mpc_disableintr disables packet arrival interrupts on the node on which it is executed.

MP_ENABLEINTR and mpc_enableintr enables interrupts on the node on which it is executed.

MP_FLUSH and mpc_flush flushes output buffers to STDOUT. This is a synchronizing call across all parallel tasks.

MP_MARKER and mpc_marker requests that the numeric and text data passed in the call be forwarded to the Program Marker Array for display.

MP_NLIGHTS and mpc_nlights returns the number of Program Marker Array lights defined for this session.

MP_QUERYINTR and mpc_queryintr returns the state of interrupts on a node.

MP_QUERYINTRDELAY and mpc_queryintrdelay returns, in microseconds, the current interrupt delay time.

MP_SETINTRDELAY and mpc_setintrdelay sets the delay parameter to the specified value in **val**. This call can be made multiple times in a program with different values being passed to it each time.

MP_STDOUT_MODE and mpc_stdout_mode requests that STDOUT be set to single, ordered, or unordered mode. In single mode, only one task output is displayed. In unordered mode, output is displayed in the order received at the home node. In ordered mode, each parallel task writes output data to its own buffer; when a flush request is made all the task buffers are flushed, in order of the task id, to home node's STDOUT.

MP_STDOUTMODE_QUERY and mpc_stdoutmode_query returns the mode to which STDOUT is currently set.

mpc_isatty determines if a device is a terminal on the home node.

© Copyright IBM Corp. 1997

For more information on the Program Marker Array, or on controlling STDIN and STDOUT using POE, refer to *IBM Parallel Environment for AIX: Operation and Use, Volume 1*.

MP_DISABLEINTR, mpc_disableintr

MP_DISABLEINTR, mpc_disableintr- Disables message arrival interrupts on the node on which it is executed.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_disableintr();
```

FORTRAN SYNOPSIS

```
MP_DISABLEINTR(INTEGER RC)
```

PARAMETERS

In FORTRAN, **rc** contains the values as described below in RETURN VALUES.

DESCRIPTION

This Parallel Utility function disables interrupts on the individual node it is executed. Use this function to dynamically control masking interrupts on a node.

NOTES

- This function overrides the setting of the environment variable MP_CSS_INTERRUPT.
- Inappropriate use of the interrupt control functions may reduce performance.
- This function can be used for IP and US protocols.
- This function is thread safe.

RETURN VALUES

0 indicates successful completion

-1 indicates that an error occurred. A message describing the error will be issued.

EXAMPLES

C Example

```
/*
 * Running this program, after compiling with mpcc,
 * without setting the MP_CSS_INTERRUPT environment variable,
 * and without using the "-css_interrupt" command-line option,
 * produces the following output:
 *
 *    Interrupts are DISABLED
 *    About to enable interrupts..
```

```
 *    Interrupts are ENABLED
 *    About to disable interrupts...
 *    Interrupts are DISABLED
 */

#include "pm_util.h"

#define QUERY if (intr = mpc_queryintr()) {\
   printf("Interrupts are ENABLED\n");\
  } else {\
   printf("Interrupts are DISABLED\n");\
  }

main()
{
 int intr;

 QUERY

 printf("About to enable interrupts...\n");
 mpc_enableintr();

 QUERY

 printf("About to disable interrupts...\n");
 mpc_disableintr();

 QUERY
}
```

FORTRAN Example

Running the following program, after compiling with mpxlf, without setting the MP_CSS_INTERRUPT environment variable, and without using the "-css_interrupt" command-line option, produces the following output:

```
      Interrupts are DISABLED
      About to enable interrupts..
      Interrupts are ENABLED
      About to disable interrupts...
      Interrupts are DISABLED



      PROGRAM INTR_EXAMPLE

      INTEGER RC

      CALL MP_QUERYINTR(RC)
      IF (RC .EQ. 0) THEN
         WRITE(6,*)'Interrupts are DISABLED'
      ELSE
         WRITE(6,*)'Interrupts are ENABLED'
      ENDIF
```

```
WRITE(6,*)'About to enable interrupts...'
CALL MP_ENABLEINTR(RC)

CALL MP_QUERYINTR(RC)
IF (RC .EQ. 0) THEN
   WRITE(6,*)'Interrupts are DISABLED'
ELSE
   WRITE(6,*)'Interrupts are ENABLED'
ENDIF

WRITE(6,*)'About to disable interrupts...'
CALL MP_DISABLEINTR(RC)

CALL MP_QUERYINTR(RC)
IF (RC .EQ. 0) THEN
   WRITE(6,*)'Interrupts are DISABLED'
ELSE
   WRITE(6,*)'Interrupts are ENABLED'
ENDIF

STOP
END
```

RELATED INFORMATION

Functions: MP_QUERYINTR(3), mpc_queryintr(3), MP_ENABLEINTR(3), mpc_enableintr(3), MP_SETINTRDELAY(3), mpc_setintrdelay(3), MP_QUERYINTRDELAY(3), mpc_queryintrdelay(3)

NAME

MP_ENABLEINTR, mpc_enableintr-Enables message arrival interrupts on the node on which it is executed.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_enableintr();
```

FORTRAN SYNOPSIS

```
MP_ENABLEINTR(INTEGER RC)
```

PARAMETERS

In FORTRAN, **rc** contains the values as described below in RETURN VALUES.

DESCRIPTION

This Parallel Utility function enables interrupts on the individual node on which it is executed. Use this function to dynamically control masking interrupts on a node.

NOTES

- This function overrides the setting of the environment variable MP_CSS_INTERRUPT.
- Inappropriate use of the interrupt control functions may reduce performance.
- This function can be used for IP and US protocols.
- This function is thread safe.

RETURN VALUES

0 indicates successful completion

-1 indicates that an error occurred. A message describing the error will be issued.

EXAMPLES

C Example

```
/*
 * Running this program, after compiling with mpcc,
 * without setting the MP_CSS_INTERRUPT environment variable,
 * and without using the "-css_interrupt" command-line option,
 * produces the following output:
 *
 *    Interrupts are DISABLED
 *    About to enable interrupts..
```

```
 *    Interrupts are ENABLED
 *    About to disable interrupts...
 *    Interrupts are DISABLED
 */

#include "pm_util.h"

#define QUERY if (intr = mpc_queryintr()) {\
   printf("Interrupts are ENABLED\n");\
  } else {\
   printf("Interrupts are DISABLED\n");\
  }

main()
{
 int intr;

 QUERY

 printf("About to enable interrupts...\n");
 mpc_enableintr();

 QUERY

 printf("About to disable interrupts...\n");
 mpc_disableintr();

 QUERY
}
```

FORTRAN Example

Running this program, after compiling with mpxlf, without setting the MP_CSS_INTERRUPT environment variable, and without using the "-css_interrupt" command-line option, produces the following output:

```
      Interrupts are DISABLED
      About to enable interrupts..
      Interrupts are ENABLED
      About to disable interrupts...
      Interrupts are DISABLED



      PROGRAM INTR_EXAMPLE

      INTEGER RC

      CALL MP_QUERYINTR(RC)
      IF (RC .EQ. 0) THEN
         WRITE(6,*)'Interrupts are DISABLED'
      ELSE
         WRITE(6,*)'Interrupts are ENABLED'
      ENDIF
```

```
WRITE(6,*)'About to enable interrupts...'
CALL MP_ENABLEINTR(RC)

CALL MP_QUERYINTR(RC)
IF (RC .EQ. 0) THEN
   WRITE(6,*)'Interrupts are DISABLED'
ELSE
   WRITE(6,*)'Interrupts are ENABLED'
ENDIF

WRITE(6,*)'About to disable interrupts...'
CALL MP_DISABLEINTR(RC)

CALL MP_QUERYINTR(RC)
IF (RC .EQ. 0) THEN
   WRITE(6,*)'Interrupts are DISABLED'
ELSE
   WRITE(6,*)'Interrupts are ENABLED'
ENDIF

STOP
END
```

RELATED INFORMATION

Functions: MP_QUERYINTR(3), mpc_queryintr(3), MP_DISABLEINTR(3), mpc_disableintr(3), MP_SETINTRDELAY(3), mpc_setintrdelay(3), MP_QUERYINTRDELAY(3), mpc_queryintrdelay(3)

NAME

MP_FLUSH, mpc_flush-Flushes task output buffers to STDOUT.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_flush(int option);
```

FORTRAN SYNOPSIS

```
MP_FLUSH(INTEGER OPTION)
```

PARAMETERS

option is an AIX file descriptor. The only valid value is:

1 to flush STDOUT buffers.

DESCRIPTION

This Parallel Utility function flushes output buffers from all the parallel tasks to STDOUT at the home node. This is a synchronizing call across all parallel tasks.

If the current STDOUT mode is ordered, then when all tasks have issued this call or when any of the output buffers are full:

1. all STDOUT buffers are flushed and put out to the user screen (or redirected) in task order.
2. an acknowledgement is sent to all tasks and control is returned to the user.

If current STDOUT mode is unordered and all tasks have issued this call, all output buffers are flushed and put out to the user screen (or redirected).

If the current STDOUT mode is single and all tasks have issued this call, the output buffer for the current single task is flushed and put out to the user screen (or redirected).

NOTES

- This is a synchronizing call regardless of the current STDOUT mode.
- All STDOUT buffers are flushed at the end of the parallel job.
- If mpc_flush is not used, standard output streams not terminated with a new-line character are buffered, even if a subsequent read to standard input is made. This may cause prompt message to appear only after input has been read.
- This function is thread safe.

RETURN VALUES

In C and C++ calls, the following applies:

0 indicates successful completion

-1 indicates that an error occurred. A message describing the error will be issued.

EXAMPLES

C Example

The following program uses **poe** with the **-labelio yes** option and three tasks:

```
    #include <pm_util.h>

main()
{
 mpc_stdout_mode(STDIO_ORDERED);
 printf("These lines will appear in task order\n");
 /*
  * Call mpc_flush here to make sure that one task
  * doesn't change the mode before all tasks have
  * sent the previous printf string to the home node.
  */
 mpc_flush(1);
 mpc_stdout_mode(STDIO_UNORDERED);
 printf("These lines will appear in the order received by the home node\n");
 /*
  * Since synchronization is not used here, one task could actually
  * execute the next statement before one of the other tasks has
  * executed the previous statement, causing one of the unordered
  * lines not to print.
  */
 mpc_stdout_mode(1);
 printf("Only 1 copy of this line will appear from task 1\n");
}
```

Running the above C program produces the following output (task order of lines 4-6 may differ):

```
0 : These lines will appear in task order.
1 : These lines will appear in task order.
2 : These lines will appear in task order.
1 : These lines will appear in the order received by the home node.
2 : These lines will appear in the order received by the home node.
0 : These lines will appear in the order received by the home node.
1 : Only 1 copy of this line will appear from task 1.
```

FORTRAN Example

```
 CALL MP_STDOUT_MODE(-2)
 WRITE(6, *) 'These lines will appear in task order'
 CALL MP_FLUSH(1)
 CALL MP_STDOUT_MODE(-3)
 WRITE(6, *) 'These lines will appear in the order received by the
xhome node'
 CALL MP_STDOUT_MODE(1)
 WRITE(6, *) 'Only 1 copy of this line will appear from task 1'
 END
```

RELATED INFORMATION

Functions: mpc_stdout_mode(3), MP_STDOUT_MODE(3), mpc_stdoutmode_query(3), MP_STDOUTMODE_QUERY(3)

NAME

MP_MARKER, mpc_marker-Passes numeric and text data to the Program Marker Array for display.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <mp_marker.h>
void mpc_marker(int light, int color, char *string);
```

FORTRAN SYNOPSIS

```
MP_MARKER(INTEGER LIGHT, INTEGER COLOR,
          CHARACTER STRING)
```

PARAMETERS

light is the light number to be colored. The lights in each task row are numbered, left to right, from 0 to one less than the number of lights. The row on which the light is colored is that of the calling task.

If the value of light is out of range, the parameter is ignored. No light is colored by the subroutine. Setting the light to a negative number allows you to update the string only.

color is the color you want to make the light. Supported values range from 0 to 102. The range 0 to 99 is, roughly, a spectrum starting with black and going through brown, green, blue, purple, red, orange, yellow, and ending with white. 100, 101, and 102 are three shades of gray growing increasingly dark. These are approximations, as the actual colors used are requested from the default X-Windows color map. If the X-Server is not capable of providing the colors in the RGB intensities requested, it colors the light with a close approximation.

If the value of color is out of range, the parameter is ignored. The subroutine does not give the light a new color.

string is the output string to be passed to the PM Array. In C programs the string must be null-terminated. This is not necessary in FORTRAN programs. The string can be any length, although only the first 80 characters can display. A null string can be passed.

DESCRIPTION

This Parallel Utility function requests that the numeric and text data passed in the call be forwarded to the Program Marker Array for display. This call waits for a specific acknowledgement from the Partition Manager if the number of lights (specified by the **MP_PMLIGHTS** environment variable) is positive. The program returns only after the message has been acknowledged by the POE home node. Hence,

this call will slow down the user's application and approximately synchronize it with the Program Marker Array.

If **MP_PMLIGHTS** is set to *0*, no message is sent.

NOTES

- Creates a Visualization Tool(VT) trace record (marker event).
- Sends a message only if PMLIGHTS is greater than 0. Each call waits for an acknowledgement from the home node.
- The Program Marker Array X-Windows display routine is distributed as a sample program.
- This function is thread safe.

EXAMPLES

C Example

The C statement:

```
   #include <mp_marker.h>
mpc_marker(2, 0, "Starting task 2");
```

gives the third light of the calling task the color black and after the lights on that task's row, prints the string "Starting task 2".

FORTRAN Example

The FORTRAN statements:

```
CALL MP_MARKER(2, 0, 'Starting task 2')
```

gives the third light of the calling task the color black and after the lights on that task's row, prints the string 'Starting task 2'.

RELATED INFORMATION

Functions: mpc_nlights(3), MP_NLIGHTS(3)

Commands: pmarray(1)

NAME

MP_NLIGHTS, mpc_nlights-Gets the number of Program Marker Array lights defined for this session.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <mp_marker.h>
int mpc_nlights();
```

FORTRAN SYNOPSIS

```
MP_NLIGHTS(INTEGER NLIGHT)
```

PARAMETERS

In FORTRAN, the number of Program Marker Array lights defined for this session is returned in the variable **NLIGHT**.

DESCRIPTION

This Parallel Utility function returns the number of Program Marker Array lights defined for this session.

NOTES

- You can set the number of Program Marker Array lights by using the environment variable **MP_PMLIGHTS** or the command line option **-pmlights**.
- This function is thread safe.

RETURN VALUES

In C, the current number of the Program Marker Array lights is the return value.

EXAMPLES

C Example

The following program uses poe with the **-pmlights 3** option and 1 task:

```
   #include <mp_marker.h>

main()
{
 printf("Number of Program Marker Array lights for this session is %d\n",
 mpc_nlights());
}
```

Running the above C program produces the following output:

```
Number of Program Marker Array lights for this session is 3.
```

FORTRAN Example

The following program uses poe with the **-pmlights 3** option and 1 task:

```
      INTEGER LIGHTS

      CALL MP_NLIGHTS(LIGHTS)
      WRITE(6, *) 'Number of Program Marker Array lights for this
     xsession is ', LIGHTS
      END
```

Running the above produces the following output:

```
Number of Program Marker Array lights for this session is 3.
```

RELATED INFORMATION

Functions: mpc_marker(3), MP_MARKER(3)

Commands: pmarray(1)

NAME

MP_QUERYINTR, mpc_queryintr-Returns the state of interrupts on a node.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_queryintr();
```

FORTRAN SYNOPSIS

```
MP_QUERYINTR(INTEGER RC)
```

PARAMETERS

In FORTRAN, **rc** contains the values as described below in RETURN VALUES.

DESCRIPTION

This Parallel Utility function returns the state of interrupts on a node.

NOTES

- This function is thread safe.

RETURN VALUES

0 indicates that interrupts are disabled on the node from which this function is called.

1 indicates that interrupts are enabled on the node from which this function is called.

EXAMPLES

C Example

```
/*
 * Running this program, after compiling with mpcc,
 * without setting the MP_CSS_INTERRUPT environment variable,
 * and without using the "-css_interrupt" command-line option,
 * produces the following output:
 *
 *    Interrupts are DISABLED
 *    About to enable interrupts..
 *    Interrupts are ENABLED
 *    About to disable interrupts...
 *    Interrupts are DISABLED
 */

#include "pm_util.h"
```

```
#define QUERY if (intr = mpc_queryintr()) {\
   printf("Interrupts are ENABLED\n");\
  } else {\
   printf("Interrupts are DISABLED\n");\
  }

main()
{
 int intr;

 QUERY

 printf("About to enable interrupts...\n");
 mpc_enableintr();

 QUERY

 printf("About to disable interrupts...\n");
 mpc_disableintr();

 QUERY
}
```

FORTRAN Example

Running this program, after compiling with mpxlf, without setting the MP_CSS_INTERRUPT environment variable, and without using the "-css_interrupt" command-line option, produces the following output:

```
      Interrupts are DISABLED
      About to enable interrupts..
      Interrupts are ENABLED
      About to disable interrupts...
      Interrupts are DISABLED
```

```
      PROGRAM INTR_EXAMPLE

      INTEGER RC

      CALL MP_QUERYINTR(RC)
      IF (RC .EQ. 0) THEN
         WRITE(6,*)'Interrupts are DISABLED'
      ELSE
         WRITE(6,*)'Interrupts are ENABLED'
      ENDIF

      WRITE(6,*)'About to enable interrupts...'
      CALL MP_ENABLEINTR(RC)

      CALL MP_QUERYINTR(RC)
      IF (RC .EQ. 0) THEN
         WRITE(6,*)'Interrupts are DISABLED'
      ELSE
```

```
   WRITE(6,*)'Interrupts are ENABLED'
ENDIF

WRITE(6,*)'About to disable interrupts...'
CALL MP_DISABLEINTR(RC)

CALL MP_QUERYINTR(RC)
IF (RC .EQ. 0) THEN
   WRITE(6,*)'Interrupts are DISABLED'
ELSE
   WRITE(6,*)'Interrupts are ENABLED'
ENDIF

STOP
END
```

RELATED INFORMATION

Functions: MP_DISABLEINTR(3), mpc_disableintr(3), MP_ENABLEINTR(3), mpc_enableintr(3), MP_SETINTRDELAY(3), mpc_setintrdelay(3), MP_QUERYINTRDELAY(3), mpc_queryintrdelay(3)

NAME

MP_QUERYINTRDELAY, mpc_queryintrdelay-Returns, in microseconds, the current interrupt delay time.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_queryintrdelay();
```

FORTRAN SYNOPSIS

```
MP_QUERYINTRDELAY(INTEGER RC)
```

PARAMETERS

In FORTRAN, **rc** contains the values as described below in RETURN VALUES.

DESCRIPTION

This Parallel Utility function returns, in microseconds, the current interrupt delay time.

NOTES

- The default interrupt delay time is 35 microseconds for TB2 and 1 microsecond for TB3.
- This function is thread safe.

RETURN VALUES

The current interrupt delay time in microseconds.

EXAMPLES

C Example

```
/*
 * Running this program, after compiling with mpcc,
 * without setting the MP_INTRDELAY environment variable,
 * and without using the "-intrdelay" command-line option,
 * produces the following output:
 *
 *    Current interrupt delay time is 35
 *    About to set interrupt delay time to 100...
 *    Current interrupt delay time is 100
 */

#include "pm_util.h"
```

```
main()
{
 printf("Current interrupt delay time is %d\n", mpc_queryintrdelay());

 printf("About to set interrupt delay time to 100...\n");
 mpc_setintrdelay(100);

 printf("Current interrupt delay time is %d\n", mpc_queryintrdelay());
}
```

FORTRAN Example

Running this program, after compiling with mpxlf, without setting the MP_INTRDELAY environment variable, and without using the "-intrdelay" command-line option, produces the following output:

```
Current interrupt delay time is 35
About to set interrupt delay time to 100...
Current interrupt delay time is 100
```

```
PROGRAM INTRDELAY_EXAMPLE

INTEGER DELAY, RC

CALL  MP_QUERYINTRDEALY(DELAY)
WRITE(6,*)'Current interrupt delay time is', delay

WRITE(6,*)'About to set interrupt delay time to 100...'
DELAY = 100
CALL MP_SETINTRDELAY(DELAY, RC)

CALL  MP_QUERYINTRDELAY(DELAY)
WRITE(6,*)'Current interrupt delay time is', delay

STOP
END
```

RELATED INFORMATION

Functions: MP_QUERYINTR(3), mpc_queryintr(3), MP_DISABLEINTR(3), mpc_disableintr(3), MP_ENABLEINTR(3), mpc_enableintr(3), MP_SETINTRDELAY(3), mpc_setintrdelay(3)

NAME

MP_SETINTRDELAY, mpc_setintrdelay-Sets the delay parameter to the specified value in **val**.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_setintrdelay(integer val);
```

FORTRAN SYNOPSIS

```
MP_SETINTRDELAY(INTEGER VAL, INTEGER RC)
```

PARAMETERS

val is the delay parameter in microseconds

rc in FORTRAN, **rc** contains the values as described below in RETURN VALUES.

DESCRIPTION

This Parallel Utility function sets the delay parameter to the specified value in **val**. This call can be made multiple times in a program with different values being passed to it each time.

You can use the environment variable MP_INTERDELAY to set an integer value before running your program. In this way, you can tune your delay parameter without having to recompile existing applications.

It should be noted that the cost of servicing an interrupt is quite high. For an application with few nodes exchanging small messages, it will help latency if the interrupt delay is kept small. For an application with a large number of nodes or one which exchanges large messages, keeping the delay parameter large will help the bandwidth. This allows multiple read transmissions to occur in a single read cycle. You should experiment with different values functions to achieve the desired performance depending on the communication pattern.

NOTES

- The default interrupt delay time is 35 microseconds for TB2 and 1 microsecond for TB3.
- Overrides the setting of the environment variable MP_INTERDELAY.
- This function is thread safe.

RETURN VALUES

0 indicates successful completion

-1 indicates that an error occurred. A message describing the error will be issued.

EXAMPLES

C Example

```
/*
 * Running this program, after compiling with mpcc,
 * without setting the MP_INTRDELAY environment variable,
 * and without using the "-intrdelay" command-line option,
 * produces the following output:
 *
 *    Current interrupt delay time is 35
 *    About to set interrupt delay time to 100...
 *    Current interrupt delay time is 100
 */

#include "pm_util.h"

main()
{
 printf("Current interrupt delay time is %d\n", mpc_queryintrdelay());

 printf("About to set interrupt delay time to 100...\n");
 mpc_setintrdelay(100);

 printf("Current interrupt delay time is %d\n", mpc_queryintrdelay());
}
```

FORTRAN Example

Running this program, after compiling with mpxlf, without setting the MP_INTRDELAY environment variable, and without using the "-intrdelay" command-line option, produces the following output:

```
      Current interrupt delay time is 35
      About to set interrupt delay time to 100...
      Current interrupt delay time is 100



      PROGRAM INTRDELAY_EXAMPLE

      INTEGER DELAY, RC

      CALL  MP_QUERYINTRDELAY(DELAY)
      WRITE(6,*)'Current interrupt delay time is', delay

      WRITE(6,*)'About to set interrupt delay time to 100...'
      DELAY = 100
```

```
CALL MP_SETINTRDELAY(DELAY, RC)

CALL  MP_QUERYINTRDELAY(DELAY)
WRITE(6,*)'Current interrupt delay time is', delay

STOP
END
```

RELATED INFORMATION

Functions: MP_QUERYINTR(3), mpc_queryintr(3), MP_DISABLEINTR(3), mpc_disableintr(3), MP_ENABLEINTR(3), mpc_enableintr(3), MP_QUERYINTRDELAY(3), mpc_queryintrdelay(3) MP_ENABLEINTR(3), mpc_enableintr(3), MP_DISABLEINTR(3), mpc_disableintr(3)

NAME

MP_STDOUT_MODE, mpc_stdout_mode-Requests that STDOUT be set to single, ordered, or unordered mode.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_stdout_mode(int mode);
```

FORTRAN SYNOPSIS

```
MP_STDOUT_MODE(INTEGER MODE)
```

PARAMETERS

mode is the mode to which STDOUT is to be set. The valid values are:

- **taskid** specifies single mode for STDOUT, where taskid is the task identifier of the new single task. This value must be between 0 and n-1, where n is the total of tasks in the current partition. The taskid requested does not have to be the issuing task.
- **-2** specifies ordered mode for STDOUT. The macro STDIO_ORDERED is supplied for use in C programs.
- **-3** specifies unordered mode for STDOUT. The macro STDIO_UNORDERED is supplied for use in C programs.

DESCRIPTION

This Parallel Utility function requests that STDOUT be set to single, ordered, or unordered mode. In single mode, only one task output is displayed. In unordered mode, output is displayed in the order received at the home node. In ordered mode, each parallel task writes output data to its own buffer. When a flush request is made all the task buffers are flushed, in order of task id, to STDOUT home node.

NOTES

- All current STDOUT buffers are flushed before the new STDOUT mode is established.
- The initial mode for STDOUT is set by using the environment variable **MP_STDOUTMODE**, or by using the poe command line option **-stdoutmode**, with the latter overriding the former. The default STDOUT mode is unordered.
- This function is implemented with a half second sleep interval to ensure that the mode change request is processed before subsequent writes to STDOUT.
- This function is thread safe.

RETURN VALUES

In C and C++ calls, the following applies:

0 indicates successful completion.

-1 indicates that an error occurred. A message describing the error will be issued.

EXAMPLES

C Example

The following program uses poe with the **-labelio yes** option and three tasks:

```
    #include <pm_util.h>

main()
{
 mpc_stdout_mode(STDIO_ORDERED);
 printf("These lines will appear in task order\n");
 /*
  * Call mpc_flush here to make sure that one task
  * doesn't change the mode before all tasks have
  * sent the previous printf string to the home node.
  */
 mpc_flush(1);
 mpc_stdout_mode(STDIO_UNORDERED);
 printf("These lines will appear in the order received by the home node\n");
 /*
  * Since synchronization is not used here, one task could actually
  * execute the next statement before one of the other tasks has
  * executed the previous statement, causing one of the unordered
  * lines not to print.
  */
 mpc_stdout_mode(1);
 printf("Only 1 copy of this line will appear from task 1\n");
}
```

Running the above C program produces the following output (task order of lines 4-6 may differ):

```
0 : These lines will appear in task order.
1 : These lines will appear in task order.
2 : These lines will appear in task order.
1 : These lines will appear in the order received by the home node.
2 : These lines will appear in the order received by the home node.
0 : These lines will appear in the order received by the home node.
1 : Only 1 copy of this line will appear from task 1.
```

FORTRAN Example

```
 CALL MP_STDOUT_MODE(-2)
 WRITE(6, *) 'These lines will appear in task order'
 CALL MP_FLUSH(1)
 CALL MP_STDOUT_MODE(-3)
 WRITE(6, *) 'These lines will appear in the order received by the
xhome node'
 CALL MP_STDOUT_MODE(1)
 WRITE(6, *) 'Only 1 copy of this line will appear from task 1'
 END
```

Running the above program produces the following output (task order of lines 4-6 may differ):

```
0 : These lines will appear in task order.
1 : These lines will appear in task order.
2 : These lines will appear in task order.
1 : These lines will appear in the order received by the home node.
2 : These lines will appear in the order received by the home node.
0 : These lines will appear in the order received by the home node.
1 : Only 1 copy of this line will appear from task 1.
```

RELATED INFORMATION

Functions: mpc_flush(3), MP_FLUSH(3), mpc_stdoutmode_query(3), MP_STDOUTMODE_QUERY(3), mpc_synch(3), MP_SYNCH(3), mpcc(1), mpCC(1), mpxlf(1)

NAME

MP_STDOUTMODE_QUERY, mpc_stdoutmode_query-Queries current STDOUT mode setting.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_stdoutmode_query(int *mode);
```

FORTRAN SYNOPSIS

```
MP_STDOUTMODE_QUERY(INTEGER MODE)
```

PARAMETERS

mode is the address of an integer in which the current STDOUT mode setting will be returned. Possible return values are:

- **taskid** indicates that the current STDOUT mode is single, i.e. output for only task taskid is displayed.
- **-2** indicates that the current STDOUT mode is ordered. The macro STDIO_ORDERED is supplied for use in C programs.
- **-3** indicates that the current STDOUT mode is unordered. The macro STDIO_UNORDERED is supplied for use in C programs.

DESCRIPTION

This Parallel Utility function returns the mode to which STDOUT is currently set.

NOTES

- Between the time one task issues a mode query request and receives a response, it is possible that another task can change the STDOUT mode setting to another value unless proper synchronization is used.
- This function is thread safe.

RETURN VALUES

In C and C++ calls, the following applies:

- **0** indicates successful completion
- **-1** indicates that an error occurred. A message describing the error will be issued.

EXAMPLES

C Example

The following program uses poe with one task:

```
#include <pm_util.h>

main()
{
 int mode;

 mpc_stdoutmode_query(&mode);
 printf("Initial (default) STDOUT mode is %d\n", mode);
 mpc_stdout_mode(STDIO_ORDERED);
 mpc_stdoutmode_query(&mode);
 printf("New STDOUT mode is %d\n", mode);
}
```

Running the above program produces the following output:

```
Initial (default) STDOUT mode is -3
New STDOUT mode is -2
```

FORTRAN Example

The following program uses poe with one task:

```
      INTEGER MODE

      CALL MP_STDOUTMODE_QUERY(mode)
      WRITE(6, *) 'Initial (default) STDOUT mode is', mode
      CALL MP_STDOUT_MODE(-2)
      CALL MP_STDOUTMODE_QUERY(mode)
      WRITE(6, *) 'New STDOUT mode is', mode
      END
```

Running the above program produces the following output:

```
Initial (default) STDOUT mode is -3
New STDOUT mode is -2
```

RELATED INFORMATION

Functions: mpc_flush(3), MP_FLUSH(3), mp_stdout-mode(3), MP_STDOUT_MODE(3), mpc_synch(3), MP_SYNCH(3), mpcc(1), mpCC(1), mpxlf(1)

NAME

mpc_isatty-Determines if a device is a terminal on the home node.

LIBRARY

libmpi.a

C SYNOPSIS

```
#include <pm_util.h>
int mpc_isatty(int FileDescriptor);
```

FORTRAN SYNOPSIS

A FORTRAN version of this function is not available.

PARAMETERS

FileDescriptor is the file descriptor number of the device. Valid values are:

0 or STDIN specifies STDIN as the device to be checked.

1 or STDOUT specifies STDOUT as the device to be checked.

2 or STDERR specifies STDERR as the device to be checked.

DESCRIPTION

This Parallel Utility function determines if the file descriptor specified by the FileDescriptor parameter is associated with a terminal device on the home node. In a Parallel Operating Environment partition, these three file descriptors are implemented as pipes to the Partition Manager Daemon. Therefore, the AIX isatty() function will always return FALSE for each of them. This function is provided for use by remote tasks that may want to know whether one of these devices is actually a terminal on the home node, for example, to determine whether or not to output a prompt.

NOTES

- This function is thread safe.

RETURN VALUES

In C and C++ calls, the following applies:

0 indicates that the device IS NOT associated with a terminal on the home.

1 indicates that the device IS associated with a terminal on the home.

-1 indicates an invalid FileDescriptor parameter.

EXAMPLES

C Example

```
/*
 * Running this program, after compiling with mpcc,
 * without redirecting STDIN, produces the following output:
 *
 *     isatty() reports STDIN as a non-terminal device
 *     mpc_isatty() reports STDIN as a terminal device
 */

#include "pm_util.h"

main()
{
 if (isatty(STDIN)) {
  printf("isatty() reports STDIN as a terminal device\n");
 } else {
  printf("isatty() reports STDIN as a non-terminal device\n");
  if (mpc_isatty(STDIN)) {
   printf("mpc_isatty() reports STDIN as a terminal device\n");
  } else {
   printf("mpc_isatty() reports STDIN as a non-terminal device\n");
  }
 }
}
```

Appendix F. Tracing Routines

This chapter contains the syntax man pages for the tracing routines used in the Visualization Tool Tracing Subsystem Tracing Subsystem. These user-callable routines allow the programmer to customize the trace collection from within the application being traced. Included are functions for:

- modifying trace collection parameters
- starting and stopping the trace from within the program, dynamically and selectively
- writing trace records back to the home node.

There is a C and FORTRAN version for each of the routines.

The tracing routines are:

VT_TRC_FLUSH and VT_trc_flush_c flushes the memory buffer to the trace collection directory.

VT_TRC_SET_PARAMS and VT_trc_set_params_c sets certain tracing parameters at which the dig should sample kernel statistics.

VT_TRC_START and VT_trc_start_c requests a trace from within the program, dynamically and selectively.

VT_TRC_STOP and VT_trc_stop_c requests that the collecting of trace events be discontinued.

For more information, refer to *IBM Parallel Environment for AIX: Operation and Use, Volume 2*.

© Copyright IBM Corp. 1997

NAME

VT_TRC_FLUSH, VT_trc_flush_c-Flushes the memory buffer to the trace collection directory.

LIBRARY

- C Library (libvtd.a)
- FORTRAN Library (libvtd.a)

  The above are automatically included by the POE functions.

C SYNOPSIS

```
#include <VT_trc.h>
int VT_trc_flush_c();
```

FORTRAN SYNOPSIS

```
VT_TRC_FLUSH(INTEGER RETURN_CODE)
```

PARAMETERS

In FORTRAN, the following applies:

return_code integer which receives the return value.

- **0** indicates successful completion
- **-1** indicates an error occurred.

DESCRIPTION

This routine may be called from the application to flush the memory buffer to the trace collection directory. This is specified by the MP_TRACE directory environment variable or -tracedir parameter and defaults to the directory from which the application was stored.

RETURN VALUES

In C and C++ calls, the following applies:

- **0** indicates successful completion
- **-1** indicates that an error occurred. A message describing the error will be issued.

ERRORS

- 0033-3090 write failed to the dig daemon, ..., error is ...

  The application could not contact the kernel statistics sampling daemon.

For more information about error conditions, refer to *IBM Parallel Environment for AIX: Messages*.

EXAMPLES

C Example

```
#include <stdlib.h>
#include <stdio.h>
#include <mpi.h>
#include <VT_trc.h>

#define COUNT 1024

int main(int argc, char **argv)
{
  int i;
  int numtask;
  int taskid;
  int msg_in[COUNT], msg_out[COUNT];
  int src, dest, len, send_msgid, recv_msgid, nBytes;

  /* Find out number of tasks/nodes. */
  MPI_Init( &argc, &argv);
  MPI_Comm_size( MPI_COMM_WORLD, &numtask);
  MPI_Comm_rank( MPI_COMM_WORLD, &taskid);

  /*************************************************

    Set tracing parameters as follows:
      Memory buffer = 500K
      Temporary trace file = 1M
      Sample every second
      Do not use wrap around buffer

   *************************************************/

  if(0!=VT_trc_set_params_c(500000,1000000,1000,0))
    {
      printf("Could not reset trace parameters!\n");
    }
  /* Disable tracing while the message buffer is being initialized */
  if(0!=VT_trc_stop_c())
    {
      printf("Could not stop tracing!\n");
    }
  /* Flush the memory buffer to disk */
  if(0!=VT_trc_flush_c())
    {
      printf("Could not flush trace buffer!\n");
    }
  for(i=0;i<COUNT;i++)
    msg_out[i]=taskid;

  /* Re-enable tracing.  Level 9 asks for everything */
  /* but only events enabled by the command line or  */
  /* environment variable will be re-enabled)        */
  if(0!=VT_trc_start_c(9))
    {
      printf("Could not restart tracing!\n");
```

```
    }

  dest = (taskid<(numtask-1))?(taskid+1):0;
  MPI_Send(msg_out, COUNT, MPI_INT, dest, 0, MPI_COMM_WORLD);

  src = (taskid>0)?(taskid-1):(numtask-1);
  MPI_Recv(msg_in, COUNT, MPI_INT, src, 0, MPI_COMM_WORLD);

  MPI_Finalize();
}
```

FORTRAN Example

```
        PROGRAM TRCDEMO
C
        INCLUDE "mpif.h"
        IMPLICIT   NONE
        INTEGER    COUNT, I
        INTEGER    BUFFSZ, FILESZ
        INTEGER    SMPL, WRAP
        PARAMETER  ( COUNT = 1024 )
        PARAMETER  ( BUFFSZ = 500000, FILESZ = 1000000 )
        PARAMETER  ( SMPL = 1000, WRAP = 0 )
        INTEGER    MSG_IN(COUNT), MSG_OUT(COUNT)
        INTEGER    NUMTASK, TASKID
        INTEGER    SRC, DEST
        INTEGER    RC
C
C FIND OUT NUMBER OF TASKS

        CALL MPI_INIT( RC )
        CALL MPI_COMM_SIZE( MPI_COMM_WORLD, NUMTASK, RC )
        CALL MPI_COMM_RANK( MPI_COMM_WORLD, TASKID, RC )

C
C
C    SET TRACING PARAMETERS AS FOLLOWS:
C      MEMORY BUFFER = 500K
C      TEMPORARY TRACE FILE = 1M
C      SAMPLE EVERY SECOND
C      DO NOT USE WRAP AROUND BUFFER
C
C


        CALL VT_TRC_SET_PARAMS(BUFFSZ, FILESZ, SMPL, WRAP, RC)
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not reset trace parameters!'
        ENDIF

C DISABLE TRACING WHILE THE MESSAGE BUFFER IS BEING INITIALIZED
        CALL VT_TRC_STOP( RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not stop tracing!'
        ENDIF

C FLUSH THE MEMORY BUFFER TO DISK
        CALL VT_TRC_FLUSH( RC )
```

```
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not flush trace buffer!'
        ENDIF

        DO I = 1,COUNT
          MSG_OUT(I) = TASKID
        ENDDO

C RE-ENABLE TRACING.  LEVEL 9 ASKS FOR EVERYTHING
C BUT ONLY EVENTS ENABLED BY THE COMMAND LINE OR
C ENVIRONMENT VARIABLE WILL BE RE-ENABLED)
        CALL VT_TRC_START( 9, RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not restart tracing!'
        ENDIF

        IF( TASKID .GE. NUMTASK-1 ) THEN
          DEST = 0
        ELSE
          DEST = TASKID + 1
        ENDIF

        CALL MPI_SEND(MSG_OUT, COUNT, MPI_INTEGER, DEST, 0,
     +            MPI_COMM_WORLD, RC )

        IF( TASKID .LE. 0 ) THEN
          SRC = NUMTASK - 1
        ELSE
          SRC = TASKID - 1
        ENDIF

        CALL MPI_RECV(MSG_IN, COUNT, MPI_INTEGER, SRC, 0,
     +           MPI_COMM_WORLD, RC)

        CALL MPI_FINALIZE( RC )

        STOP
        END
C
```

RELATED INFORMATION

Functions: VT_trc_set_params_c, VT_TRC_SET_PARAMS, VT_trc_start_c, VT_TRC_START, VT_trc_stop_c, VT_TRC_STOP

For information about VT tracing philosophy, refer to *IBM Parallel Environment for AIX: Operation and Use, Volume 2*.

NAME

VT_TRC_SET_PARAMS, VT_trc_set_params_c-Allows applications to modify certain tracing parameters during execution.

LIBRARY

- C Library (libvtd.a) functions.
- FORTRAN Library (libvtd.a)

  The above are automatically included by the POE.

C SYNOPSIS

```
#include <VT_trc.h>
int VT_trc_set_params_c(size_t buff_size, size_t file_size,
                        int sampling_freq, int wrap_around_flag);
```

FORTRAN SYNOPSIS

```
VT_TRC_SET_PARAMS(INTEGER BUFF_SIZE, INTEGER FILE_SIZE,
                  INTEGER SAMPLING_FREQ, INTEGER WRAP_AROUND_FLAG,
                  INTEGER RETURN_CODE)
```

PARAMETERS

buff_size the size of the memory buffer

file_size the size of the node trace files

sampling_freq the interval with which to sample AIX units (microseconds)

wrap_around_flag if non-zero, a wrap around memory buffer will be used.

return_code integer which receives the return value.

- **0** indicates successful completion
- **-1** indicates an error occurred.

DESCRIPTION

This routine allows applications to set certain tracing parameters such as the size of the memory buffer and trace files, and the kernel statistics sampling interval.

NOTES

- Parameter values replace command line values.

RETURN VALUES

In C and C++ calls, the following applies:

0 indicates successful completion

-1 indicates that an error occurred. A message describing the error will be issued.

ERRORS

- 0033-3101 Setting Temp File size to threshold size. Set Size =, Minimum size = "

  You tried to set the size of the temporary file less than the minimum threshold size. The program automatically set the size to the minimum and continued trace generation.

  Change the parameters in VT_TRC_SET_PARAMS call above the threshold value.

- 0033-3103 Setting Buffer size to Threshold size. Set size = , Minimum size =

  You tried to set the size of the trace buffer less than the minimum threshold size. The program automatically set the size to the threshold size and continued trace generation.

  Change the parameters in VT_TRC_SET_PARAMS call above the threshold value.

- 0033-3104 Setting system statistics sampling frequency to the threshold value. Set Size = , Minimum size =

  You tried to set the sampling frequency to be less than the minimum threshold value. The program automatically set the size to the threshold value and continued trace generation.

  Change the parameters in VT_TRC_SET_PARAMS call above the threshold value.

- 0033-3011 VT_trc_set-params(), reallocation of ... bytes failed

  The application could not obtain the memory required to change parameters.

- 0033-3090 write failed to the dig daemon, ..., error is ...

  The application could not contact the kernel statistics sampling daemon.

For more information about error conditions, refer to *IBM Parallel Environment for AIX: Messages*.

EXAMPLES

C Example

```
#include <stdlib.h>
#include <stdio.h>
#include <mpi.h>
#include <VT_trc.h>

#define COUNT 1024

int main(int argc, char **argv)
{
  int i;
  int numtask;
  int taskid;
  int msg_in[COUNT], msg_out[COUNT];
  int src, dest, len, send_msgid, recv_msgid, nBytes;

  /* Find out number of tasks/nodes. */
```

```
  MPI_Init( &argc, &argv);
  MPI_Comm_size( MPI_COMM_WORLD, &numtask);
  MPI_Comm_rank( MPI_COMM_WORLD, &taskid);

  /************************************************

    Set tracing parameters as follows:
      Memory buffer = 500K
      Temporary trace file = 1M
      Sample every second
      Do not use wrap around buffer

   ************************************************/

  if(0!=VT_trc_set_params_c(500000,1000000,1000,0))
    {
      printf("Could not reset trace parameters!\n");
    }
  /* Disable tracing while the message buffer is being initialized */
  if(0!=VT_trc_stop_c())
    {
      printf("Could not stop tracing!\n");
    }
  /* Flush the memory buffer to disk */
  if(0!=VT_trc_flush_c())
    {
      printf("Could not flush trace buffer!\n");
    }
  for(i=0;i<COUNT;i++)
    msg_out[i]=taskid;

  /* Re-enable tracing.  Level 9 asks for everything */
  /* but only events enabled by the command line or  */
  /* environment variable will be re-enabled)        */
  if(0!=VT_trc_start_c(9))
    {
      printf("Could not restart tracing!\n");
    }

  dest = (taskid<(numtask-1))?(taskid+1):0;
  MPI_Send(msg_out, COUNT, MPI_INT, dest, 0, MPI_COMM_WORLD);

  src = (taskid>0)?(taskid-1):(numtask-1);
  MPI_Recv(msg_in, COUNT, MPI_INT, src, 0, MPI_COMM_WORLD);

  MPI_Finalize();
}
```

FORTRAN Example

```
        PROGRAM TRCDEMO
C
        INCLUDE "mpif.h"
        IMPLICIT   NONE
        INTEGER    COUNT, I
        INTEGER    BUFFSZ, FILESZ
        INTEGER    SMPL, WRAP
```

```
        PARAMETER  ( COUNT = 1024 )
        PARAMETER  ( BUFFSZ = 500000, FILESZ = 1000000 )
        PARAMETER  ( SMPL = 1000, WRAP = 0 )
        INTEGER    MSG_IN(COUNT), MSG_OUT(COUNT)
        INTEGER    NUMTASK, TASKID
        INTEGER    SRC, DEST
        INTEGER    RC
C
C FIND OUT NUMBER OF TASKS

        CALL MPI_INIT( RC )
        CALL MPI_COMM_SIZE( MPI_COMM_WORLD, NUMTASK, RC )
        CALL MPI_COMM_RANK( MPI_COMM_WORLD, TASKID, RC )

C
C
C    SET TRACING PARAMETERS AS FOLLOWS:
C      MEMORY BUFFER = 500K
C      TEMPORARY TRACE FILE = 1M
C      SAMPLE EVERY SECOND
C      DO NOT USE WRAP AROUND BUFFER
C
C

        CALL VT_TRC_SET_PARAMS(BUFFSZ, FILESZ, SMPL, WRAP, RC)
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not reset trace parameters!'
        ENDIF

C DISABLE TRACING WHILE THE MESSAGE BUFFER IS BEING INITIALIZED
        CALL VT_TRC_STOP( RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not stop tracing!'
        ENDIF

C FLUSH THE MEMORY BUFFER TO DISK
        CALL VT_TRC_FLUSH( RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not flush trace buffer!'
        ENDIF

        DO I = 1,COUNT
          MSG_OUT(I) = TASKID
        ENDDO

C RE-ENABLE TRACING.  LEVEL 9 ASKS FOR EVERYTHING
C BUT ONLY EVENTS ENABLED BY THE COMMAND LINE OR
C ENVIRONMENT VARIABLE WILL BE RE-ENABLED)
        CALL VT_TRC_START( 9, RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not restart tracing!'
        ENDIF

        IF( TASKID .GE. NUMTASK-1 ) THEN
          DEST = 0
        ELSE
          DEST = TASKID + 1
        ENDIF
```

```
      CALL MPI_SEND(MSG_OUT, COUNT, MPI_INTEGER, DEST, 0,
     +           MPI_COMM_WORLD, RC )

      IF( TASKID .LE. 0 ) THEN
        SRC = NUMTASK - 1
      ELSE
        SRC = TASKID - 1
      ENDIF

      CALL MPI_RECV(MSG_IN, COUNT, MPI_INTEGER, SRC, 0,
     +           MPI_COMM_WORLD, RC)

      CALL MPI_FINALIZE( RC )

      STOP
      END
C
```

RELATED INFORMATION

Functions: VT_trc_flush_c, VT_TRC_FLUSH, VT_trc_start_c, VT_TRC_START, VT_trc_stop_c, VT_TRC_STOP

For information about VT tracing philosophy, refer to *IBM Parallel Environment for AIX: Operation and Use, Volume 2*.

NAME

VT_TRC_START, VT_trc_start_c-Allows you to start the program trace from within the program, dynamically and selectively.

LIBRARY

- C Library (libvtd.a)
- FORTRAN Library (libvtd.a)

The above is automatically included by the POE functions.

C SYNOPSIS

```
#include <VT_trc.h>
int VT_trc_start_c(int flag);
```

FORTRAN SYNOPSIS

```
VT_TRC_START(INTEGER FLAG, INTEGER RETURN_CODE)
```

PARAMETERS

flag the trace control flags that define what information is traced.

return_code integer which receives the return value.

- **0** indicates successful completion
- **-1** indicates an error occurred.

DESCRIPTION

This routine may be called by the application program. This gives the program capability to start the trace from within the program, dynamically and selectively. The flag has the same meaning as the trace flag passed by PM upon startup. If the flag indicates VT_AIX_TRACE, this routine also sends a message to the monitoring daemon process.

NOTES

- Tracing is initially on when the application starts.
- If a trace event, such as kernel statistics or communication events, was not selected at the start of the application it cannot be enabled by this routine.

RETURN VALUES

In C and C++ calls, the following applies:

0 indicates successful completion

-1 indicates that an error occurred. A message describing the error will be issued.

ERRORS

- 0033-3090 write failed to the dig daemon ...., error is ...

  The application could not contact the kernel statistics sampling daemon.

For more information about error conditions, refer to *IBM Parallel Environment for AIX: Messages*.

EXAMPLES

C Example

```
#include <stdlib.h>
#include <stdio.h>
#include <mpi.h>
#include <VT_trc.h>

#define COUNT 1024

int main(int argc, char **argv)
{
  int i;
  int numtask;
  int taskid;
  int msg_in[COUNT], msg_out[COUNT];
  int src, dest, len, send_msgid, recv_msgid, nBytes;

  /* Find out number of tasks/nodes. */
  MPI_Init( &argc, &argv);
  MPI_Comm_size( MPI_COMM_WORLD, &numtask);
  MPI_Comm_rank( MPI_COMM_WORLD, &taskid);

  /************************************************

    Set tracing parameters as follows:
      Memory buffer = 500K
      Temporary trace file = 1M
      Sample every second
      Do not use wrap around buffer

   ************************************************/

  if(0!=VT_trc_set_params_c(500000,1000000,1000,0))
    {
      printf("Could not reset trace parameters!\n");
    }
  /* Disable tracing while the message buffer is being initialized */
  if(0!=VT_trc_stop_c())
    {
      printf("Could not stop tracing!\n");
    }
  /* Flush the memory buffer to disk */
  if(0!=VT_trc_flush_c())
    {
      printf("Could not flush trace buffer!\n");
    }
  for(i=0;i<COUNT;i++)
```

```
    msg_out[i]=taskid;

  /* Re-enable tracing.  Level 9 asks for everything */
  /* but only events enabled by the command line or  */
  /* environment variable will be re-enabled)        */
  if(0!=VT_trc_start_c(9))
    {
      printf("Could not restart tracing!\n");
    }

  dest = (taskid<(numtask-1))?(taskid+1):0;
  MPI_Send(msg_out, COUNT, MPI_INT, dest, 0, MPI_COMM_WORLD);

  src = (taskid>0)?(taskid-1):(numtask-1);
  MPI_Recv(msg_in, COUNT, MPI_INT, src, 0, MPI_COMM_WORLD);

  MPI_Finalize();
}
```

FORTRAN Example

```
        PROGRAM TRCDEMO
C
        INCLUDE "mpif.h"
        IMPLICIT   NONE
        INTEGER    COUNT, I
        INTEGER    BUFFSZ, FILESZ
        INTEGER    SMPL, WRAP
        PARAMETER  ( COUNT = 1024 )
        PARAMETER  ( BUFFSZ = 500000, FILESZ = 1000000 )
        PARAMETER  ( SMPL = 1000, WRAP = 0 )
        INTEGER    MSG_IN(COUNT), MSG_OUT(COUNT)
        INTEGER    NUMTASK, TASKID
        INTEGER    SRC, DEST
        INTEGER    RC
C
C FIND OUT NUMBER OF TASKS

        CALL MPI_INIT( RC )
        CALL MPI_COMM_SIZE( MPI_COMM_WORLD, NUMTASK, RC )
        CALL MPI_COMM_RANK( MPI_COMM_WORLD, TASKID, RC )

C
C
C    SET TRACING PARAMETERS AS FOLLOWS:
C      MEMORY BUFFER = 500K
C      TEMPORARY TRACE FILE = 1M
C      SAMPLE EVERY SECOND
C      DO NOT USE WRAP AROUND BUFFER
C
C

        CALL VT_TRC_SET_PARAMS(BUFFSZ, FILESZ, SMPL, WRAP, RC)
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not reset trace parameters!'
        ENDIF
```

```
C DISABLE TRACING WHILE THE MESSAGE BUFFER IS BEING INITIALIZED
        CALL VT_TRC_STOP( RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not stop tracing!'
        ENDIF

C FLUSH THE MEMORY BUFFER TO DISK
        CALL VT_TRC_FLUSH( RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not flush trace buffer!'
        ENDIF

        DO I = 1,COUNT
          MSG_OUT(I) = TASKID
        ENDDO

C RE-ENABLE TRACING.  LEVEL 9 ASKS FOR EVERYTHING
C BUT ONLY EVENTS ENABLED BY THE COMMAND LINE OR
C ENVIRONMENT VARIABLE WILL BE RE-ENABLED)
        CALL VT_TRC_START( 9, RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not restart tracing!'
        ENDIF

        IF( TASKID .GE. NUMTASK-1 ) THEN
          DEST = 0
        ELSE
          DEST = TASKID + 1
        ENDIF

        CALL MPI_SEND(MSG_OUT, COUNT, MPI_INTEGER, DEST, 0,
     +           MPI_COMM_WORLD, RC )

        IF( TASKID .LE. 0 ) THEN
          SRC = NUMTASK - 1
        ELSE
          SRC = TASKID - 1
        ENDIF

        CALL MPI_RECV(MSG_IN, COUNT, MPI_INTEGER, SRC, 0,
     +          MPI_COMM_WORLD, RC)

        CALL MPI_FINALIZE( RC )

        STOP
        END
C
```

RELATED INFORMATION

Functions: VT_trc_flush_c, VT_TRC_FLUSH, VT_trc_set_params_c, VT_TRC_SET_PARAMS, VT_trc_start_c, VT_TRC_START, VT_trc_stop_c, VT_TRC_STOP

For information about VT tracing philosophy, refer to IBM Parallel Environment for AIX: Operation and Use, Volume 2.

NAME

VT_TRC_STOP, VT_trc_stop_c-Stops collecting trace events.

LIBRARY

- C Library (libvtd.a)
- FORTRAN Library (libvtd.a)

The above is automatically included by the POE functions.

C SYNOPSIS

```
#include <VT_trc.h>
int VT_trc_stop_c();
```

FORTRAN SYNOPSIS

```
VT_TRC_STOP(INTEGER RETURN_CODE)
```

PARAMETERS

return_code integer which receives the return value.

0 indicates successful completion

-1 indicates an error occurred.

DESCRIPTION

This routine may be called by the application program to stop tracing.

RETURN VALUES

In C and C++ calls, the following applies:

0 indicates successful completion

-1 indicates that an error occurred. A message describing the error will be issued.

ERRORS

- 0033-3090 write failed to the dig daemon ...., error is ...

  The application could not contact the kernel statistics sampling daemon.

For more information about error conditions, refer to *IBM Parallel Environment for AIX: Messages*.

EXAMPLES

C Example

```
#include <stdlib.h>
#include <stdio.h>
#include <mpi.h>
#include <VT_trc.h>

#define COUNT 1024

int main(int argc, char **argv)
{
  int i;
  int numtask;
  int taskid;
  int msg_in[COUNT], msg_out[COUNT];
  int src, dest, len, send_msgid, recv_msgid, nBytes;

  /* Find out number of tasks/nodes. */
  MPI_Init( &argc, &argv);
  MPI_Comm_size( MPI_COMM_WORLD, &numtask);
  MPI_Comm_rank( MPI_COMM_WORLD, &taskid);

  /************************************************

    Set tracing parameters as follows:
      Memory buffer = 500K
      Temporary trace file = 1M
      Sample every second
      Do not use wrap around buffer

   ************************************************/

  if(0!=VT_trc_set_params_c(500000,1000000,1000,0))
    {
      printf("Could not reset trace parameters!\n");
    }
  /* Disable tracing while the message buffer is being initialized */
  if(0!=VT_trc_stop_c())
    {
      printf("Could not stop tracing!\n");
    }
  /* Flush the memory buffer to disk */
  if(0!=VT_trc_flush_c())
    {
      printf("Could not flush trace buffer!\n");
    }
  for(i=0;i<COUNT;i++)
    msg_out[i]=taskid;

  /* Re-enable tracing.  Level 9 asks for everything */
  /* but only events enabled by the command line or  */
  /* environment variable will be re-enabled)        */
  if(0!=VT_trc_start_c(9))
    {
      printf("Could not restart tracing!\n");
    }

  dest = (taskid<(numtask-1))?(taskid+1):0;
  MPI_Send(msg_out, COUNT, MPI_INT, dest, 0, MPI_COMM_WORLD);
```

```
  src = (taskid>0)?(taskid-1):(numtask-1);
  MPI_Recv(msg_in, COUNT, MPI_INT, src, 0, MPI_COMM_WORLD);

  MPI_Finalize();
}
```

FORTRAN Example

```
        PROGRAM TRCDEMO
C
        INCLUDE "mpif.h"
        IMPLICIT   NONE
        INTEGER    COUNT, I
        INTEGER    BUFFSZ, FILESZ
        INTEGER    SMPL, WRAP
        PARAMETER  ( COUNT = 1024 )
        PARAMETER  ( BUFFSZ = 500000, FILESZ = 1000000 )
        PARAMETER  ( SMPL = 1000, WRAP = 0 )
        INTEGER    MSG_IN(COUNT), MSG_OUT(COUNT)
        INTEGER    NUMTASK, TASKID
        INTEGER    SRC, DEST
        INTEGER    RC
C
C FIND OUT NUMBER OF TASKS

        CALL MPI_INIT( RC )
        CALL MPI_COMM_SIZE( MPI_COMM_WORLD, NUMTASK, RC )
        CALL MPI_COMM_RANK( MPI_COMM_WORLD, TASKID, RC )

C
C
C    SET TRACING PARAMETERS AS FOLLOWS:
C      MEMORY BUFFER = 500K
C      TEMPORARY TRACE FILE = 1M
C      SAMPLE EVERY SECOND
C      DO NOT USE WRAP AROUND BUFFER
C
C


        CALL VT_TRC_SET_PARAMS(BUFFSZ, FILESZ, SMPL, WRAP, RC)
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not reset trace parameters!'
        ENDIF

C DISABLE TRACING WHILE THE MESSAGE BUFFER IS BEING INITIALIZED
        CALL VT_TRC_STOP( RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not stop tracing!'
        ENDIF

C FLUSH THE MEMORY BUFFER TO DISK
        CALL VT_TRC_FLUSH( RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not flush trace buffer!'
        ENDIF
```

```
        DO I = 1,COUNT
          MSG_OUT(I) = TASKID
        ENDDO

C RE-ENABLE TRACING.  LEVEL 9 ASKS FOR EVERYTHING
C BUT ONLY EVENTS ENABLED BY THE COMMAND LINE OR
C ENVIRONMENT VARIABLE WILL BE RE-ENABLED)
        CALL VT_TRC_START( 9, RC )
        IF(RC .NE. 0) THEN
          WRITE(6,*)'Could not restart tracing!'
        ENDIF

        IF( TASKID .GE. NUMTASK-1 ) THEN
          DEST = 0
        ELSE
          DEST = TASKID + 1
        ENDIF

        CALL MPI_SEND(MSG_OUT, COUNT, MPI_INTEGER, DEST, 0,
     +            MPI_COMM_WORLD, RC )

        IF( TASKID .LE. 0 ) THEN
          SRC = NUMTASK - 1
        ELSE
          SRC = TASKID - 1
        ENDIF

        CALL MPI_RECV(MSG_IN, COUNT, MPI_INTEGER, SRC, 0,
     +           MPI_COMM_WORLD, RC)

        CALL MPI_FINALIZE( RC )

        STOP
        END
C
```

RELATED INFORMATION

Functions: VT_trc_flush_c, VT_TRC_FLUSH, VT_trc_set_params_c, VT_TRC_SET_PARAMS, VT_trc_start_c, VT_TRC_START

For information about VT tracing philosophy, refer to *IBM Parallel Environment for AIX: Operation and Use, Volume 2*.

Appendix G. Programming Considerations for User Applications in POE

This appendix documents various limitations, restrictions, and programming considerations for user applications written to run under the IBM Parallel Environment for AIX(PE) program product.

PE includes two versions of the message passing libraries. These are called the signal-handling library and the threaded library.

- The signal-handling library uses AIX signals as an asynchronous way to move data in and out of message buffers. They also ensure that message packets are acknowledged and retransmitted when necessary. It supports both MPL and MPI calls.
- The threaded library uses AIX kernel threads for the same message passing tasks. It supports MPI only. The threaded library also supports message passing on user-created threads. The threaded library is required if MPI coexists with the other user space protocols, for example, the LAPI interface on the RS/6000 Scalable POWERparallel Systems(SP).

The subsections following list the programming considerations common to both libraries, as well as those unique to either the signal-handling or threaded library. There is also a subsection on using POE and the FORTRAN compiler. Specifically, the sections are as follows:

- "MPI Signal-Handling And MPI Threaded Library Considerations"
- "MPI Signal-Handling Library Considerations" on page 337
- "MPI Threaded Library Considerations" on page 340
- "FORTRAN Considerations" on page 344

MPI Signal-Handling And MPI Threaded Library Considerations

The information in this section pertains to both the (MPL/MPI) signal-handling library and the MPI threaded library.

Environment Overview

As the end user, you are encouraged to think of the Parallel Operating Environment(POE) (also referred to as the **poe** command) as an ordinary (serial) command. It accepts redirected I/O, can be run under the **nice** and **time** commands, interprets command flags, and can be invoked in shell scripts.

An n-task parallel job running in the Parallel Operating Environment actually consists of the n user tasks, an equal number (n) of instances of the IBM Parallel Environment for AIX **pmd** daemon (which is the parent process of the user's task), and the POE "home node" process in which the **poe** command runs. A **pmd** daemon is started by the POE home node on each machine on which each user task runs, and serves as the point of contact between the home node and the user's tasks.

The POE home node routes standard input, standard output and standard error streams between the home node and the user's tasks via the **pmd** daemon, using TCP/IP sockets for this purpose. The sockets are created when the POE home

© Copyright IBM Corp. 1997

node starts the **pmd** daemon for each task of a parallel job. The POE home node and **pmd** also use the sockets to exchange control messages to provide task synchronization, exit status and signaling. These capabilities do not depend upon the message passing library and are available to control any parallel program run by the **poe** command.

Exit Status

Exit status is a value between 0 and 255 inclusive. It is returned from POE on the home node reflecting the composite exit status of your parallel application, as follows:

- If MPI_ABORT(comm,nn>0,ierror) or MPI_Abort(comm,nn>0) is called, the exit status is nn (mod 256).
- If MP_STOPALL(nn>=0) or mpc_stopall(nn>=0) is called, the exit status is nn (mod 256). This does not apply to threaded libraries.
- If all tasks terminate via exit(MM>=0) or STOP MM>=0 and MM is not equal to 1 and is <128 for all nodes, then POE provides a synchronization barrier at the exit. The exit status is the largest value of MM from any parallel job (mod 256).
- If any task terminates via exit(MM =1) or STOP MM =1, then POE will immediately terminate the parallel job, as if MP_STOPALL(1) or MPI_Abort(MPI_COMM_WORLD,1) had been called. This may also occur if a FORTRAN I/O library error occurs.
- If any task terminates via a signal (for example, a segment violation), the exit status is 128+signal and the entire job is immediately terminated.
- If POE terminates before the start of the user's application, the exit status is =1.
- If the user's application cannot be loaded or fails before the user's main() is called, the exit status is =255.
- You should explicitly call exit(MM) or STOP MM to set the desired exit code. A program exiting without an explicit exit value returns unpredictable status, and may result in causing premature termination of the parallel application.

POE Job Step Function

The POE job-step function is intended for the execution of a sequence of separate yet inter-related dependent programs. Therefore, it provides you with a job control mechanism that allows both job-step progression and job-step termination. The job control mechanism is the program's exit code.

- Job-step progression:

  POE continues the job-step sequence if the task exit code is 0 or in the range of 2 - 127.

- Job-step termination:

  POE terminates the parallel job, and does not execute any remaining user programs in the job-step list if the task exit code is 1 or greater than 127.

- Default termination:

  Any POE infrastructure detected failure (such as failure to open pipes to the child process or an exec failure to start the user's executable) terminates the parallel job, and does not execute any remaining user programs in the job-step queue.

POE Additions To The User Executable

POE links in the following routines when your executable is compiled with any of the POE compilation scripts (mpcc, mpcc_r, mpxlf,etc.).

Signal Handlers

POE installs signal handlers for most signals that cause program termination in order to notify the other tasks of termination and to complete the VT trace file, if enabled. POE then causes the program to exit normally with a code of (128+signal). If you install a signal handler for any of these signals, it should call the POE registered signal handler if the process decides to terminate. (See "Let POE Handle Signals When Possible" on page 332.)

The following are the signals specifically handled by POE or the message passing library:

- SIGHUP

  Caught and exits with an exit code of 128+SIGHUP.

- SIGINT

  Caught and exits with an exit code of 128+SIGINT.

  Note: This signal may be caught by user or by dbx, in which case this usage is ignored.

- SIGQUIT

  Caught, sets the default signal handler and calls exit handler with an exit code of 128+SIGQUIT. The exit handler dumps the user's context and takes the default signal action.

- SIGFPE

  Caught, sets the default signal handler and calls exit handler with an exit code of 128+SIGFPE. The exit handler dumps the user's context and takes the default signal action.

- SIGSEGV

  Caught, sets the default signal handler and calls exit handler with an exit code of 128+SIGSEGV. The exit handler dumps the user's context and takes the default signal action.

- SIGBUS

  Caught, sets the default signal handler and calls exit handler with an exit code of 128+SIGBUS. The exit handler dumps the user's context and takes the default signal action.

- SIGTERM

  Caught and exits with an exit code of 128+SIGTERM. This is also used by POE to signal orderly termination of a parallel job. If it must be caught by the user, please read carefully the section on program termination (below).

- SIGSTOP

  Default action (cannot be caught)

- SIGTSTP

  Default action

- SIGCONT

  Default action

- SIGPWR

  Caught, sets the default signal handler and calls exit handler with an exit code of 128+SIGPWR. The exit handler dumps the user's context and takes the default signal action.

- SIGDANGER

  Caught and exits with an exit code of 128+SIGDANGER.

The signal-handling library uses SIGIO, SIGALRM and SIGPIPE for its operations and it also handles these signals. For more information about the signal-handling library, see the subsection “MPI Signal-Handling Library Considerations” on page 337.

Replacement exit/atexit

POE requires its own versions of the library exit()/atexit() functions, and expects to load them dynamically from its own version of **libc.a** (or **libc_r.a**) in **/usr/lpp/ppe.poe/lib**; therefore, do not code your own exit function to override the library function. This is to synchronize profiling and to provide barrier synchronization upon exit.

Let POE Handle Signals When Possible

Programs that handle signals must coordinate with POE's handling of most of the common signals (see above).

DO NOT issue message passing calls from signal handlers. Also, many AIX library calls are not "signal safe", and should not be issued from signal handlers. Check the AIX Technical Reference (function sigaction()) for a list of AIX functions callable from signal handlers.

POE sets up signal handlers for all the signals that normally terminate program execution. It does this so that it can terminate the entire parallel job in an orderly fashion if one task terminates abnormally (via signal). A user program may install a handler for any or all of these signals, but should save the address of the POE signal handler. If the user program decides to terminate, it should call the POE signal handler. If the user program decides not to terminate, it should just return to the interrupted code. SIGTERM is used by POE to shutdown the parallel job in a variety of abnormal circumstances, and should be allowed to terminate the job.

The POE home node converts a user's SIGTSTP signal (Ctrl-z) to a SIGSTOP signal to all the remote nodes, and passes the SIGCONT signal sent by the **fg** or **bg** command to all the remote nodes to restart the job.

Don't Hard Code File Descriptor Numbers

Do not use hard coded file descriptor numbers beyond those specified by STDIN, STDOUT and STDERR.

POE opens several files and uses file descriptors as message passing handles. These are allocated before the user gets control, so the first file descriptor allocated to a user is unpredictable.

Termination Of A Parallel Job

POE provides for orderly termination of a parallel job, so that all tasks terminate at the same time. This is accomplished in the atexit routine registered at program initialization. For normal exits (codes 0, 2-127), the atexit routine sends a control message to the POE home node, and waits for a positive response. For abnormal exits and those which don't go through the atexit routine, the **pmd** daemon catches the exit code and sends a control message to the POE home node.

For normal exits, when POE gets a control message for every task, it responds to each node, allowing that node to exit normally with its individual exit code. The **pmd** daemon monitors the exit code and passes it back to the POE home node for presentation to the user.

For abnormal exits and those detected by **pmd**, POE sends a message to each **pmd** asking that it send a SIGTERM signal to its task, thereby terminating the task. When the task finally exits, **pmd** sends its exit code back to the POE home node and exits itself.

User-initiated termination of the POE home node via SIGINT (Ctrl-c) and/or SIGQUIT (Ctrl-\) causes a message to be sent to **pmd** asking that the appropriate signal be sent to the parallel task. Again, **pmd** waits for the task to die then terminates itself.

For programs that use the MPI threaded libraries, pthread_exit should not be used to terminate the POE job. POE's atexit handler is not executed when pthread_exit() is called. If the atexit handler is not executed, it is possible that one or more of POE's threads will not be terminated, causing the job to hang.

Your Program Can't Run As Root

To prevent uncontrolled root access to the entire parallel job computation resource, POE checks to see that the user is not root as part of its authentication.

AIX Function Limitations

The use of the following AIX functions may be limited, but no formal testing has been done:

- wide character sets
- shared memory - the message passing library uses shared memory for adapter mapping. You can use the remaining data segments as desired.
- getuinfo does not show terminal information, since the user program running in the parallel partition does not have an attached terminal.

Shell Execution

You can have POE run a shell script which is loaded and run on the remote nodes as if it were a binary file.

If the POE home node process is not started under the Korn shell, mounted file system names may not be mapped correctly to the names defined for the automount daemon or AIX equivalent running on the RS/6000 Scalable POWERparallel Systems(SP). See the *IBM Parallel Environment for AIX: Operation and Use, Volume 1* for a discussion of alternative name mapping techniques.

The program executed by POE on the parallel nodes does not run under a shell on those nodes. Redirection and piping of STDIO applies to the POE home node (**poe** binary), and not the user's code. If shell processing of a command line is desired on the remote nodes, invoke a shell script on the remote nodes to provide the desired preprocessing before the user's application is executed.

Do Not Rewind stdin, stdout Or stderr

The partition manager daemon uses pipes to direct stdin, stdout and stderr to the user's program, therefore, do not rewind these files.

Ensuring String Arguments Are Passed To Your Program Correctly

Quotation marks, either single or double, used as argument delimiters are stripped away by the shell and are never "seen" by poe. Therefore, the quotation marks must be escaped to allow the quoted string to be passed correctly to the remote task(s) as one argument. For example, if you want to pass the following string to the user program (including the imbedded blank)

```
a b
```

then you need to enter the following:

```
poe user_program \"a b\"
```

user_program is passed the following argument as one token:

```
a b
```

Without the backslashes, the string would have been treated as two arguments (a *and* b).

POE behaves like rsh when arguments are passed to POE. Therefore, the following:

```
poe user_program "a b"
```

is equivalent to:

```
rsh some_machine user_program "a b"
```

In order to pass the string argument as one token, the quotes have to be escaped.

Network Tuning Considerations

Programs generating large volumes of STDOUT or STDERR may overload the home node. As described previously, standard output and standard error files generated by a user's program are piped to **pmd**, then forwarded to the poe binary via a TCP/IP socket. It is possible to generate so much data that the IP message buffers on the home node are exhausted, the poe binary hangs and possibly the entire node may hang). Note that the option -stdoutmode (environment variable

MP_STDOUTMODE) controls which output stream is displayed by the poe binary, but does not limit the standard output traffic received from the remote nodes, even if set to display the output of just one node.

The POE environment variable MP_SNDBUF can be used to override the default network settings for the size of the TCP buffers used.

If you have large volumes of standard I/O, work with your network administrator to establish appropriate TCP/IP tuning parameters. You may also want to examine if using named pipes is appropriate for your application.

Standard I/O Requires Special Attention

When your program runs on the remote nodes, it has no controlling terminal. STDIN and STDOUT, STDERR are always piped.

Programs that depend on piping standard input or standard output as part of a processing sequence may wish to bypass the home node poe binary. Running the **poe** command (or starting a program compiled with one of the POE compile scripts) causes the poe binary to be loaded on the machine on which you typed the command (the POE home node). The poe binary, in turn, starts a daemon named **pmd** on each parallel node assigned to run the job, and then requests **pmd** to run your executable (via fork and exec). The poe binary reads STDIN and passes it to each of the parallel tasks via a TCP/IP socket connection to the **pmd** daemon, which pipes it to the user. Similarly, STDOUT and STDERR from the user are piped to **pmd** and sent on the socket back to the home node, where it is written to the poe binary's STDOUT and STDERR descriptors. If you know that the task reading STDIN or writing STDOUT must be on the same node (processor) as the poe binary (the poe home node), named pipes can be used to bypass poe's reading and forwarding STDIN and STDOUT.

If STDIN is piped or redirected to the poe binary (via ordinary pipes), and your application is linked with the signal handling message passing library, (via mpcc, mpxlf, or mpCC), then set the environment variable MP_HOLD_STDIN to "yes". This lets poe initialize the signal-handling library before handling the STDIN file.

If your application is linked with the threaded library, see “Standard I/O Requires Special Attention” on page 342 for more information.

STDIN/STDOUT Piping Example

The following two scripts show how STDIN and STDOUT can be piped directly between prior and post processing steps, bypassing the POE home node process. This example assumes that parallel task 0 is known or forced to be on the same node as the POE home node.

The script compute_home runs on the home node; the script compute_parallel runs on the parallel nodes (those running tasks 0 through n-1).

```
compute_home:
#! /bin/ksh
# Example script compute_home runs three processes:
#    data_generator creates/gets data and writes to stdout
#    data_processor is a parallel program that reads data
#      from stdin, processes it in parallel, and writes
#      the results to stdout.
#    data_consumer reads data from stdin and summarizes it
#
mkfifo poe_in_$$
mkfifo poe_out_$$
export MP_STDOUTMODE=0
export MP_STDINMODE=0
data_generator >poe_in_$$ |
     poe compute_parallel poe_in_$$ poe_out_$$ data_processor |
     data_consumer <poe_out_$$
 rc=$?
 rm poe_in_$$
 rm poe_out_$$
 exit rc


compute_parallel:
#! /bin/ksh
# Example script compute_parallel is a shell script that
#    takes the following arguments:
#    1) name of input named pipe (stdin)
#    2) name of output named pipe (stdout)
#    3) name of program to be run (and arguments)
#
poe_in=$1
poe_out=$2
shift 2
$*   <$poe_in   >$poe_out
```

Reserved Environment Variables

Environment variables starting with MP_ are intended for use by POE, and should be set only as instructed in the documentation. POE also uses a handful of MP_... environment variables for internal purposes, which should not be interfered with.

AIX Message Catalog Considerations

POE assumes that NLSPATH contains the appropriate POE message catalogs, even if LANG is set to "C" or is unset. Duplicate message catalogs are provided for languages "En_US", "en_US", and "C".

Language Bindings

The FORTRAN, C and C++ bindings for MPI are contained in the same library and can be freely intermixed.

- libmpi.a for the signal-handling version
- libmpi_r.a for the threaded version

Refer to “FORTRAN Considerations” on page 344 for more information about the FORTRAN compiler.

The AIX compilers support the flag -qarch. This option allows you to target code generation to a particular processor architecture. While this option can provide performance enhancements on specific platforms, it inhibits portability, particularly between the Power and PowerPC machines. The MPI library is not targeted to a specific architecture and is the same on PowerPc and Power nodes.

MPI Signal-Handling Library Considerations

The information in this subsection provides you with specific additional programming considerations for when you are using POE and the MPL/MPI signal-handling library.

POE Gets Control First And Handles Process Initialization

POE sets up its environment environment via the entry point mp_main(). mp_main() initializes the message passing library, sets up signal handlers, sets up an atexit routine, and initializes VT trace data collection before calling your main program.

Using Message Passing Handlers

Only a subset of MPL message passing is allowed on handlers created by the MPL Receive and Call function (mpc_rcvncall or MP_RCNVCALL). MPI calls on these handlers are not supported.

POE Additions To The User Executable

POE links in the following routines when your executable is compiled with mpcc, mpxlf or mpCC. These are routines specific for the signal handling environment.

Message Passing Initialization Module

POE initializes the parallel message passing library and determines that all nodes can communicate successfully before the user main() program gains control. As a result, any program compiled with the POE compiler scripts must be run under the control of POE and is not suitable as a serial program.

If communication initialization fails, the parallel task is terminated with an appropriate exit code.

Signal Handlers

The message passing library sets up signal handlers for SIGALRM, SIGIO and SIGPIPE to manage message passing activity. A user program may install a handler for any or all of these signals, but should save the address of and invoke the POE signal handler before returning to the interrupted code. The sigaction() function returns the required structure. Also, set SA_RESTART as well as the mask so all signals are masked when the signal handler is running.

The following are the signals used and specifically handled by the message passing library in a signal handling environment:

- SIGPIPE

  Caught by the non-threaded User Space message passing library to manage the RS/6000 SP switch. If your application catches this signal, it should call the registered message passing signal handler before returning to the main code.

Do not block this signal for more than a few milliseconds.

- SIGALRM

  Caught by message passing library to manage message traffic. If you provide your own interval timing mechanism, then you should arrange to call the POE signal handler approximately every 200-800 milliseconds. Message passing calls from user programs may be blocked until the POE signal handler is called.

  If the user application catches this signal but doesn't do interval timing, it should call the registered message passing signal handler before returning to the main code.

- SIGIO

  Caught by the user space message passing library to manage message traffic. If your application catches this signal, it should call the registered message passing signal handler before returning to the main code.

Interrupted System Calls

The message passing library uses an interval timer to manage message traffic, specifically to ensure that messages progress even when message passing calls are not being made. When this interval timer expires, a SIGALRM signal is sent to the program, interrupting whatever computation is in progress. The message passing library has a signal handler set, and normally handles the signal and returns to the user's program without the program's knowledge. However, the following library and system calls are interrupted and do not complete normally. The user is responsible for testing whether an interrupt occurred and recovering from the interrupt. In many cases, this is accomplished by just retrying the call.

- sleep(see note below)/usleep/nsleep
- select
- open/close/fopen/fclose
- pause
- sigpause
- accept
- connect
- recv/recvfrom/recvmsg
- send/sendto/sendmsg
- aio_read/aio_write/aio_suspend
- fork
- system
- exec/execv/...
- msem_lock/semop
- AIX msg... routines
- poll

Note: The normal timer interval is less than 500 milliseconds. So a sleep call (with time specified in seconds) returns to the original sleep interval, due to rounding, and can't be used to determine how much time remains in the interval. You should use the functions usleep and nsleep instead. See also the “Sample Replacement Sleep Program” on page 345 in Appendix H, “Using Signals and the IBM PE Programs” on page 345.

With the exception of sleep, system and exec, the routines listed above set the system error indicator (the variable errno) to EINTR, which can be tested by the

user's program. See the “Sample Replacement Select Program” on page 345 in Appendix H, “Using Signals and the IBM PE Programs” on page 345.

Normal file read and write are restarted automatically by AIX, and should not require any special treatment by the user.

The **system** and **fork** calls create a new process in which the interval timer is still running. If a fork is followed by an exec (which is what **system** does), the signal handler for the timer is overlaid, and the process is terminated when the interval timer expires.

To handle this for the **system** call, temporarily turn the interval timer off (using the **alarm**(0) call) before the call, and turn it on again (**ualarm**(500000, 500000) will do) after the **system** call.

To handle the interval timer for a forked child, merely turn off the interval timer via **alarm**(0) in the child.

There are other restrictions on fork described below.

Forks Are Limited

As described earlier, if a process forks, the forked child inherits the running timer. The timer should be turned off before forking another program. If the forked child does not exec another program, it should be aware that an atexit routine has been registered for the parent which is also inherited by the child. In most cases, the atexit routine will request POE to terminate the task (parent). A forked child should terminate with an _exit(0) system call to prevent the atexit routine from being called. Also, if the forked parent terminates before the child, the child process will not be cleaned up by POE.

A forked child MUST NOT call the message passing library.

MPI Threaded Library Considerations

When programming in a threaded environment specific skills and considerations are required. The information in this subsection provides you with specific programming considerations when using POE and the MPI threaded library. It assumes you are familiar with POSIX threads in general including mutexes, thread condition waiting, thread-specific storage, thread creation and termination.

POE Gets Control First And Handles Process Initialization

POE sets up its environment via the entry point mp_main_r(). mp_main_r() sets up signal handlers, initializes VT, and sets up an atexit routine before calling your main program.

Note: In the threaded library, message passing initialization takes place when MPI_INIT is called and not by mp_main_r. The threaded library and the signal-handling library differ significantly in this regard.

Language Bindings

The FORTRAN, C and C++ bindings for MPI are contained in the same library (libmpi_r.a) and can be freely intermixed.

Refer to “FORTRAN Considerations” on page 344 for more information about running FORTRAN programs in a threaded environment.

Use of AIX Signals

The threaded MPI library creates a thread to handle the following asynchronous signals:

- SIGQUIT
- SIGPWR
- SIGDANGER
- SIGTSTP
- SIGTERM
- SIGHUP
- SIGINT

These signals are handled with sigwait(). A signal handler must not be invoked by any signal that is being handled by sigwait().

SIGALRM

The threaded library does not use SIGALRM and *long* system calls such as sleep() are not interrupted by the message passing library. For example, sleep() runs its entire duration unless interrupted by a user-generated event.

SIGIO

PE blocks SIGIO before calling your program. SIGIO is used in the IP version of the library to notify you of an I/O event or the arrival of a message packet. This notification is enabled via the environment variable MP_CSS_INTERRUPT. If this environment variable is set to YES, the message packet arrival dispatches the interrupt service thread to process the packet.

The User Space version of the library receives notification of an arriving packet via an AIX kernel event and does not use SIGO. You may unblock it or use sigwait to process SIGIO signals.

If you've registered a signal handler (via sigaction) for SIGIO before MPI_INIT is called, the function is added to the interrupt service thread and is executed each time the service thread is dispatched. Although registered as a signal handler, the function is not required to be signal safe because it is executed on a thread. You can use pthread calls to communicate with other threads. You cannot call MPI functions in this handler.

After MPI_FINALIZE is called, your signal handler is restored but you need to unblock SIGIO in order to receive subsequent SIGIO signals.

If you register or change the SIGIO signal handler after calling MPI_INIT, your changes are ignored by the MPI library but your changes are not undone by MPI_FINALIZE.

SIGPIPE

Neither the threaded or non-threaded IP libraries use SIGPIPE. The threaded User Space library polls a variable set by the AIX kernel to determine if the switch has faulted and needs to be restarted. As a result, it does not use SIGPIPE.

Limitations In Setting The Thread Stacksize

The main thread stacksize is the same as the stacksize used for non-threaded applications. If you write your own MPI reduce functions to use with nonblocking collective communications or a SIGIO handler that will be executed on one of the library service threads, you are limited to a stacksize of 64KB by default. To increase your thread stacksize, use the environment variable MP_THREAD_STACKSIZE.

Forks Are Limited

If a process forks, only the thread that forked exists in the child process. Therefore, the message passing library will not operate properly. Also, if the forked child does not exec another program, it should be aware that an atexit routine has been registered for the parent which is also inherited by the child. In most cases, the atexit routine requests that POE terminate the task (parent). A forked child should terminate with an _exit(0) system call to prevent the atexit routine from being called. Also, if the forked parent terminates before the child, the child process will not be cleaned up by POE.

A forked child MUST NOT call the message passing library.

Standard I/O Requires Special Attention

When your program runs on the remote nodes, it has no controlling terminal. STDIN and STDOUT, STDERR are always piped.

If your threaded MPI program processes STDIN from a large file on the home node, you must do one of the following:

- Invoke MPI_Init() before performing any STDIN processing, or
- Ensure that all STDIN has been processed (EOF) before invoking MPI_Init().

This also includes programs which may not explicitly use MPI.

If STDIN is piped (or redirected) to the poe binary (via ordinary pipes) and your application is linked with the threaded library, then handle STDIN in the following way:

- If all of STDIN is read by your program before MPI_Init is called, set the environment variable MP_HOLD_STDIN=NO.
- If none of STDIN is read before MPI_Init is called, set the environment variable MP_HOLD_STDIN=YES.
- If STDIN is less than approximately 4000 bytes in length, set MP_HOLD_STDIN=NO.
- If none of the above applies, it may not be possible to run your program correctly, and you will have to devise some other mechanism for providing data to your program.

Thread-Safe Libraries

AIX provides thread-safe versions of some libraries, such as libc_r.a. However, not all libraries have a thread-safe version. It is your responsibility to determine whether the libraries you use can be safely called by more than one thread.

Program And Thread Termination

MPI_FINALIZE terminates the MPI service threads but does not affect user-created threads. Use pthread_exit to terminate any user-created threads, and exit(m) to terminate the main program (initial thread). The value of m is used to set POE's exit status as explained on “Exit Status” on page 330.

Other Thread-Specific Considerations

Order Requirement For System Includes

For threaded programs, AIX requires that the system include <pthread.h> must be first with <stdio.h> or other system includes following it. <pthread.h> defines some conditional compile variables that modify the code generation of subsequent includes, particularly <stdio.h>. Please note that <pthread.h> is not required unless your file uses thread-related calls or data.

MPI_INIT

Call MPI_INIT once per task not once per thread. MPI_INIT does not have to be called on the main thread but MPI_INIT and MPI_FINALIZE must be called on the same thread.

MPI calls on other threads must adhere to the MPI standard in regard to the following:

- A thread cannot make MPI calls until MPI_INIT has been called.
- A thread cannot make MPI calls after MPI_FINALIZE has been called.
- Unless there is a specific thread protocol programmed, you cannot rely on any specific order or speed of thread processing.

Collective Communications

Collective communications must meet the MPI standard requirement that all participating tasks execute collective communications in the same order. If collective communications calls are made on multiple threads, it is your responsibility to ensure the proper sequencing.

FORTRAN Considerations

The information in this subsection provides you with some specific programming considerations for when you are using POE and the FORTRAN compiler.

Fortran 90 and MPI

Incompatibilities exist between Fortran 90 and MPI which may effect the ability to use such programs. Refer to the information in

/usr/lpp/ppe.poe/samples/mpif90/README.mpif90

for further details. PE, Version 2, Release 2 provided the header file mpif90.h for use with Fortran 90. The file is still available in PE, Version 2, Release 3 but should not be used by new code. The mpif.h header file is formatted to work with either **mpxlf90** or **mpxlf** compilation.

FORTRAN and Threads

The current version (4.1) of the AIX XLF FORTRAN compiler is not thread-safe. However, XLF version 4.1.0.1 provides a partial thread support XLF runtime library. It supports multi-threaded applications that have one FORTRAN thread. Be sure you thoroughly test such use.

The partial thread-support library is **libxlf90_t.a** and is installed as **/usr/lib/libxlf90_t.a**. When you use the mpxlf_r command, this library is automatically included.

Restrictions

When you use **libxlf90_t.a** the following restrictions apply. Therefore, only one FORTRAN thread in a multi-threaded application may use the library.

- Routines in the library are not thread-reentrant.
- Use of routines in the math library (**libm.a**) by more than one thread may produce unpredictable results.

Appendix H. Using Signals and the IBM PE Programs

This section applies to the signal-handling version of the Message Passing library. Any AIX function that is interruptible by a signal may not behave as expected because the message passing subsystem uses timer signals to manage message traffic. For example, the user's program does not sleep for the full time but returns quickly with an error code of EINTR. This indicates the sleep was interrupted by a signal. This happens for select system call as well.

The following are some sample programs to replace sleep and select.

Sample Replacement Sleep Program

The following sample replacement program for sleep guarantees to sleep for the specified interval, even if interrupted.

Sleep Example

```
#include <errno.h>
#include <sys/time.h>
int SLEEP(int amount)
{
 struct timestruc_t Requested, Remaining;
 double famount = amount;
 int rc;
 while (famount > 0.0) {
  Requested.tv_sec = (int) famount;
  Requested.tv_nsec =
   (int) ((famount - Requested.tv_sec)*1000000000.);
  rc = nsleep ( &Requested, &Remaining );
  if ((rc == -1) && (errno == EINTR)) {
   /* Sleep interrupted.  Resume it */
   famount = Remaining.tv_sec + Remaining.tv_nsec /
       1000000000.;
   continue;
  }
  else /* Completed sleep.  Set return to zero */
  {
    return (0);
  }
 }    /* end of while */

 /* famount = 0; exit */
 return (0);
}
```

Sample Replacement Select Program

The following is a sample replacement program for select. SELECT restores the status of the file descriptor bit masks and handles the remaining time after an interrupt.

Select Example

© Copyright IBM Corp. 1997

```
#include <stdio.h>
#include <sys/select.h>
#include <sys/types.h>
#include <sys/time.h>
#include <errno.h>

int SELECT(int maxfds, fd_set *reads, fd_set *writes, fd_set *errors,
  struct timeval *timeout)
{
 struct timestruc_t Timer1, Timer2;
 struct timeval timetogo;
 static fd_set readcopy;
 static fd_set writecopy;
 static fd_set errcopy;
 int rc;
 double worktime;
 double remaining;

 /* If we get interrupted, will need to restore select bits */
 if (reads) bcopy(reads,&readcopy,sizeof(fd_set));
 if (writes) bcopy(writes,&writecopy,sizeof(fd_set));
 if (errors) bcopy(errors,&errcopy,sizeof(fd_set));

 /* two cases: if timeout specifies a time structure, we
  need to worry about timeouts.  Otherwise, we can
  ignore it */

 if (timeout == NULL) {
  while (TRUE) {
   rc = select(maxfds,reads,writes,errors,NULL);
   if ((rc == -1) && (errno == EINTR)) { /* interrupted */
    if (reads) bcopy(&readcopy,reads,sizeof(fd_set));
    if (writes) bcopy(&writecopy,writes,sizeof(fd_set));
    if (errors) bcopy(&errcopy,errors,sizeof(fd_set));
    continue;
   }
   else return(rc);
  }
 }
 else  { /* timeout is not null */
  timetogo.tv_sec = timeout->tv_sec;
  timetogo.tv_usec = timeout->tv_usec;
  remaining = timetogo.tv_sec + timetogo.tv_usec/1000000.;
/*
  fprintf(stderr,"remaining time = %f\n",remaining);
  fflush(stderr);
*/
  gettimer(TIMEOFDAY, &Timer2);
  while (TRUE) {
   Timer1.tv_sec = Timer2.tv_sec;
   Timer1.tv_nsec = Timer2.tv_nsec;
   rc = select(maxfds,reads,writes,errors,&timetogo);
   if ((rc == -1) && (errno == EINTR)) { /* interrupted */
    gettimer(TIMEOFDAY, &Timer2);
    /* compute amount remaining */
    worktime = (Timer2.tv_sec - Timer1.tv_sec) +
     (Timer2.tv_nsec - Timer1.tv_nsec)/1000000000.;
    remaining = remaining - worktime;
```

```
    timetogo.tv_sec = (int) remaining;
    timetogo.tv_usec = (int) ((remaining - timetogo.tv_sec)*
        1000000.);
    /* restore the select bits */
    if (reads) bcopy(&readcopy,reads,sizeof(fd_set));
    if (writes) bcopy(&writecopy,writes,sizeof(fd_set));
    if (errors) bcopy(&errcopy,errors,sizeof(fd_set));
    continue;
   }
   else return(rc);
  }
 }
}
```

Appendix I. Predefined Datatypes

The following is a list of the various predefined MPI datatypes by category that you can use with the signal-handling library.

Special Purpose

Datatype	Description
MPI_LB	Explicit lower bound marker
MPI_UB	Explicit upper bound marker
MPI_BYTE	Untyped byte data
MPI_PACKED	Packed data (byte)

For C Language Bindings

Datatype	Description
MPI_CHAR	8 bit character
MPI_UNSIGNED_CHAR	8 bit unsigned character
MPI_SIGNED_CHAR	8 bit signed character
MPI_SHORT	16 bit integer
MPI_INT	32 bit integer
MPI_LONG	32 bit integer
MPI_UNSIGNED_SHORT	16 bit unsigned integer
MPI_UNSIGNED	32 bit unsigned integer
MPI_UNSIGNED_LONG	32 bit unsigned integer
MPI_FLOAT	32 bit floating point
MPI_DOUBLE	64 bit floating point
MPI_LONG_DOUBLE	64 bit floating point

© Copyright IBM Corp. 1997

For FORTRAN Language Bindings

Datatype	Description
MPI_INTEGER1	8 bit integer
MPI_INTEGER2	16 bit integer
MPI_INTEGER4	32 bit integer
MPI_INTEGER	32 bit integer
MPI_REAL4	32 bit floating point
MPI_REAL	32 bit floating point
MPI_REAL8	64 bit floating point
MPI_DOUBLE_PRECISION	64 bit floating point
MPI_REAL16	128 bit floating point
MPI_COMPLEX8	32 bit float real, 32 bit float imag.
MPI_COMPLEX	32 bit float real, 32 bit float imag.
MPI_COMPLEX16	64 bit float real, 64 bit float imag.
MPI_DOUBLE_COMPLEX	64 bit float real, 64 bit float imag.
MPI_COMPLEX32	128 bit float real, 128 bit float imag.
MPI_LOGICAL1	8 bit logical
MPI_LOGICAL2	16 bit logical
MPI_LOGICAL4	32 bit logical
MPI_LOGICAL	32 bit logical
MPI_CHARACTER	8 bit character

For Reduction Functions (C Reduction Types)

Datatype	Description
MPI_FLOAT_INT	{MPI_FLOAT, MPI_INT}
MPI_DOUBLE_INT	{MPI_DOUBLE, MPI_INT}
MPI_LONG_INT	{MPI_LONG, MPI_INT}
MPI_2INT	{MPI_INT, MPI_INT}
MPI_SHORT_INT	{MPI_SHORT, MPI_INT}
MPI_LONG_DOUBLE_INT	{MPI_LONG_DOUBLE, MPI_INT}

For Reduction Functions (FORTRAN Reduction Types)

Datatype	Description
MPI_2REAL	{MPI_REAL, MPI_REAL}
MPI_2DOUBLE_PRECISION	{MPI_DOUBLE_PRECISION, MPI_DOUBLE_PRECISION}
MPI_2INTEGER	{MPI_INTEGER, MPI_INTEGER}
MPI_2COMPLEX	{MPI_COMPLEX, MPI_COMPLEX}

Index

© Copyright IBM Corp. 1997