

Sun™ MPI 4.1 Programming and Reference Guide



THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 Fax 650 969-9131

Part No. 806-3734-10
March 2000, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: (c) Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, SunStore, AnswerBook2, docs.sun.com, Solaris, Sun HPC ClusterTools, Prism, Sun Performance WorkShop Fortran, Sun Performance Library, Sun WorkShop Compilers C, Sun WorkShop Compilers C++, Sun WorkShop Compilers Fortran, Sun Visual WorkShop, and UltraSPARC are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303-4900 U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd. La notice suivante est applicable à Netscape Communicator™: (c) Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Solaris, Sun HPC ClusterTools, Prism, Sun Performance WorkShop Fortran, Sun Performance Library, Sun WorkShop Compilers C, Sun WorkShop Compilers C++, Sun WorkShop Compilers Fortran, Sun Visual WorkShop, et UltraSPARC sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Adobe PostScript

Contents

Preface vii

1. Introduction to Sun MPI 1

What Is Sun MPI? 1

Background: The MPI Standard 1

Sun MPI Features 2

The Sun MPI Library 2

MPI I/O 3

Using Sun MPI 3

2. The Sun MPI Library 5

The Libraries 5

Sun MPI Routines 6

Point-to-Point Routines 6

Collective Communication 7

Managing Groups, Contexts, and Communicators 8

Data Types 9

Persistent Communication Requests 10

Managing Process Topologies 10

Environmental Inquiry Functions 11

Programming With Sun MPI	11
Fortran Support	12
Recommendations for All-to-All and All-to-One Communication	12
Signals and MPI	13
Multithreaded Programming	13
Guidelines for Thread-Safe Programming	14
Error Handlers	16
Profiling Interface	16
MPE: Extensions to the Library	18
3. Getting Started	21
Header Files	21
Sample Code	22
Compiling and Linking	24
Choosing a Library Path	26
Stubbing Thread Calls	27
Basic Job Execution	28
Executing With the CRE	28
Executing With LSF Suite 3.2.3	28
Debugging	29
Debugging With the Prism Environment	29
Debugging With MPE	30
4. Programming With Sun MPI I/O	31
Using Sun MPI I/O	31
Data Partitioning and Data Types	32
Definitions	32
Routines	33
Sample Code	41

For More Information	48
A. Sun MPI and Sun MPI I/O Routines	49
Sun MPI Routines	49
Sun MPI I/O Routines	74
B. Environment Variables	83
Informational	83
General Performance Tuning	84
Tuning Memory for Point-to-Point Performance	85
Numerics	87
Tuning Rendezvous	88
Miscellaneous	89
C. Troubleshooting	93
MPI Messages	93
MPI I/O Error Handling	96
D. TNF Probes	99
TNF Probes for MPI	100
TNF Probes for MPI I/O	118

Preface

Sun™ MPI Programming and Reference Guide describes the Sun MPI library of message-passing routines and explains how to develop an MPI (message-passing interface) program on a Sun HPC system. For information about using MPI with the Sun HPC ClusterTools software, see the *Sun HPC ClusterTools User's Guide*.

For the most part, this guide does not repeat information that is available in detail elsewhere; it focuses instead on what is specific to the Sun MPI implementation. References to more general source materials are provided in the section “Related Publications” on page ix of this preface.

The reader is assumed to be familiar with programming in C or Fortran. Some familiarity with parallel programming and with the message-passing model is also required.

Before You Read This Book

For general information about writing MPI programs, refer to any of the several MPI source documents cited in the section “Related Publications” on page ix. Sun MPI is part of the Sun HPC ClusterTools suite of software. For more information about running Sun MPI jobs, see the *Sun HPC ClusterTools User's Guide*. Product notes for Sun MPI are included in *Sun HPC ClusterTools Product Notes*.

Using UNIX® Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- AnswerBook2™ online documentation for the Solaris™ software environment
- Other software documentation that you received with your system

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Unless otherwise specified, examples are presented in C-shell syntax.

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Publications

This book focuses on Sun MPI and assumes familiarity with the MPI standard. The following materials provide useful background about using Sun MPI and about the MPI standard.

Books

Among the documents included with Sun HPC ClusterTools software, you may want to pay particular attention to these:

Application	Title	Part Number
Sun HPC ClusterTools software	<i>Sun HPC ClusterTools Product Notes</i>	806-4182-10
Running Sun MPI jobs	Sun HPC ClusterTools User's Guide	806-3733-10
Scalable Coherent Interface (SCI)	<i>Sun HPC SCI Guide</i>	806-4183-10
Sun HPC ClusterTools performance tuning	<i>Sun HPC ClusterTools Performance Guide</i>	806-3732-10

Application	Title	Part Number
Sun HPC ClusterTools administration	<i>Sun HPC ClusterTools Administrator's Guide</i>	806-3731-10
Prism™ development environment	<i>Prism User's Guide</i>	806-3736-10
Prism development environment	<i>Prism Reference Manual</i>	806-3737-10

In addition, if you are using Platform Computing's Load Sharing Facility (LSF) Suite, you will want to consult Platform's documentation for LSF.

LSF Suite	<i>LSF Batch User's Guide</i>
LSF Suite	<i>LSF Parallel User's Guide</i>
LSF Suite	<i>LSF Batch Programmer's Guide</i>
LSF Suite	<i>LSF Batch User's Quick Reference</i>

These documents are available from Platform Computing. See their web site for more information:

<http://www.platform.com>

The following books, which are not provided by Sun, should be available at your local computer bookstore:

- *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, by William Gropp, Anthony Skjellum, and Ewing Lusk (Cambridge: MIT Press, 1999).
- *Using MPI-2: Advanced Features of the Message-Passing Interface*, by William Gropp, Rajeev Thakur, et al.(Cambridge: MIT Press, 1999).
- *MPI: The Complete Reference, Volume 1, The MPI Core*, by Marc Snir, Steve W. Otto, Steven Huss-Lederman, David W. Walker, and Jack Dongarra (Cambridge, MA: MIT Press, 1999).
- *MPI: The Complete Reference, Volume 2, The MPI Extensions*, by William Gropp, Steven Huss-Lederman, Andrew Lumsdaine, Ewing Lusk, Bill Nitzberg, William Saphir, and Marc Snir (Cambridge, MA: MIT Press, 1999).
- *Parallel Programming with MPI*, by Peter S. Pacheco (San Francisco: Morgan Kaufmann Publishers, Inc., 1997).

Man Pages

Man pages are also available online for all the Sun MPI and MPI I/O routines and are accessible via the Solaris™ `man` command. These man pages are usually installed in `/opt/SUNWhpc/man`. You may need to ask your system administrator for their location at your site.

On the World Wide Web

There is a wealth of documentation on MPI available on the World Wide Web. Here are a few URLs for Web sites:

- The MPI home page, with links to specifications for the MPI-2 standard:

<http://www.mpi-forum.org>

- Additional Web sites that provide links to papers, talks, the standard, implementations, information about MPI-2, plus pointers to many other sources:

<http://www.erc.msstate.edu/mpi/>

<http://www.arc.unm.edu/homepage/TutorialWorkshop.html>

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

<http://www1.fatbrain.com/documentation/sun>

Accessing Sun Documentation Online

The `docs.sun.com`SM web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

<http://docs.sun.com>

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

`docfeedback@sun.com`

Please include the part number (806-3734-10) of your document in the subject line of your email.

Introduction to Sun MPI

What Is Sun MPI?

Sun MPI is Sun Microsystems' implementation of MPI (message-passing interface), the industry-standard specification for writing message-passing programs. Message passing is a programming model that gives the programmer explicit control over interprocess communication.

Background: The MPI Standard

The MPI specification was developed by the MPI Forum, a group of software developers, computer vendors, academics, and computer-science researchers whose goal was to develop a standard for writing message-passing programs that would be efficient, flexible, and portable.

The outcome, known as the MPI Standard, was first published in 1993; its most recent version (MPI-2) was published in July 1997. It was well received, and there are several implementations available publicly.

Sun MPI Features

Sun MPI provides the following features:

- Integration with the Sun Cluster Runtime Environment (CRE).
- Integration with Platform Computing's Load Sharing Facility (LSF) Suite.
- Support for multithreaded programming.
- Seamless use of different network protocols; for example, code compiled on a Sun HPC cluster that has a Scalable Coherent Interface (SCI) network, can be run without change on a cluster that has an ATM network.
- Multiprotocol support such that MPI picks the fastest available medium for each type of connection (such as shared memory, SCI, or ATM).
- Communication via shared memory for fast performance on clusters of SMPs.
- Finely tunable shared memory communication.
- Optimized collectives for symmetric multiprocessors (SMPs) and clusters of SMPs.
- MPI I/O support for parallel file I/O.
- *Prism support* – Users can develop, run, and debug programs in the Prism programming environment.
- *Implicit coscheduling* – The Sun HPC spind daemon enables certain processes of a given MPI job on a shared-memory system to be scheduled at approximately the same time as other related processes. This coscheduling reduces the load on the processors, thus reducing the effect that MPI jobs have on each other.
- Limited support of one-sided communication routines.
- Sun MPI is a dynamic library.
- MPI-2 dynamic support.

Sun MPI and MPI I/O provide full F77, C, and C++ support, as well as Basic F90 support.

The Sun MPI Library

Sun MPI is a library of message-passing routines, including all MPI 1.1-compliant routines and a subset of the MPI 2-compliant routines. Man pages for Sun MPI routines are available online, and the routines are listed in Appendix A. Chapter 2 describes the Sun MPI library.

MPI I/O

File I/O in Sun MPI comprises MPI 2-compliant routines for parallel file I/O. Chapter 4 describes these routines. Their man pages are provided online, and the routines are listed in Appendix A.

Using Sun MPI

The current release of Sun MPI is optimized to run with Sun HPC ClusterTools software using C, C++, FORTRAN 77, or Fortran 90.

To get started developing, executing, and debugging a Sun MPI program, see Chapter 3. The *Sun HPC ClusterTools User's Guide* describes using Sun MPI in more detail.

The Sun MPI Library

This chapter describes the Sun MPI library:

- “The Libraries” on page 5
- “Sun MPI Routines” on page 6
- “Programming With Sun MPI” on page 11
- “Multithreaded Programming” on page 13
- “Profiling Interface” on page 16
- “MPE: Extensions to the Library” on page 18

Note – Sun MPI I/O is described separately, in Chapter 4.

The Libraries

Sun MPI comprises eight MPI libraries: four 32-bit versions and four 64-bit versions:

- *32- and 64-bit libraries* – If you want to take advantage of the 64-bit capabilities of Sun MPI, you must explicitly link to the 64-bit libraries. The 32-bit libraries are the default in each category.

Note – The 64-bit libraries are installed only when the installation system is running Solaris 7 or 8.

- *Thread-safe and non-thread-safe libraries* – For multithreaded programs, the user must link with the thread-safe library in the appropriate category unless the program has only one thread calling MPI. For programs that are not multithreaded, the user can link against either the thread-safe or the default (non-thread-safe) library. However, non-multithreaded programs will have better

performance using the default library, as it does not incur the extra overhead of providing thread-safety. Therefore, you should use the default libraries whenever possible for maximum performance.

- *Standard and trace libraries* – The trace libraries are used to take advantage of Prism’s MPI performance analysis features and to provide enhanced error reporting. For example, they check arguments, the validity of structure pointers, and nonnegative count values, and so on. These libraries are intended for development purposes only, as the overhead involved in their aggressive parameter-checking and probes degrades performance compared with the standard libraries.

The 32-bit libraries are the default, as are the standard (nontrace) libraries within the 32- or 64-bit categories. Within any given category (32- or 64-bit, standard or trace library), the non-thread-safe library is the default. For full information about linking to libraries, see “Compiling and Linking” on page 24.

Sun MPI Routines

This section gives a brief description of the routines in the Sun MPI library. All the Sun MPI routines are listed in Appendix A with brief descriptions and their C syntax. For detailed descriptions of individual routines, see the man pages. For more complete information, see the MPI standard (see “Related Publications” on page ix of the preface).

Point-to-Point Routines

Point-to-point routines include the basic send and receive routines in both blocking and nonblocking forms and in four modes.

A *blocking send* blocks until its message buffer can be written with a new message. A *blocking receive* blocks until the received message is in the receive buffer.

Nonblocking sends and receives differ from blocking sends and receives in that they return immediately and their completion must be waited or tested for. It is expected that eventually nonblocking send and receive calls will allow the overlap of communication and computation.

MPI’s four modes for point-to-point communication are:

- *Standard*, in which the completion of a send implies that the message either is buffered internally or has been received. Users are free to overwrite the buffer that they passed in with any of the blocking send or receive routines, after the routine returns.

- *Buffered*, in which the user guarantees a certain amount of buffering space.
- *Synchronous*, in which rendezvous semantics occur between sender and receiver; that is, a send blocks until the corresponding receive has occurred.
- *Ready*, in which a send can be started only if the matching receive is already posted. The ready mode for sends is a way for the programmer to notify the system that the receive has been posted, so that the underlying system can use a faster protocol if it is available.

Collective Communication

Collective communication routines are blocking routines that involve all processes in a communicator. Collective communication includes broadcasts and scatters, reductions and gathers, all-gathers and all-to-alls, scans, and a synchronizing barrier call.

TABLE 2-1 Collective Communication Routines

<code>MPI_Bcast()</code>	Broadcasts from one process to all others in a communicator.
<code>MPI_Scatter()</code>	Scatters from one process to all others in a communicator.
<code>MPI_Reduce()</code>	Reduces from all to one in a communicator.
<code>MPI_Allreduce()</code>	Reduces, then broadcasts result to all nodes in a communicator.
<code>MPI_Reduce_scatter()</code>	Scatters a vector that contains results across the nodes in a communicator.
<code>MPI_Gather()</code>	Gathers from all to one in a communicator.
<code>MPI_Allgather()</code>	Gathers, then broadcasts the results of the gather in a communicator.
<code>MPI_Alltoall()</code>	Performs a set of gathers in which each process receives a specific result in a communicator.
<code>MPI_Scan()</code>	Scans (parallel prefix) across processes in a communicator.
<code>MPI_Barrier()</code>	Synchronizes processes in a communicator (no data is transmitted).

Many of the collective communication calls have alternative vector forms, with which different amounts of data can be sent to or received from different processes.

The syntax and semantics of these routines are basically consistent with the point-to-point routines (upon which they are built), but there are restrictions to keep them from getting too complicated:

- The amount of data sent must exactly match the amount of data specified by the receiver.
- There is only one mode, a mode analogous to the standard mode of point-to-point routines.

Managing Groups, Contexts, and Communicators

A distinguishing feature of the MPI standard is that it includes a mechanism for creating separate worlds of communication, accomplished through *communicators*, *contexts*, and *groups*.

A *communicator* specifies a group of processes that will conduct communication operations within a specified context without affecting or being affected by operations occurring in other groups or contexts elsewhere in the program. A communicator also guarantees that, within any group and context, point-to-point and collective communication are isolated from each other.

A *group* is an ordered collection of processes. Each process has a rank in the group; the rank runs from 0 to $n-1$. A process can belong to more than one group; its rank in one group has nothing to do with its rank in any other group.

A *context* is the internal mechanism by which a communicator guarantees safe communication space to the group.

At program startup, two default communicators are defined: `MPI_COMM_WORLD`, which has as a process group all the processes of the job; and `MPI_COMM_SELF`, which is equivalent to an identity communicator. The process group that corresponds to `MPI_COMM_WORLD` is not predefined, but can be accessed using `MPI_COMM_GROUP`. One `MPI_COMM_SELF` communicator is defined for each process, each of which has rank zero in its own communicator. For many programs, these are the only communicators needed.

Communicators are of two kinds: *intracommunicators*, which conduct operations within a given group of processes; and *intercommunicators*, which conduct operations between two groups of processes.

Communicators provide a *caching* mechanism, which allows an application to attach attributes to communicators. Attributes can be user data or any other kind of information.

New groups and new communicators are constructed from existing ones. Group constructor routines are local, and their execution does not require interprocessor communication. Communicator constructor routines are collective, and their execution may require interprocess communication.

Note – Users who do not need any communicator other than the default `MPI_COMM_WORLD` communicator — that is, who do not need any sub- or supersets of processes — can simply plug in `MPI_COMM_WORLD` wherever a communicator argument is requested. In these circumstances, users can ignore this section and the associated routines. (These routines can be identified from the listing in Appendix A.)

Data Types

All Sun MPI communication routines have a data type argument. These may be primitive data types, such as integers or floating-point numbers, or they may be user-defined, derived data types, which are specified in terms of primitive types.

Derived data types allow users to specify more general, mixed, and noncontiguous communication buffers, such as array sections and structures that contain combinations of primitive data types.

The basic data types that can be specified for the data-type argument correspond to the basic data types of the host language. Values for the data-type argument for Fortran and the corresponding Fortran types are listed in the following table.

TABLE 2-2 Possible Values for the Data Type Argument for Fortran

MPI Data Type	Fortran Data Type
<code>MPI_INTEGER</code>	<code>INTEGER</code>
<code>MPI_REAL</code>	<code>REAL</code>
<code>MPI_DOUBLE_PRECISION</code>	<code>DOUBLE PRECISION</code>
<code>MPI_COMPLEX</code>	<code>COMPLEX</code>
<code>MPI_LOGICAL</code>	<code>LOGICAL</code>
<code>MPI_CHARACTER</code>	<code>CHARACTER(1)</code>
<code>MPI_DOUBLE_COMPLEX</code>	<code>DOUBLE COMPLEX</code>
<code>MPI_REAL4</code>	<code>REAL*4</code>
<code>MPI_REAL8</code>	<code>REAL*8</code>
<code>MPI_INTEGER2</code>	<code>INTEGER*2</code>
<code>MPI_INTEGER4</code>	<code>INTEGER*4</code>
<code>MPI_BYTE</code>	
<code>MPI_PACKED</code>	

Values for the data-type argument in C and the corresponding C types are listed in the following table.

TABLE 2-3 Possible Values for the Data Type Argument for C

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

The data types `MPI_BYTE` and `MPI_PACKED` have no corresponding Fortran or C data types.

Persistent Communication Requests

Sometimes within an inner loop of a parallel computation, a communication with the same argument list is executed repeatedly. The communication can be slightly improved by using a *persistent* communication request, which reduces the overhead for communication between the process and the communication controller. A persistent request can be thought of as a communication port or “half-channel.”

Managing Process Topologies

Process topologies are associated with communicators; they are optional attributes that can be given to an intracommunicator (not to an intercommunicator).

Recall that processes in a group are ranked from 0 to $n-1$. This linear ranking often reflects nothing of the logical communication pattern of the processes, which may be, for instance, a 2- or 3-dimensional grid. The logical communication pattern is referred to as a *virtual topology* (separate and distinct from any hardware topology). In MPI, there are two types of virtual topologies that can be created: Cartesian (grid) topology and graph topology.

You can use virtual topologies in your programs by taking physical processor organization into account to provide a ranking of processors that optimizes communications.

Environmental Inquiry Functions

Environmental inquiry functions include routines for starting up and shutting down, error-handling routines, and timers.

Few MPI routines may be called before `MPI_Init()` or after `MPI_Finalize()`. Examples include `MPI_Initialized()` and `MPI_Version()`. `MPI_Finalize()` may be called only if there are no outstanding communications involving that process.

The set of errors handled by MPI is dependent upon the implementation. See Appendix C for tables listing the Sun MPI error classes.

Programming With Sun MPI

Although there are about 190 (non-I/O) routines in the Sun MPI library, you can write programs for a wide range of problems using only six routines:

TABLE 2-4 Six Basic MPI Routines

<code>MPI_Init()</code>	Initializes the MPI library.
<code>MPI_Finalize()</code>	Finalizes the MPI library. This includes releasing resources used by the library.
<code>MPI_Comm_size()</code>	Determines the number of processes in a specified communicator.
<code>MPI_Comm_rank()</code>	Determines the rank of calling process within a communicator.
<code>MPI_Send()</code>	Sends a message.
<code>MPI_Recv()</code>	Receives a message.

This set of six routines includes the basic send and receive routines. Programs that depend heavily on collective communication may also include `MPI_Bcast()` and `MPI_Reduce()`.

The functionality of these routines means you can have the benefit of parallel operations without having to learn the whole library at once. As you become more familiar with programming for message passing, you can start learning the more complex and esoteric routines and add them to your programs as needed.

See “Sample Code” on page 22, for two simple Sun MPI code samples, one in C and one in Fortran. See “Sun MPI Routines” on page 49, for a complete list of Sun MPI routines.

Fortran Support

Sun MPI provides basic Fortran support, as described in section 10.2 of the MPI-2 standard. Essentially, Fortran bindings and an `mpif.h` file are provided, as specified in the MPI-1 standard. The `mpif.h` file is valid for both fixed- and free-source form, as specified in the MPI-2 standard.

The MPI interface is known to violate the Fortran standard in several ways, which cause few problems for FORTRAN 77 programs. These standard violations can cause more significant problems for Fortran 90 programs, however, if you do not follow the guidelines recommended in the standard. If you are programming in Fortran, and particularly if you are using Fortran 90, you should consult section 10.2 of the MPI-2 standard for detailed information about basic Fortran support in an MPI implementation.

Recommendations for All-to-All and All-to-One Communication

The Sun MPI library uses the TCP protocol to communicate over a variety of networks. MPI depends on TCP to ensure reliable, correct data flow. TCP’s reliability compensates for unreliability in the underlying network, as the TCP retransmission algorithms will handle any segments that are lost or corrupted. In most cases, this works well with good performance characteristics. However, when doing all-to-all and all-to-one communication over certain networks, a large number of TCP segments may be lost, resulting in poor performance.

You can compensate for this diminished performance over TCP in these ways:

- When writing your own algorithms, avoid flooding one node with a lot of data.

- If you need to do all-to-all or all-to-one communication, use one of the Sun MPI routines to do so. They are implemented in a way that avoids congesting a single node with lots of data. The following routines fall into this category:
 - `MPI_Alltoall()` and `MPI_Alltoallv()` – These have been implemented using a pairwise communication pattern, so that every rank is communicating with only one other rank at a given time.
 - `MPI_Gather()/MPI_Gatherv()` – The root process sends ready-to-send packets to each nonroot-rank process to tell the processes to send their data. In this way, the root process can regulate how much data it is receiving at any one time. Using this ready-to-send method is, however associated with a minor performance cost. For this reason, you can override this method by setting the `MPI_TCPSAFEGATHER` environment variable to 0. (See Appendix B for information about environment variables.)

Signals and MPI

When running the MPI library over TCP, nonfatal `SIGPIPE` signals may be generated. To handle them, the library sets the signal handler for `SIGPIPE` to `ignore`, overriding the default setting (terminate the process). In this way, the MPI library can recover in certain situations. You should therefore avoid changing the `SIGPIPE` signal handler.

The Sun MPI Fortran and C++ bindings are implemented as wrappers on top of the C bindings. The profiling interface is implemented using weak symbols. This means a profiling library need contain only a profiled version of C bindings.

The `SIGPIPE`s may occur when a process first starts communicating over TCP. This happens because the MPI library creates connections over TCP only when processes actually communicate with one another. There are some unavoidable conditions where `SIGPIPE`s may be generated when two processes establish a connection. If you want to avoid any `SIGPIPE`s, set the environment variable `MPI_FULLCONNINIT`, which creates all connections during `MPI_Init()` and avoids any situations which may generate a `SIGPIPE`. For more information about environment variables, see Appendix B.

Multithreaded Programming

When you are linked to one of the thread-safe libraries, Sun MPI calls are thread safe, in accordance with basic tenets of thread safety for MPI mentioned in the MPI-2 specification¹. This means that:

- When two concurrently running threads make MPI calls, the outcome will be as if the calls executed in some order.
- Blocking MPI calls will block the calling thread only. A blocked calling thread will not prevent progress of other runnable threads on the same process, nor will it prevent them from executing MPI calls. Thus, multiple sends and receives are concurrent.

Guidelines for Thread-Safe Programming

Each thread within an MPI process may issue MPI calls; however, threads are not separately addressable. That is, the rank of a send or receive call identifies a process, not a thread, meaning that no order is defined for the case where two threads call `MPI_Recv()` with the same tag and communicator. Such threads are said to be *in conflict*.

If threads within the same application post conflicting communication calls, data races will result. You can prevent such data races by using distinct communicators or tags for each thread.

In general, you will need to adhere to these guidelines:

- You must not have a request serviced by more than one thread. Although you may have an operation posted in one thread and then completed in another, you may not have the operation completed in more than one thread.
- A data type or communicator must not be freed by one thread while it is in use by another thread.
- Once `MPI_Finalize()` has been called, subsequent calls in any thread will fail.
- You must ensure that a sufficient number of lightweight processes (LWPs) are available for your multithreaded program. Failure to do so may degrade performance or even result in deadlock.
- You cannot stub the thread calls in your multithreaded program by omitting the threads libraries in the link line. The `libmpi.so` library automatically calls in the threads libraries, which effectively overrides any stubs.

The following sections describe more specific guidelines that apply for some routines. They also include some general considerations for collective calls and communicator operations that you should be aware of.

1. *Document for a Standard Message-Passing Interface*. Please see the preface of this document for more information about this and other recommended reference material.

```
MPI_Wait(), MPI_Waitall(), MPI_Waitany(),  
MPI_Waitsome()
```

In a program where two or more threads call one of these routines, you must ensure that they are not waiting for the same request. Similarly, the same request cannot appear in the array of requests of multiple concurrent wait calls.

```
MPI_Cancel()
```

One thread must not cancel a request while that request is being serviced by another thread.

```
MPI_Probe(), MPI_Iprobe()
```

A call to `MPI_Probe()` or `MPI_Iprobe()` from one thread on a given communicator should not have a source rank and tags that match those of any other probes or receives on the same communicator. Otherwise, correct matching of message to probe call may not occur.

Collective Calls

Collective calls are matched on a communicator according to the order in which the calls are issued at each processor. All the processes on a given communicator must make the same collective call. You can avoid the effects of this restriction on the threads on a given processor by using a different communicator for each thread.

No process that belongs to the communicator may omit making a particular collective call; that is, none should be left “dangling.”

Communicator Operations

Each of the communicator functions operates simultaneously with each of the noncommunicator functions, regardless of what the parameters are and of whether the functions are on the same or different communicators. However, if you are using multiple instances of the same communicator function on the same communicator, where all parameters are the same, it cannot be determined which threads belong to which resultant communicator. Therefore, when concurrent threads issue such calls, you must assure that the calls are synchronized in such a way that threads in different processes participating in the same communicator operation are grouped together. Do this either by using a different base communicator for each call or by making the calls in single-thread mode before actually using them within the separate threads.

Please note also these special situations:

- If you are using multiple instances of the same function with differing parameters and multiple threads, you must use different communicators. You must not use multiple instances of the same function on the same communicator with other differing parameters.
- When using splits with multiple instances of the same function with the same parameters, but with different threads at the split, you must use different communicators.

For example, suppose you wish to produce several communicators in different sets of threads by performing `MPI_Comm_split()` on some base communicator. To ensure proper, thread-safe operation, you should replicate the base communicator via `MPI_Comm_dup()` (in the root thread or in one thread) and then perform `MPI_Comm_split()` on the resulting duplicate communicators.

- Do not free a communicator in one thread if it is still being used by another thread.

Error Handlers

When an error occurs as a result of an MPI call, the handler may not run on the same thread as the thread that made the error-raising call. In other words, you cannot assume that the error handler will execute in the local context of the thread that made the error-raising call. The error handler may be executed by another thread on the same process, distinct from the one that returns the error code. Therefore, you cannot rely on local variables for error handling in threads; instead, use global variables from the process.

Profiling Interface

The Prism development environment, a component of Sun HPC ClusterTools software, can be used in conjunction with the TNF probes and libraries included with Sun MPI for profiling your code. (TNF is included with the Solaris operating environment.) See Appendix D for information about the TNF probes and “Choosing a Library Path” on page 26 for information about linking to the “trace” or TNF libraries. See the *Prism User's Guide* for more information about the TNF viewer built into the Prism environment.

Sun MPI also meets the requirements of the profiling interface described in Chapter 8 of the MPI-1 Standard. You may write your own profiling library or choose from a number of available profiling libraries, such as those included with the multiprocessing environment (MPE) from Argonne National Laboratory. (See “MPE:

Extensions to the Library” on page 18 for more information.) The *User’s Guide for mpich, a Portable Implementation of MPI*, includes more detailed information about using profiling libraries. For information about this and other MPI- and MPICH-related publications, see “Related Publications” on page ix.

FIGURE 2-1 on page 17 illustrates how the software fits together. In this example, the user is linking against a profiling library that collects information on `MPI_Send()`. No profiling information is being collected for `MPI_Recv()`.

To compile the program, the user’s link line would look like this:

```
# cc ..... -library-name -lmpi
```

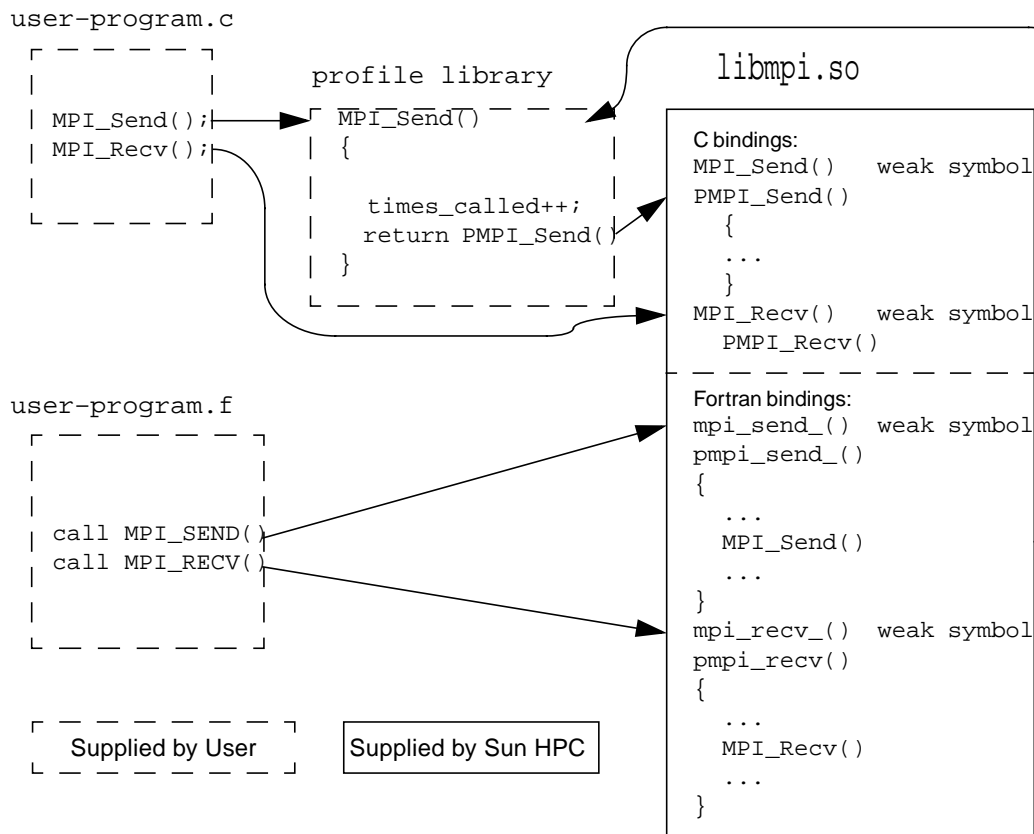


FIGURE 2-1 Sun MPI Profiling Interface

MPE: Extensions to the Library

Although the Sun MPI library does not include or support the multiprocessing environment (MPE) available from Argonne National Laboratory (ANL), it is compatible with MPE. In case you would like to use these extensions to the MPI library, we have included some instructions for downloading it from ANL and building it yourself. Note that these procedures may change if ANL makes changes to MPE.

▼ To Obtain and Build MPE

The MPE software is available from Argonne National Laboratory.

1. Use `ftp` to obtain the file.

```
ftp://ftp.mcs.anl.gov/pub/mpi/misc/mpe.tar.gz
```

The `mpe.tar.gz` file is about 240 Kbytes.

2. Use `gunzip` and `tar` to decompress the software.

```
# gunzip mpe.tar.gz
# tar xvf mpe.tar
```

3. Change your current working directory to the `mpe` directory, and execute `configure` with the arguments shown.

```
# cd mpe
# configure -cc=cc -fc=f77 -opt=-I/opt/SUNWhpc/include
```

4. Execute a `make`.

```
# make
```

This will build several libraries.

Note – Sun MPI does not include the MPE error handlers. You must call the debug routines `MPE_Errors_call_dbx_in_xterm()` and `MPE_Signals_call_debugger()` yourself.

Please refer to the *User's Guide for mpich, a Portable Implementation of MPI*, for information on how to use MPE. It is available at the Argonne National Laboratory web site:

<http://www.mcs.anl.gov/mpi/mpich/>

Getting Started

This chapter describes the rudiments of developing, compiling and linking, executing, and debugging a Sun MPI program. The chapter focuses on what is specific to the Sun MPI implementation and, for the most part, does not repeat information that can be found in related documents. Information about programming with the Sun MPI I/O routines is in Chapter 4.

For complete information about developing MPI programs, see some of the MPI publications listed in the preface. For complete information about executing programs Sun HPC ClusterTools software, see the *Sun HPC ClusterTools User's Guide*.

Header Files

Include syntax must be placed at the top of any program that calls Sun MPI routines.

- For C, use

```
#include <mpi.h>
```
- For C++, use

```
#include <mpi.h>
```
- For Fortran, use

```
INCLUDE 'mpif.h'
```

These lines allow the program to access the Sun MPI version of the `mpi` header file, which contains the definitions, macros, and function prototypes required when compiling the program. Ensure that you are referencing the *Sun MPI* include file.

The include files are usually found in `/opt/SUNWhpc/include/` or `/opt/SUNWhpc/include/v9/`. If the compiler cannot find them, check that they exist and are accessible from the machine on which you are compiling your code. The location of the include file is specified by a compiler option (see “Compiling and Linking” on page 24).

Sample Code

Two simple Sun MPI programs are available in `/opt/SUNWhpc/examples/mpi` and are included here in their entirety. In the same directory you will find the `Readme` file; which provides instructions for using the examples, and the `make` file, `Makefile`.

TABLE 3-1 Simple Sun MPI Program in C: `connectivity.c`

```
/*
 * Test the connectivity between all processes.
 */

#pragma ident "@(#)connectivity.c 1.1 99/02/02"

#include <errno.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <netdb.h>
#include <unistd.h>

#include <mpi.h>

int
main(int argc, char **argv)
{
    MPI_Status  status;
    int         verbose = 0;
    int         rank;
    int         np;           /* number of processes in job */
    int         peer;
    int         i;
    int         j;

    MPI_Init(&argc, &argv);
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &np);

    if (argc>1 && strcmp(argv[1], "-v")==0)
        verbose = 1;
}
```

TABLE 3-1 Simple Sun MPI Program in C: connectivity.c (Continued)

```

for (i=0; i<np; i++) {
    if (rank==i) {
        /* rank i sends to and receives from each higher rank */
        for(j=i+1; j<np; j++) {
            if (verbose)
                printf("checking connection %4d <-> %-4d\n", i, j);
            MPI_Send(&rank, 1, MPI_INT, j, rank, MPI_COMM_WORLD);
            MPI_Recv(&peer, 1, MPI_INT, j, j, MPI_COMM_WORLD, &status);
        }
    } else if (rank>i) {
        /* receive from and reply to rank i */
        MPI_Recv(&peer, 1, MPI_INT, i, i, MPI_COMM_WORLD, &status);
        MPI_Send(&rank, 1, MPI_INT, i, rank, MPI_COMM_WORLD);
    }
}

MPI_Barrier(MPI_COMM_WORLD);
if (rank==0)
    printf("Connectivity test on %d processes PASSED.\n", np);

MPI_Finalize();
return 0;
}

```

TABLE 3-2 Simple Sun MPI Program in Fortran: monte.f

```

!
! Estimate pi via Monte-Carlo method.
!
! Each process sums how many of samplesize random points generated
! in the square (-1,-1),(-1,1),(1,1),(1,-1) fall in the circle of
! radius 1 and center (0,0), and then estimates pi from the formula
! pi = (4 * sum) / samplesize.
! The final estimate of pi is calculated at rank 0 as the average of
! all the estimates.
!
    program monte

        include 'mpif.h'

        double precision drand
        external drand

        double precision x, y, pi, pisum
        integer*4 ierr, rank, np
        integer*4 incircle, samplesize

```

TABLE 3-2 Simple Sun MPI Program in Fortran: `monte.f` (Continued)

```
parameter(samplesize=2000000)

call MPI_INIT(ierr)
call MPI_COMM_RANK(MPI_COMM_WORLD, rank, ierr)
call MPI_COMM_SIZE(MPI_COMM_WORLD, np, ierr)

! seed random number generator
x = drand(2 + 11*rank)

incircle = 0
do i = 1, samplesize
  x = drand(0)*2.0d0 - 1.0d0 ! generate a random point
  y = drand(0)*2.0d0 - 1.0d0

  if ((x*x + y*y) .lt. 1.0d0) then
    incircle = incircle+1 ! point is in the circle
  endif
end do

pi = 4.0d0 * DBLE(incircle) / DBLE(samplesize)

! sum estimates at rank 0
call MPI_REDUCE(pi, pisum, 1, MPI_DOUBLE_PRECISION, MPI_SUM,
& 0, MPI_COMM_WORLD, ierr)

if (rank .eq. 0) then
! final estimate is the average
pi = pisum / DBLE(np)
print '(A,I4,A,F8.6,A)', 'Monte-Carlo estimate of pi by ', np,
& ' processes is ', pi, '.'
endif

call MPI_FINALIZE(ierr)
end
```

Compiling and Linking

Sun MPI programs are compiled with ordinary C, C++, or Fortran compilers, just like any other C, C++, or Fortran program, and linked with the Sun MPI library.

The `mpf77`, `mpf90`, `mpcc`, and `mpCC` utilities may be used to compile Fortran, C, and C++ programs, respectively. For example, you might use

```
% mpf77 -fast -xarch=v8plusa -o a.out a.f -lmpi
```

to compile an f77 program that uses Sun MPI. See the man pages for these routines for more information.

For performance, the single most important compilation switch is `-fast`. This is a macro that expands to settings appropriate for high performance for a general set of circumstances. Because its expansion varies from one compiler release to another, you may prefer to specify the underlying switches explicitly. To see what `-fast` expands to, use `-v` for “verbose” compilation output. Also, `-fast` assumes native compilation, so you should compile on UltraSPARC processors.

The next important compilation switch is `-xarch`. While `-fast` picks many performance-oriented settings by default, optimizations specific to UltraSPARC must be specified explicitly to override certain binary-compatible defaults. Specify

```
-xarch=v8plusa
```

or

```
-xarch=v9a
```

after `-fast` for 32-bit or 64-bit binaries, respectively. To run 64-bit binaries, you must use Solaris 7.

For more information, see the *Sun HPC ClusterTools Performance Guide* and the documents that came with your compiler.

If you will be using the Prism debugger, you must compile your program with Sun WorkShop™ Compilers C/C++ or Sun WorkShop Compilers Fortran, either v4.2 or v5.0. (See “Debugging” on page 29.)

TABLE 3-3 Compile and Link Line Options for Sun MPI and Sun MPI I/O

When using . . .	Use . . .
C (nonthreaded example)	Use <code>mpcc</code> (below), or, if you prefer: <pre>% cc filename.c -o filename \ -I/opt/SUNWhpc/include -L/opt/SUNWhpc/lib \ -R/opt/SUNWhpc/lib -lmpi</pre>
C++ Note that <code>x.y</code> represents the version of your C++ compiler.	Use <code>mpCC</code> (below), or, if you prefer: <pre>% CC filename.cc -o filename \ -I/opt/SUNWhpc/include -L/opt/SUNWhpc/lib \ -R/opt/SUNWhpc/lib -L/opt/SUNWhpc/lib/SCx.y \ -R/opt/SUNWhpc/lib/SCx.y -mt -lmpi++ -lmpi</pre>
<code>mpcc</code> , <code>mpCC</code>	<pre>% mpcc -o filename filename.c -lmpi % mpCC -o filename filename.cc -mt -lmpi</pre>

TABLE 3-3 Compile and Link Line Options for Sun MPI and Sun MPI I/O (Continued)

When using . . .	Use . . .
FORTRAN 77 (nonthreaded example)	Use <code>mpf77</code> (below), or, if you prefer: <pre>% f77 -dalign filename.f -o filename \ -I/opt/SUNWhpc/include -L/opt/SUNWhpc/lib \ -R/opt/SUNWhpc/lib -lmpi</pre>
Fortran on a 64-bit system	<pre>% f77 -dalign filename.f -o filename \ -I/opt/SUNWhpc/include/v9 \ -L/opt/SUNWhpc/lib/sparcv9 \ -R/opt/SUNWhpc/lib/sparcv9 -lmpi</pre>
Fortran 90 <code>mpf77</code> , <code>mpf90</code>	Replace <code>mpf77</code> with <code>mpf90</code> , or <code>f77</code> with <code>f90</code> . <pre>% mpf77 -o -dalign filename filename.f -lmpi % mpf90 -o -dalign filename filename.f -lmpi</pre>
Multithreaded programs and programs containing nonblocking MPI I/O routines	Replace <code>-lmpi</code> with <code>-lmpi_mt</code> .

Note – For the Fortran interface, the `-dalign` option is necessary to avoid the possibility of bus errors. (The underlying C routines in Sun MPI internals assume that parameters and buffer types passed as REALs are double-aligned.)

Note – If your program has previously been linked to any static libraries, you will have to relink it to `libmpi.so` before executing it.

Choosing a Library Path

The eight Sun MPI libraries are described in “The Libraries” on page 5. The paths for each of these libraries, which you must specify when you are compiling and linking your program, are listed in the following table.

TABLE 3-4 Sun MPI Libraries

Category	Description	Path: <code>/opt/SUNWhpc/lib/...</code>
32-Bit Libraries	Default, not thread-safe	<code>libmpi.so</code>
	C++ (in addition to <code>libmpi.so</code>)	<code>SC4.2/libmpi++.so</code>
	Thread-safe	<code>libmpi_mt.so</code>

TABLE 3-4 Sun MPI Libraries (Continued)

Category	Description	Path: /opt/SUNWhpc/lib/...
Trace	Trace, not thread-safe	tnf/libmpi.so
	Trace, thread-safe	tnf/libmpi_mt.so
64-Bit Libraries	Non-thread-safe	sparcv9/libmpi.so
	C++ (in addition to sparcv9/libmpi.so)	SC5.0/libmpi++.so
	Thread-safe	sparcv9/libmpi_mt.so
Trace	Trace, not thread-safe	tnf/sparcv9/libmpi.so
	Trace, thread-safe	tnf/sparcv9/libmpi_mt.so

Overriding the Runtime Library

As shown in the sample compile and link lines in TABLE 3-3 on page 25, you use the `-R` flag in the compile and link line to specify the path for a runtime library when you are compiling. At run time, you can override the library specified in the `-R` argument by setting the `LD_LIBRARY_PATH` environment variable. For example, to link to the 32-bit trace libraries before running your program, do this:

```
% setenv LD_LIBRARY_PATH /opt/SUNWhpc/lib/tnf
```

(This is a C shell example.)

Stubbing Thread Calls

The `libthread.so` libraries are automatically linked into the respective `libmpi.so` libraries. This means that any thread-function calls in your program will be resolved by the `libthread.so` library. Simply omitting `libthread.so` from the link line will not cause thread calls to be stubbed out — you must remove the thread calls yourself. For more information about the `libthread.so` library, see its man page. (For the location of Solaris man pages at your site, see your system administrator.)

Basic Job Execution

The *Sun HPC ClusterTools User's Guide* and the `mprun` man page provide detailed information about running jobs with the CRE. Likewise, the *LSF Batch User's Guide*, the *Sun HPC ClusterTools User's Guide*, and the `lsfintro` and `bsub` man pages provide thorough instructions for executing jobs with the LSF Suite. In this section you will find some basic information about executing jobs with either resource manager.

Before starting your job, you may want to set one or more environment variables, which are also described in Appendix B and in the *Sun HPC ClusterTools Performance Guide*.

Executing With the CRE

Note – Running parallel jobs with the CRE is supported on up to 1024 processors and up to 64 nodes.

When using the CRE, parallel jobs are launched using the `mprun` command. For example, to start a job with 6 processes named `mpijob`, use this command:

```
% mprun -np 6 mpijob
```

Executing With LSF Suite 3.2.3

Note – Running parallel jobs with LSF Suite is supported on up to 1024 processors and up to 64 nodes.

Parallel jobs can either be launched by LSF's Parallel Application Manager (PAM) or be submitted in queues configured to run PAM as the parallel job starter. LSF's `bsub` command launches both parallel interactive and batch jobs. For example, to start a batch job named `mpijob` on four CPUs, use this command:

```
% bsub -n 4 pam mpijob
```

To launch an interactive job, add the `-I` argument to the command line. For example, to launch an interactive job named `earth` on a single CPU in the queue named `sun`, which is configured to launch jobs with PAM):


```
% bsub -q sun -Ip -n 1 earth
```

Debugging

Debugging parallel programs is notoriously difficult, since you are in effect debugging a program potentially made up of many distinct programs executing simultaneously. Even if the application is an SPMD one (single process, multiple data), each instance may be executing a different line of code at any instant. The Prism development environment eases the debugging process considerably and is recommended for debugging with Sun HPC ClusterTools software.

Debugging With the Prism Environment

Note – To run the graphical version of the Prism environment, you must be running the Solaris 2.6, 7, or 8 operating environment with either OpenWindows™ or the Common Desktop Environment (CDE), and with your `DISPLAY` environment variable set correctly. See the *Prism User's Guide* for information.

This section provides a brief introduction to the Prism development environment. For complete information about this environment, see the *Prism User's Guide*.

Prism can debug only one Sun MPI job at a time. Therefore, if an MPI job spawns or connects to another job (using `MPI_Comm_accept()` and `MPI_Comm_connect()` to implement client/server communication, for example, or `MPI_Comm_spawn()` to spawn jobs), the Prism session nonetheless has control of only the original MPI job to which it is attached. For example, a Prism session debugging a server job cannot also debug the clients of that job.

To debug a Sun MPI program in the Prism environment, the program must be written in the SPMD (single process, multiple data) style — that is, all processes that make up a Sun MPI program must be running the same executable. The program must also have been compiled using one of the compilers included in either the Sun Performance WorkShop Fortran or Sun Performance WorkShop C++/C suite of tools.

Note – `MPI_Comm_spawn_multiple()` can create multiple executables with only one job id. You cannot use Prism to debug jobs with different executables that have been spawned with this command.

Starting Up Prism

To start Prism on a Sun MPI program, use the `-np` option to `mprun` to specify how many processes you want to start. For example,

```
% prism -np 4 foo
```

launches Prism on executable `foo` with four processes.

This starts up a graphical version of Prism with your program loaded. You can then debug and visualize data in your Sun MPI program.

You can also attach Prism to running processes. First determine the job id (not the individual process id), or *jid*, using `mpps`. (See the *Sun HPC ClusterTools User's Guide* for further information about `mpps`.) Then specify the *jid* at the command line with the `-np` option:

```
% prism -np 4 foo 12345
```

This will launch Prism and attach it to the processes running in job 12345.

One important feature of the Prism environment is that it lets you debug the Sun MPI program at any level of detail. You can look at the program as a whole or at subsets of processes within the program (for example, those that have an error condition), or at individual processes, all within the same debugging session. For complete information, see the *Prism User's Guide*.

Debugging With MPE

The multiprocessing environment (MPE) available from Argonne National Laboratory includes a debugger that can also be used for debugging at the thread level. For information about obtaining and building MPE, see “MPE: Extensions to the Library” on page 18.

Programming With Sun MPI I/O

File I/O in Sun MPI is fully MPI-2 compliant. MPI I/O is specified as part of that standard, which was published in July, 1997. Its goal is to provide a library of routines featuring a portable parallel file system interface that is an extension of the MPI framework. See “Related Publications” on page ix for more information about the MPI-2 standard.

The closest thing to a standard in file I/O is the UNIX file interface, but UNIX does not provide efficient coordination among multiple simultaneous accesses to a file, particularly when those accesses originate on multiple machines in a cluster. Another drawback of the UNIX file interface is its single-offset interface, that is, its lack of aggregate requests, which can also lead to inefficient access. The MPI I/O library provides routines that accomplish this coordination. Furthermore, MPI I/O allows multiple simultaneous access requests to be made to take advantage of Sun HPC’s parallel file system, PFS. It is currently the only application programming interface through which users can access Sun HPC’s PFS. For more information about PFS, see the *Sun HPC ClusterTools Administrator’s Guide* and the `pfsstat` man page.

Note – A direct interface to Sun HPC’s PFS (parallel file system) is not available to the user in this release. Currently, the only way to access PFS is through Sun’s implementation of MPI I/O or Solaris command-line utilities.

Using Sun MPI I/O

MPI I/O models file I/O on message passing; that is, writing to a file is analogous to sending a message, and reading from a file is analogous to receiving a message. The MPI library provides a high-level way of partitioning data among processes, which

saves you from having to specify the details involved in making sure that the right pieces of data go to the right processes. This section describes basic MPI I/O concepts and the Sun MPI I/O routines.

Data Partitioning and Data Types

MPI I/O uses the MPI model of communicators and derived data types to describe communication between processes and I/O devices. MPI I/O determines which processes are communicating with a particular I/O device. Derived data types can be used to define the layout of data in memory and of data in a file on the I/O device. (For more information about derived data types, see “Data Types” on page 9.) Because MPI I/O builds on MPI concepts, it’s easy for a knowledgeable MPI programmer to add MPI I/O code to a program.

Data is stored in memory and in the file according to MPI data types. Herein lies one of MPI and MPI I/O’s advantages: Because they provide a mechanism whereby you can create your own data types, you have more freedom and flexibility in specifying data layout in memory and in the file.

The library also simplifies the task of describing how your data moves from processor memory to file and back again. You create derived data types that describe how the data is arranged in each process’s memory and how it should be arranged in that process’s part of the disk file.

The Sun MPI I/O routines are described in “Routines” on page 33. But first, to be able to define a data layout, you will need to understand some basic MPI I/O data-layout concepts. The next section explains some of the fundamental terms and concepts.

Definitions

The following terms are used to describe partitioning data among processes. FIGURE 4-1 on page 33 illustrates some of these concepts.

- An *elementary data type* (or `etype`) is the unit of data access and positioning. It can be any MPI basic or derived data type. Data access is performed in elementary-data-type units, and offsets (see below) are expressed as a count of elementary data types.
- The *file type* (or `filetype`) is used to partition a file among processes; that is, a file type defines a template for accessing the file. It is either a single elementary data type or a derived MPI data type constructed from elementary data types. A file type may contain “holes,” or extents of bytes that will not be accessed by this process.

- A file *displacement* (or `disp`) is an absolute byte position relative to the beginning of a file. The displacement defines the location where a view begins (see below).
- A *view* defines the current set of data visible and accessible by a process from an open file in terms of a displacement, an elementary data type, and a file type. The pattern described by a file type is repeated, beginning at the displacement, to define the view.
- An *offset* is a position relative to the current view, expressed as a count of elementary data types. Holes in the view's file type are ignored when calculating this position.

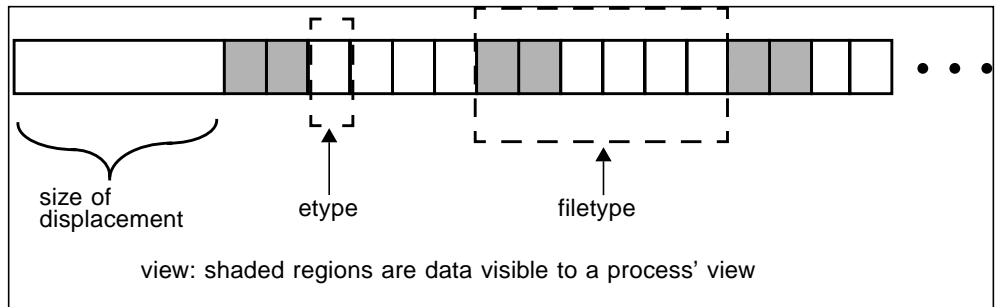


FIGURE 4-1 Displacement, the Elementary Data Type, the File Type, and the View

For a more detailed description of MPI I/O, see Chapter 9, “I/O,” of the MPI-2 standard.

Note for Fortran Users

When writing a Fortran program, you must declare the variable `ADDRESS` as

```
INTEGER*MPI_ADDRESS_KIND ADDRESS
```

`MPI_ADDRESS_KIND` is a constant defined in `mpi.h`. This constant defines the length of the declared integer.

Routines

This release of Sun MPI includes all the MPI I/O routines, which are defined in Chapter 9, “I/O,” of the MPI-2 specification. (See the preface for information about this specification.)

Code samples that use many of these routines are provided in “Sample Code” on page 41.

File Manipulation

Collective coordination	Noncollective coordination
<code>MPI_File_open()</code>	<code>MPI_File_delete()</code>
<code>MPI_File_close()</code>	<code>MPI_File_get_size()</code>
<code>MPI_File_set_size()</code>	<code>MPI_File_get_group()</code>
<code>MPI_File_preallocate()</code>	<code>MPI_File_get_amode()</code>

`MPI_File_open()` and `MPI_File_close()` are collective operations that open and close a file, respectively — that is, all processes in a communicator group must together open or close a file. To achieve a single-user, UNIX-like open, set the communicator to `MPI_COMM_SELF`.

`MPI_File_delete()` deletes a specified file.

The routines `MPI_File_set_size()`, `MPI_File_get_size()`, `MPI_File_get_group()`, and `MPI_File_get_amode()` get and set information about a file. When using the collective routine `MPI_File_set_size()` on a UNIX file, if the size that is set is smaller than the current file size, the file is truncated at the position defined by *size*. If *size* is set to be larger than the current file size, the file size becomes *size*.

When the file size is increased this way with `MPI_File_set_size()`, new regions are created in the file with displacements between the old file size and the larger, newly set file size. Sun MPI I/O does not necessarily allocate file space for such new regions. You may reserve file space either by using `MPI_File_preallocate()` or by performing a read or write to unallocated bytes. `MPI_File_preallocate()` ensures that storage space is allocated for a set quantity of bytes for the specified file; however, its use is very “expensive” in terms of performance and disk space.

The routine `MPI_File_get_group()` returns a communicator group, but it does not free the group.

File Info

Noncollective coordination	Collective coordination
<code>MPI_File_get_info()</code>	<code>MPI_File_set_info()</code>

The opaque `info` object allows you to provide hints for optimization of your code, making it run faster or more efficiently, for example. These hints are set for each file, using the `MPI_File_open()`, `MPI_File_set_view()`, `MPI_File_set_info()`,

and `MPI_File_delete()` routines. `MPI_File_set_info()` sets new values for the specified file's hints. `MPI_File_get_info()` returns all the hints that the system currently associates with the specified file.

When using UNIX files, Sun MPI I/O provides four hints for controlling how much buffer space it uses to satisfy I/O requests: `noncoll_read_bufsize`, `noncoll_write_bufsize`, `coll_read_bufsize`, and `coll_write_bufsize`. These hints may be tuned for your particular hardware configuration and application to improve performance for both noncollective and collective data accesses. For example, if your application uses a single MPI I/O call to request multiple noncontiguous chunks that form a regular strided pattern in the file, you may want to adjust the `noncoll_write_bufsize` to match the size of the stride. Note that these hints limit the size of MPI I/O's underlying buffers but do not limit the size of how much data a user can read or write in a single request.

File Views

Noncollective coordination	Collective coordination
<code>MPI_File_get_view()</code>	<code>MPI_File_set_view()</code>

The `MPI_File_set_view()` routine changes the process's view of the data in the file, specifying its displacement, elementary data type, and file type, as well as setting the individual file pointers and shared file pointer to 0.

`MPI_File_set_view()` is a collective routine; all processes in the group must pass identical values for the file handle and the elementary data type, although the values for the displacement, the file type, and the info object may vary. However, if you use the data-access routines that use file positioning with a shared file pointer, you must also give the displacement and the file type identical values. The data types passed in as the elementary data type and the file type must be committed.

You can also specify the type of data representation for the file. See “File Interoperability” on page 39 for information about registering data representation identifiers.

Note – Displacements within the file type and the elementary data type must be monotonically nondecreasing.

Data Access

The 35 data-access routines are categorized according to file positioning. Data access can be achieved by any of these methods of file positioning:

- By explicit offset
- By individual file pointer
- By shared file pointer

In the following subsections, each of these methods is discussed in more detail.

While *blocking* I/O calls will not return until the request is completed, *nonblocking* calls do not wait for the I/O request to complete. A separate “request complete” call, such as `MPI_Test()` or `MPI_Wait()`, is needed to confirm that the buffer is ready to be used again. Nonblocking routines have the prefix `MPI_File_i`, where the `i` stands for immediate.

All the nonblocking collective routines for data access are “split” into two routines, each with `_begin` or `_end` as a suffix. These *split collective* routines are subject to the semantic rules described in Section 9.4.5 of the MPI-2 standard.

Data Access With Explicit Offsets

Synchronism	Noncollective coordination	Collective coordination
Blocking	<code>MPI_File_read_at()</code> <code>MPI_File_write_at()</code>	<code>MPI_File_read_at_all()</code> <code>MPI_File_write_at_all()</code>
Nonblocking or split collective	<code>MPI_File_iread_at()</code> <code>MPI_File_iwrite_at()</code>	<code>MPI_File_read_at_all_begin()</code> <code>MPI_File_read_at_all_end()</code> <code>MPI_File_write_at_all_begin()</code> <code>MPI_File_write_at_all_end()</code>

To access data at an explicit offset, specify the position in the file where the next data access for each process should begin. For each call to a data-access routine, a process attempts to access a specified number of file types of a specified data type (starting at the specified offset) into a specified user buffer.

The offset is measured in elementary data type units relative to the current view; moreover, “holes” are not counted when locating an offset. The data is read from (in the case of a read) or written into (in the case of a write) those parts of the file specified by the current view. These routines store the number of buffer elements of a particular data type actually read (or written) in the status object, and all the other fields associated with the status object are undefined. The number of elements that are read or written can be accessed using `MPI_Get_count()`.

`MPI_File_read_at()` attempts to read from the file via the associated file handle returned from a successful `MPI_File_open()`. Similarly, `MPI_File_write_at()` attempts to write data from a user buffer to a file. `MPI_File_iread_at()` and `MPI_File_iwrite_at()` are the nonblocking versions of `MPI_File_read_at()` and `MPI_File_write_at()`, respectively.

`MPI_File_read_at_all()` and `MPI_File_write_at_all()` are collective versions of `MPI_File_read_at()` and `MPI_File_write_at()`, in which each process provides an explicit offset. The split collective versions of these nonblocking routines are listed in the table at the beginning of this section.

Data Access With Individual File Pointers

Synchronism	Noncollective coordination	Collective coordination
Blocking	<code>MPI_File_read()</code>	<code>MPI_File_read_all()</code>
	<code>MPI_File_write()</code>	<code>MPI_File_write_all()</code>
Nonblocking or split collective	<code>MPI_File_iread()</code>	<code>MPI_File_read_all_begin()</code>
	<code>MPI_File_iwrite()</code>	<code>MPI_File_read_all_end()</code>
		<code>MPI_File_write_all_begin()</code>
		<code>MPI_File_write_all_end()</code>

For each open file, Sun MPI I/O maintains one individual file pointer per process per collective `MPI_File_open()`. For these data-access routines, MPI I/O implicitly uses the value of the individual file pointer. These routines use and update only the individual file pointers maintained by MPI I/O by pointing to the next elementary data type after the one that has most recently been accessed. The individual file pointer is updated relative to the current view of the file. The shared file pointer is neither used nor updated. (For data access with shared file pointers, please see the next section.)

These routines have similar semantics to the explicit-offset data-access routines, except that the offset is defined here to be the current value of the individual file pointer.

`MPI_File_read_all()` and `MPI_File_write_all()` are collective versions of `MPI_File_read()` and `MPI_File_write()`, with each process using its individual file pointer.

`MPI_File_iread()` and `MPI_File_iwrite()` are the nonblocking versions of `MPI_File_read()` and `MPI_File_write()`, respectively. The split collective versions of `MPI_File_read_all()` and `MPI_File_write_all()` are listed in the table at the beginning of this section.

Pointer Manipulation

`MPI_File_seek`
`MPI_File_get_position`
`MPI_File_get_byte_offset`

Each process can call the routine `MPI_File_seek()` to update its individual file pointer according to the update mode. The update mode has the following possible values:

- `MPI_SEEK_SET` – The pointer is set to the offset.
- `MPI_SEEK_CUR` – The pointer is set to the current pointer position plus the offset.
- `MPI_SEEK_END` – The pointer is set to the end of the file plus the offset.

The offset can be negative for backwards seeking, but you cannot seek to a negative position in the file. The current position is defined as the elementary data item immediately following the last-accessed data item.

`MPI_File_get_position()` returns the current position of the individual file pointer relative to the current displacement and file type.

`MPI_File_get_byte_offset()` converts the offset specified for the current view to the displacement value, or absolute byte position, for the file.

Data Access With Shared File Pointers

Synchronism	Noncollective coordination	Collective coordination
Blocking	<code>MPI_File_read_shared()</code>	<code>MPI_File_read_ordered()</code>
	<code>MPI_File_write_shared()</code>	<code>MPI_File_write_ordered()</code>
		<code>MPI_File_seek_shared()</code>
		<code>MPI_File_get_position_shared()</code>
Nonblocking or split collective	<code>MPI_File_iread_shared()</code>	<code>MPI_File_read_ordered_begin()</code>
	<code>MPI_File_iwrite_shared()</code>	<code>MPI_File_read_ordered_end()</code>
		<code>MPI_File_write_ordered_begin()</code>
		<code>MPI_File_write_ordered_end()</code>

Sun MPI I/O maintains one shared file pointer per collective `MPI_File_open()` (shared among processes in the communicator group that opened the file). As with the routines for data access with individual file pointers, you can also use the current value of the shared file pointer to specify the offset of data accesses implicitly. These routines use and update only the shared file pointer; the individual file pointers are neither used nor updated by any of these routines.

These routines have similar semantics to the explicit-offset data-access routines, except:

- The offset is defined here to be the current value of the shared file pointer.
- Multiple calls (one for each process in the communicator group) affect the shared file pointer routines as if the calls were serialized.
- All processes must use the same file view.

After a shared file pointer operation is initiated, it is updated, relative to the current view of the file, to point to the elementary data item immediately following the last one requested, regardless of the number of items actually accessed.

`MPI_File_read_shared()` and `MPI_File_write_shared()` are blocking routines that use the shared file pointer to read and write files, respectively. The order of serialization is not deterministic for these noncollective routines, so you need to use other methods of synchronization if you wish to impose a particular order.

`MPI_File_iread_shared()` and `MPI_File_iwrite_shared()` are the nonblocking versions of `MPI_File_read_shared()` and `MPI_File_write_shared()`, respectively.

`MPI_File_read_ordered()` and `MPI_File_write_ordered()` are the collective versions of `MPI_File_read_shared()` and `MPI_File_write_shared()`. They must be called by all processes in the communicator group associated with the file handle, and the accesses to the file occur in the order determined by the ranks of the processes within the group. After all the processes in the group have issued their respective calls, for each process in the group, these routines determine the position where the shared file pointer would be after all processes with ranks lower than this process's rank had accessed their data. Then data is accessed (read or written) at that position. The shared file pointer is then updated by the amount of data requested by all processes of the group.

The split collective versions of `MPI_File_read_ordered()` and `MPI_File_write_ordered()` are listed in the table at the beginning of this section.

`MPI_File_seek_shared()` is a collective routine, and all processes in the communicator group associated with the particular file handler must call `MPI_File_seek_shared()` with the same file offset and the same update mode. All the processes are synchronized with a barrier before the shared file pointer is updated.

The offset can be negative for backwards seeking, but you cannot seek to a negative position in the file. The current position is defined as the elementary data item immediately following the last-accessed data item, even if that location is a hole.

`MPI_File_get_position_shared()` returns the current position of the shared file pointer relative to the current displacement and file type.

File Interoperability

`MPI_Register_datarep()`
`MPI_File_get_type_extent()`

Sun MPI I/O supports the basic data representations described in Section 9.5 of the MPI-2 standard:

- *native* – With native representation, data is stored exactly as in memory, in other words, in Solaris/UltraSPARC data representation. This format offers the highest performance and no loss of arithmetic precision. It should be used only in a homogeneous environment, that is, on Solaris/UltraSPARC nodes running Sun ClusterTools software. It may also be used when the MPI application will perform the data type conversions itself.
- *internal* – With internal representation, data is stored in an implementation-dependent format, such as for Sun MPI.
- *external32* – With external32 representation, data is stored in a portable format, prescribed by the MPI-2 and IEEE standards.

These data representations, as well as any user-defined representations, are specified as an argument to `MPI_File_set_view()`.

You may create user-defined data representations with `MPI_Register_datarep()`. Once a data representation has been defined with this routine, you may specify it as an argument to `MPI_File_set_view()`, so that subsequent data-access operations will call the conversion functions specified with `MPI_Register_datarep()`.

If the file data representation is anything but native, you must be careful when constructing elementary data types and file types. For those functions that accept displacements in bytes, the displacements must be specified in terms of their values in the file for the file data representation being used.

`MPI_File_get_type_extent()` can be used to calculate the extents of data types in the file. The extent is the same for all processes accessing the specified file. If the current view uses a user-defined data representation, `MPI_File_get_type_extent()` uses one of the functions specified in setting the data representation to calculate the extent.

File Consistency and Semantics

Noncollective coordination	Collective coordination
<code>MPI_File_get_atomicsity()</code>	<code>MPI_File_set_atomicsity()</code>
	<code>MPI_File_sync()</code>

The routines ending in `_atomicity` allow you to set or query whether a file is in atomic or nonatomic mode. In *atomic mode*, all operations within the communicator group that opens a file are completed as if sequentialized into some serial order. In *nonatomic mode*, no such guarantee is made. In nonatomic mode, `MPI_File_sync()` can be used to assure weak consistency.

The default mode varies with the number of nodes you are using. If you are running a job on a single node, a file is in *nonatomic* mode by default when it is opened. If you are running a job on more than one node, a file is in *atomic* mode by default.

`MPI_File_set_atomicity()` is a collective call that sets the consistency semantics for data-access operations. All the processes in the group must pass identical values for both the file handle and the Boolean flag that indicates whether atomic mode is set.

`MPI_File_get_atomicity()` returns the current consistency semantics for data-access operations. Again, a Boolean flag indicates whether the atomic mode is set.

Note – In some cases, setting atomicity to `false` may provide better performance. The default atomicity value on a cluster is `true`. The lack of synchronization among the distributed caches on a cluster will often prevent your data from completing in the desired state. In these circumstances, you may suffer performance disadvantages with atomicity set to `true`, especially when the data accesses overlap.

Sample Code

In this section, we give some sample code to get you started with programming your I/O using Sun MPI. We start with an example that shows how a parallel job can partition file data among its processes. Next we explore how you can adapt our initial example to use a broad range of other I/O programming styles supported by Sun MPI I/O. Finally, we present a sample code that illustrates the use of the nonblocking MPI I/O routines.

Before we start, remember that MPI I/O is part of MPI, so you must call `MPI_Init()` before calling any MPI I/O routines and `MPI_Finalize()` at the end of your program, even if you only use MPI I/O routines.

Partitioned Writing and Reading in a Parallel Job

MPI I/O was designed to enable processes in a parallel job to request multiple data items that are noncontiguous within a file. Typically, a parallel job partitions file data among the processes.

One method of partitioning a file is to derive the offset at which to access data from the rank of the process. The rich set of MPI derived types also allows us to easily partition file data. For example, we could create an MPI vector type as the filetype passed into `MPI_File_set_view()`. Since vector types do not end with a hole, a call must be made, either to `MPI_Type_create_resized()` or to `MPI_Type_ub()`, to complete the partition. This call extends the extent to include

holes at the end of the type for processes with higher ranks. We create a partitioned file by passing different displacements to `MPI_File_set_view()`. Each of these displacements would be derived from the process' rank. Consequently, offsets would not need to be derived from the ranks because only the data in that process' portion of the partition would be visible in that process' view.

In the following example, we use the first method where we derive the file offsets directly from the process' rank. Each process writes and reads `NUM_INTS` integers starting at the offset `rank * NUM_INTS`. We pass an explicit offset to our MPI I/O data-access routines `MPI_File_write_at()` and `MPI_File_read_at()`. We call `MPI_Get_elements()` to find out how many elements were written or read. To verify that the write was successful, we compare the data written and read as well as set up an `MPI_Barrier()` before calling `MPI_File_get_size()` to verify that the file is the size that we expect upon completion of all the processes' writes.

Observe that we called `MPI_File_set_view()` to set our view of the file as essentially an array of integers instead of the UNIX-like view of the file as an array of bytes. Thus, the offsets that we pass to `MPI_File_write_at()` and `MPI_File_read_at()` are indices into an array of integers and not a byte offset.

CODE EXAMPLE 4-1 Example code in which each process writes and reads `NUM_INTS` integers to a file using `MPI_File_write_at()` and `MPI_File_read_at()`, respectively.

```
/* wr_at.c
 *
 * Example to demonstrate use of MPI_File_write_at and MPI_File_read_at
 *
 */

#include <stdio.h>
#include "mpi.h"

#define NUM_INTS 100

void sample_error(int error, char *string)
{
    fprintf(stderr, "Error %d in %s\n", error, string);
    MPI_Finalize();
    exit(-1);
}

void
main( int argc, char **argv )
{
    char filename[128];
    int i, rank, comm_size;
    int *buff1, *buff2;
    MPI_File fh;
```

CODE EXAMPLE 4-1 Example code in which each process writes and reads `NUM_INTS` integers to a file using `MPI_File_write_at()` and `MPI_File_read_at()`, respectively. *(Continued)*

```
MPI_Offset disp, offset, file_size;
MPI_Datatype etype, ftype, buftype;
MPI_Info info;
MPI_Status status;
int result, count, differs;

if(argc < 2) {
    fprintf(stdout, "Missing argument: filename\n");
    exit(-1);
}
strcpy(filename, argv[1]);

MPI_Init(&argc, &argv);

/* get this processor's rank */
result = MPI_Comm_rank(MPI_COMM_WORLD, &rank);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_Comm_rank");

result = MPI_Comm_size(MPI_COMM_WORLD, &comm_size);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_Comm_size");

/* communicator group MPI_COMM_WORLD opens file "foo"
   for reading and writing (and creating, if necessary) */
result = MPI_File_open(MPI_COMM_WORLD, filename,
    MPI_MODE_RDWR | MPI_MODE_CREATE, (int)NULL, &fh);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_File_open");

/* Set the file view which tiles the file type MPI_INT, starting
   at displacement 0. In this example, the etype is also MPI_INT. */
disp = 0;
etype = MPI_INT;
ftype = MPI_INT;
info = (MPI_Info)NULL;
result = MPI_File_set_view(fh, disp, etype, ftype, (char *)NULL, info);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_File_set_view");

/* Allocate and initialize a buffer (buff1) containing NUM_INTS integers,
   where the integer in location i is set to i. */
buff1 = (int *)malloc(NUM_INTS*sizeof(int));
for(i=0;i<NUM_INTS;i++) buff1[i] = i;

/* Set the buffer type to also be MPI_INT, then write the buffer (buff1)
```

CODE EXAMPLE 4-1 Example code in which each process writes and reads NUM_INTS integers to a file using MPI_File_write_at() and MPI_File_read_at(), respectively. (Continued)

```
    starting at offset 0, i.e., the first etype in the file. */
    buftype = MPI_INT;
    offset = rank * NUM_INTS;
    result = MPI_File_write_at(fh, offset, buff1, NUM_INTS, buftype, &status);
    if(result != MPI_SUCCESS)
        sample_error(result, "MPI_File_write_at");

    result = MPI_Get_elements(&status, MPI_BYTE, &count);
    if(result != MPI_SUCCESS)
        sample_error(result, "MPI_Get_elements");
    if(count != NUM_INTS*sizeof(int))
        fprintf(stderr, "Did not write the same number of bytes as requested\n");
    else
        fprintf(stdout, "Wrote %d bytes\n", count);

    /* Allocate another buffer (buff2) to read into, then read NUM_INTS
       integers into this buffer. */
    buff2 = (int *)malloc(NUM_INTS*sizeof(int));
    result = MPI_File_read_at(fh, offset, buff2, NUM_INTS, buftype, &status);
    if(result != MPI_SUCCESS)
        sample_error(result, "MPI_File_read_at");

    /* Find out how many bytes were read and compare to how many
       we expected */
    result = MPI_Get_elements(&status, MPI_BYTE, &count);
    if(result != MPI_SUCCESS)
        sample_error(result, "MPI_Get_elements");
    if(count != NUM_INTS*sizeof(int))
        fprintf(stderr, "Did not read the same number of bytes as requested\n");
    else
        fprintf(stdout, "Read %d bytes\n", count);

    /* Check to see that each integer read from each location is
       the same as the integer written to that location. */
    differs = 0;
    for(i=0; i<NUM_INTS; i++) {
        if(buff1[i] != buff2[i]) {
            fprintf(stderr, "Integer number %d differs\n", i);
            differs = 1;
        }
    }
    if(!differs)
        fprintf(stdout, "Wrote and read the same data\n");

    MPI_Barrier(MPI_COMM_WORLD);
```


CODE EXAMPLE 4-1 Example code in which each process writes and reads `NUM_INTS` integers to a file using `MPI_File_write_at()` and `MPI_File_read_at()`, respectively. (Continued)

```
result = MPI_File_get_size(fh, &file_size);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_File_get_size");

/* Compare the file size with what we expect */
/* To see a negative response, make the file preexist with a larger
   size than what is written by this program */
if(file_size != (comm_size * NUM_INTS * sizeof(int)))
    fprintf(stderr, "File size is not equal to the write size\n");

result = MPI_File_close(&fh);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_File_close");

MPI_Finalize();

free(buff1);
free(buff2);
}
```

Data Access Styles

We can adapt our example above to support the I/O programming style that best suits our application. Essentially, there are three dimensions on which to choose an appropriate data access routine for your particular task: file pointer type, collective or noncollective, and blocking or nonblocking.

We need to choose which file pointer type to use: explicit, individual, or shared. In the example above, we used an explicit pointer and passed it directly as the offset parameter to the `MPI_File_write_at()` and `MPI_File_read_at()` routines. Using an explicit pointer is equivalent to calling `MPI_File_seek()` to set the individual file pointer to offset, then calling `MPI_File_write()` or `MPI_File_read()`, which is directly analogous to calling UNIX `lseek()` and `write()` or `read()`. If each process accesses the file sequentially, individual file pointers save you the effort to recalculate offset for each data access. We would use a shared file pointer in situations where all the processes need to cooperatively access a file in a sequential way, for example, writing log files.

Collective data-access routines allow the user to enforce some implicit coordination among the processes in a parallel job when making data accesses. For example, if a parallel job alternately reads in a matrix and performs computation on it, but cannot progress to the next stage of computation until all processes have completed the last stage, then a coordinated effort between processes when accessing data might be more efficient. In the example above, we could easily append the suffix `_all` to `MPI_File_write_at()` and `MPI_File_read_at()` to make the accesses

collective. By coordinating the processes, we could achieve greater efficiency in the MPI library or at the file system level in buffering or caching the next matrix. In contrast, noncollective accesses are used when it is not evident that any benefit would be gained by coordinating disparate accesses by each process. UNIX file accesses are noncollective.

Overlapping I/O With Computation and Communication

MPI I/O also supports nonblocking versions of each of the data-access routines, that is, the data-access routines that have the letter *i* before write or read in the routine name (*i* stands for immediate). By definition, nonblocking I/O routines return immediately after the I/O request has been issued and does not wait until the I/O request has been completed. This functionality allows the user to perform computation and communication at the same time as the I/O. Since large I/O requests can take a long time to complete, this provides a way to more efficiently utilize your programs waiting time.

As in our example above, parallel jobs often partition large matrices stored in files. These parallel jobs may use many large matrices or matrices that are too large to fit into memory at once. Thus, each process may access the multiple and/or large matrices in stages. During each stage, a process reads in a chunk of data, then performs some computation on it (which may involve communicating with the other processes in the parallel job). While performing the computation and communication, the process could issue a nonblocking I/O read request for the next chunk of data. Similarly, once the computation on a particular chunk has completed, a nonblocking write request could be issued before performing computation and communication on the next chunk.

The following example code illustrates the use of a nonblocking data-access routine. Notice that, like nonblocking communication routines, the nonblocking I/O routines require a call to `MPI_Wait()` to wait for the nonblocking request to complete or repeated calls to `MPI_Test()` to determine when the nonblocking data access has completed. Once complete, the write or read buffer is available for use again by the program.

CODE EXAMPLE 4-2 Example code in which each process reads and writes `NUM_BYTES` bytes to a file using the nonblocking MPI I/O routines `MPI_File_iread_at()` and `MPI_File_iwrite_at()`, respectively. Note the use of `MPI_Wait()` and `MPI_Test()` to determine when the nonblocking requests have completed.

```
/* iwr_at.c
 *
 * Example to demonstrate use of MPI_File_iwrite_at and MPI_File_iread_at
 *
 */
```

CODE EXAMPLE 4-2 Example code in which each process reads and writes NUM_BYTES bytes to a file using the nonblocking MPI I/O routines MPI_File_iread_at() and MPI_File_iwrite_at(), respectively. Note the use of MPI_Wait() and MPI_Test() to determine when the nonblocking requests have completed. (Continued)

```
#include <stdio.h>
#include "mpi.h"

#define NUM_BYTES 100

void sample_error(int error, char *string)
{
    fprintf(stderr, "Error %d in %s\n", error, string);
    MPI_Finalize();
    exit(-1);
}

void
main( int argc, char **argv )
{
    char filename[128];
    char *buff;
    MPI_File fh;
    MPI_Offset offset;
    MPI_Request request;
    MPI_Status status;
    int i, rank, flag, result;

    if(argc < 2) {
        fprintf(stdout, "Missing argument: filename\n");
        exit(-1);
    }
    strcpy(filename, argv[1]);

    MPI_Init(&argc, &argv);

    result = MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    if(result != MPI_SUCCESS)
        sample_error(result, "MPI_Comm_rank");

    result = MPI_File_open(MPI_COMM_WORLD, filename,
        MPI_MODE_RDWR | MPI_MODE_CREATE,
        (MPI_Info)NULL, &fh);
    if(result != MPI_SUCCESS)
        sample_error(result, "MPI_File_open");

    buff = (char *)malloc(NUM_BYTES*sizeof(char));
    for(i=0;i<NUM_BYTES;i++) buff[i] = i;

    offset = rank * NUM_BYTES;
```

CODE EXAMPLE 4-2 Example code in which each process reads and writes NUM_BYTES bytes to a file using the nonblocking MPI I/O routines MPI_File_iread_at() and MPI_File_iwrite_at(), respectively. Note the use of MPI_Wait() and MPI_Test() to determine when the nonblocking requests have completed. (Continued)

```
result = MPI_File_iread_at(fh, offset, buff, NUM_BYTES,
    MPI_BYTE, &request);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_File_iread_at");

/* Perform some useful computation and/or communication */

result = MPI_Wait(&request, &status);

buff = (char *)malloc(NUM_BYTES*sizeof(char));
for(i=0;i<NUM_BYTES;i++) buff[i] = i;
result = MPI_File_iwrite_at(fh, offset, buff, NUM_BYTES,
    MPI_BYTE, &request);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_File_iwrite_at");

/* Perform some useful computation and/or communication */

flag = 0;
i = 0;
while(!flag) {
    result = MPI_Test(&request, &flag, &status);
    i++;
    /* Perform some more computation or communication, if possible */
}

result = MPI_File_close(&fh);
if(result != MPI_SUCCESS)
    sample_error(result, "MPI_File_close");

MPI_Finalize();

fprintf(stdout, "Successful completion\n");

free(buff);
}
```

For More Information

For more information on MPI I/O, refer to the documents listed in the section “Related Publications” on page ix of the preface.

Sun MPI and Sun MPI I/O Routines

The tables in this appendix list the routines for the Sun MPI and Sun MPI I/O routines, along with the C syntax of the routines and a brief description of each. For more information about the routines, see their online man pages, usually found in

`/opt/SUNWhpc/man`

Your system administrator can tell you where they are installed at your site.

Sun MPI Routines

TABLE A-1 on page 59 lists the Sun MPI routines in alphabetical order. The following sections list the routines by functional category.

Point-to-Point Communication

Blocking Routines

```
MPI_Send()  
MPI_Bsend()  
MPI_Ssend()  
MPI_Rsend()  
MPI_Recv()  
MPI_Sendrecv()  
MPI_Sendrecv_replace()
```

Nonblocking Routines

```
MPI_Isend()  
MPI_Ibsend()  
MPI_Issend()  
MPI_Irsend()  
MPI_Irecv()
```

Communication Buffer Allocation

```
MPI_Buffer_attach()  
MPI_Buffer_detach()
```

Status Data Structure

```
MPI_Get_count()  
MPI_Get_elements()
```

Persistent (Half-Channel) Communication

```
MPI_Send_init()  
MPI_Bsend_init()  
MPI_Rsend_init()  
MPI_Ssend_init()  
MPI_Recv_init()  
MPI_Start()  
MPI_Startall()
```

Completion Tests

```
MPI_Wait()  
MPI_Waitany()  
MPI_Waitsome()  
MPI_Waitall()  
MPI_Test()  
MPI_Testany()  
MPI_Testsome()  
MPI_Testall()  
MPI_Request_free()  
MPI_Cancel()  
MPI_Test_cancelled()
```

Probing for Messages (Blocking/Nonblocking)

```
MPI_Probe()  
MPI_Iprobe()
```

Packing and Unpacking Functions

```
MPI_Pack()  
MPI_Pack_size()  
MPI_Unpack()
```

Derived Data Type Constructors and Functions

`MPI_Address()`: *Deprecated* - Use `MPI_Get_address()`.
`MPI_Type_commit()`
`MPI_Type_contiguous()`
`MPI_Type_create_indexed_block()`
`MPI_Type_create_keyval()`
`MPI_Type_delete_attr()`
`MPI_Type_dup()`
`MPI_Type_free_keyval()`
`MPI_Type_get_attr()`
`MPI_Type_set_attr()`
`MPI_Type_get_contents()`
`MPI_Type_get_envelope()`
`MPI_Type_get_name()`
`MPI_Type_set_name()`
`MPI_Type_create_resized()`
`MPI_Type_free()`
`MPI_Type_get_true_extent()`
`MPI_Type_hvector()`: *Deprecated* - Use `MPI_Type_create_hvector()`.
`MPI_Type_indexed()`
`MPI_Type_hindexed()`: *Deprecated* - Use `MPI_Type_create_hindexed()`.
`MPI_Type_struct()`: *Deprecated* - Use `MPI_Type_create_struct()`.
`MPI_Type_lb()`: *Deprecated* - Use `MPI_Type_get_extent()`.
`MPI_Type_ub()`: *Deprecated* - Use `MPI_Type_get_extent()`.
`MPI_Type_vector()`
`MPI_Type_extent()`: *Deprecated* - Use `MPI_Type_get_extent()`.
`MPI_Type_size()`

Collective Communication

Barrier

`MPI_Barrier()`

Broadcast

`MPI_Bcast()`

Processor Gather/Scatter

```
MPI_Gather()  
MPI_Gatherv()  
MPI_Allgather()  
MPI_Allgatherv()  
MPI_Scatter()  
MPI_Scatterv()  
MPI_Alltoall()  
MPI_Alltoallv()
```

Global Reduction/Scan Operations

```
MPI_Reduce()  
MPI_Allreduce()  
MPI_Reduce_scatter()  
MPI_Scan()  
MPI_Op_create()  
MPI_Op_free()
```

Groups, Contexts, and Communicators

Group Management

Group Accessors

```
MPI_Group_size()  
MPI_Group_rank()  
MPI_Group_translate_ranks()  
MPI_Group_compare()
```

Group Constructors

```
MPI_Comm_group()  
MPI_Group_union()  
MPI_Group_intersection()  
MPI_Group_difference()  
MPI_Group_incl()  
MPI_Group_excl()  
MPI_Group_range_incl()  
MPI_Group_range_excl()  
MPI_Group_free()
```

Communicator Management

Communicator Accessors

```
MPI_Comm_size()  
MPI_Comm_rank()  
MPI_Comm_compare()
```

Communicator Constructors

```
MPI_Comm_dup()  
MPI_Comm_create()  
MPI_Comm_split()  
MPI_Comm_free()
```

Intercommunicators

```
MPI_Comm_test_inter()  
MPI_Comm_remote_group()  
MPI_Comm_remote_size()  
MPI_Intercomm_create()  
MPI_Intercomm_merge()
```

Communicator Attributes

`MPI_Keyval_create()`: *Deprecated* - Use `MPI_Comm_create_keyval()`.
`MPI_Keyval_free()`: *Deprecated* - Use `MPI_Comm_free_keyval()`.
`MPI_Attr_put()`: *Deprecated* - Use `MPI_Comm_set_attr()`.
`MPI_Attr_get()`: *Deprecated* - Use `MPI_Comm_get_attr()`.
`MPI_Attr_delete()`: *Deprecated* - Use `MPI_Comm_delete_attr()`.

Process Topologies

`MPI_Cart_create()`
`MPI_Dims_create()`
`MPI_Graph_create()`
`MPI_Topo_test()`
`MPI_Graphdims_get()`
`MPI_Graph_get()`
`MPI_Cartdim_get()`
`MPI_Cart_get()`
`MPI_Cart_rank()`
`MPI_Cart_coords()`
`MPI_Graph_neighbors()`
`MPI_Graph_neighbors_count()`
`MPI_Cart_shift()`
`MPI_Cart_sub()`
`MPI_Cart_map()`
`MPI_Graph_map()`

Process Creation and Management

Establishing Communication

`MPI_Close_port()`
`MPI_Comm_accept()`
`MPI_Comm_connect()`
`MPI_Comm_disconnect()`
`MPI_Open_port()`

Process Manager Interface

```
MPI_Comm_get_parent()  
MPI_Comm_spawn()  
MPI_Comm_spawn_multiple()
```

Environmental Inquiry Functions and Profiling

Startup and Shutdown

```
MPI_Init()  
MPI_Finalize()  
MPI_Finalized()  
MPI_Initialized()  
MPI_Abort()  
MPI_Get_processor_name()  
MPI_Get_version()
```

Error Handler Functions

```
MPI_Errhandler_create(): Deprecated - Use  
MPI_Comm_create_errhandler().  
MPI_Errhandler_set(): Deprecated - Use MPI_Comm_set_errhandler().  
MPI_Errhandler_get(): Deprecated - Use MPI_Comm_get_errhandler().  
MPI_Errhandler_free()  
MPI_Error_string()  
MPI_Error_class()
```

Info Objects

```
MPI_Info_create()  
MPI_Info_delete()  
MPI_Info_dup()  
MPI_Info_free()  
MPI_Info_get()  
MPI_Info_get_nkeys()  
MPI_Info_get_nthkey()  
MPI_Info_get_valuelen()  
MPI_Info_set()
```

Timers

```
MPI_Wtime()  
MPI_Wtick()
```

Profiling

```
MPI_Pcontrol()
```

Miscellaneous

Associating Information With Status

```
MPI_Status_set_cancelled()  
MPI_Status_set_elements()
```

Generalized Requests

```
MPI_Grequest_complete()  
MPI_Grequest_start()
```

Naming Objects

```
MPI_Comm_get_name()  
MPI_Comm_set_name()  
MPI_Type_get_name()  
MPI_Type_set_name()
```

Threads

```
MPI_Query_thread()
```

Handle Translation

```
MPI_Comm_c2f()  
MPI_Comm_f2c()  
MPI_Group_c2f()  
MPI_Group_f2c()  
MPI_Info_c2f()  
MPI_Info_f2c()  
MPI_Op_c2f()  
MPI_Op_f2c()  
MPI_Request_c2f()  
MPI_Request_f2c()  
MPI_Type_c2f()  
MPI_Type_f2c()
```

Status Conversion

```
MPI_Status_c2f()  
MPI_Status_f2c()
```

MPI Routines: Alphabetical Listing

TABLE A-1 Sun MPI Routines

Routine and C Syntax	Description
MPI_Abort (MPI_Comm <i>comm</i> , int <i>errorcode</i>)	Terminates MPI execution environment.
MPI_Accumulate (void * <i>origin_addr</i> , int <i>origin_count</i> , MPI_Datatype <i>origin_datatype</i> , int <i>target_rank</i> , MPI_Aint <i>target_disp</i> , int <i>target_count</i> , MPI_Datatype <i>target_datatype</i> , MPI_Op <i>op</i> , MPI_Win <i>win</i>)	Combines the contents of the origin buffer with that of a target buffer.
MPI_Address (void * <i>location</i> , MPI_Aint * <i>address</i>)	<i>Deprecated:</i> Use instead <code>MPI_Get_address()</code> . Gets the address of a location in memory.
MPI_Allgather (void * <i>sendbuf</i> , int <i>sendcount</i> , MPI_Datatype <i>sendtype</i> , void * <i>recvbuf</i> , int <i>recvcount</i> , MPI_Datatype <i>recvtype</i> , MPI_Comm <i>comm</i>)	Gathers data from all processes and distributes it to all.
MPI_Allgatherv (void * <i>sendbuf</i> , int <i>sendcount</i> , MPI_Datatype <i>sendtype</i> , void * <i>recvbuf</i> , int * <i>recvcount</i> , int * <i>displs</i> , MPI_Datatype <i>recvtype</i> , MPI_Comm <i>comm</i>)	Gathers data from all processes and delivers it to all. Each process may contribute a different amount of data.
MPI_Alloc_mem (MPI_Aint <i>size</i> , MPI_Info <i>info</i> , void * <i>baseptr</i>)	Allocates a specified memory segment.
MPI_Allreduce (void * <i>sendbuf</i> , void * <i>recvbuf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Op <i>op</i> , MPI_Comm <i>comm</i>)	Combines values from all processes and distributes the result back to all processes.
MPI_Alltoall (void * <i>sendbuf</i> , int <i>sendcount</i> , MPI_Datatype <i>sendtype</i> , void * <i>recvbuf</i> , int <i>recvcount</i> , MPI_Datatype <i>recvtype</i> , MPI_Comm <i>comm</i>)	Sends data from all to all processes.
MPI_Alltoallv (void * <i>sendbuf</i> , int * <i>sendcounts</i> , int * <i>sdispls</i> , MPI_Datatype <i>sendtype</i> , void * <i>recvbuf</i> , int * <i>recvcounts</i> , int * <i>rdispls</i> , MPI_Datatype <i>recvtype</i> , MPI_Comm <i>comm</i>)	Sends data from all to all processes, with a displacement. Each process may contribute a different amount of data.
MPI_Attr_delete (MPI_Comm <i>comm</i> , int <i>keyval</i>)	<i>Deprecated:</i> Use instead <code>MPI_Comm_delete_attr()</code> . Deletes attribute value associated with a key.
MPI_Attr_get (MPI_Comm <i>comm</i> , int <i>keyval</i> , void * <i>attribute_val</i> , int * <i>flag</i>)	<i>Deprecated:</i> Use instead <code>MPI_Comm_get_attr()</code> . Retrieves attribute value by key.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Attr_put (MPI_Comm <i>comm</i> , int <i>keyval</i> , void * <i>attribute_val</i>)	<i>Deprecated:</i> Use instead <code>MPI_Comm_set_attr()</code> . Stores attribute value associated with a key.
MPI_Barrier (MPI_Comm <i>comm</i>)	Blocks until all processes have reached this routine.
MPI_Bcast (void * <i>buffer</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , int <i>root</i> , MPI_Comm <i>comm</i>)	Broadcasts a message from the process with rank <code>root</code> to all other processes of the group.
MPI_Bsend (void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , int <i>dest</i> , int <i>tag</i> , MPI_Comm <i>comm</i>)	Basic send with user-specified buffering.
MPI_Bsend_init (void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , int <i>dest</i> , int <i>tag</i> , MPI_Comm <i>comm</i> , MPI_Request * <i>request</i>)	Builds a handle for a buffered send.
MPI_Buffer_attach (void * <i>buf</i> , int <i>size</i>)	Attaches a user-defined buffer for sending.
MPI_Buffer_detach (void * <i>buf</i> , int * <i>size</i>)	Removes an existing buffer (for use in <code>MPI_Bsend()</code> , etc.).
MPI_Cancel (MPI_Request * <i>request</i>)	Cancels a communication request.
MPI_Cart_coords (MPI_Comm <i>comm</i> , int <i>rank</i> , int <i>maxdims</i> , int * <i>coords</i>)	Determines process coordinates in Cartesian topology given rank in group.
MPI_Cart_create (MPI_Comm <i>comm_old</i> , int <i>ndims</i> , int * <i>dims</i> , int * <i>periods</i> , int <i>reorder</i> , MPI_Comm * <i>comm_cart</i>)	Makes a new communicator to which Cartesian topology information has been attached.
MPI_Cart_get (MPI_Comm <i>comm</i> , int <i>maxdims</i> , int * <i>dims</i> , int * <i>periods</i> , int * <i>coords</i>)	Retrieves Cartesian topology information associated with a communicator.
MPI_Cart_map (MPI_Comm <i>comm</i> , int <i>ndims</i> , int * <i>dims</i> , int * <i>periods</i> , int * <i>newrank</i>)	Maps process to Cartesian topology information.
MPI_Cart_rank (MPI_Comm <i>comm</i> , int * <i>coords</i> , int * <i>rank</i>)	Determines process rank in communicator given Cartesian location.
MPI_Cart_shift (MPI_Comm <i>comm</i> , int <i>direction</i> , int <i>disp</i> , int * <i>rank_source</i> , int * <i>rank_dest</i>)	Returns the shifted source and destination ranks, given a shift direction and amount.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Cart_sub (MPI_Comm <i>comm</i> , int <i>*remain_dims</i> , MPI_Comm <i>*comm_new</i>)	Partitions a communicator into subcommunicators, which form lower-dimensional Cartesian subgrids.
MPI_Cartdim_get (MPI_Comm <i>comm</i> , int <i>*ndims</i>)	Retrieves Cartesian topology information associated with a communicator.
MPI_Close_port (char <i>*port_name</i>)	Releases the specified network address.
MPI_Comm_accept (char <i>*port_name</i> , MPI_Info <i>info</i> , int <i>root</i> , MPI_Comm <i>comm</i> , MPI_Comm <i>*newcomm</i>)	Establishes communication with a client (collective).
MPI_Comm_c2f (MPI_Comm <i>comm</i>)	Translates a C handle into a Fortran handle.
MPI_Comm_compare (MPI_Comm <i>comm1</i> , MPI_Comm <i>comm2</i> , int <i>*result</i>)	Compares two communicators.
MPI_Comm_connect (char <i>*port_name</i> , MPI_Info <i>info</i> , int <i>root</i> , MPI_Comm <i>comm</i> , MPI_Comm <i>*newcomm</i>)	Establishes communication with a server (collective).
MPI_Comm_create (MPI_Comm <i>comm</i> , MPI_Group <i>group</i> , MPI_Comm <i>*newcomm</i>)	Creates a new communicator from a group.
MPI_Comm_create_errhandler (MPI_Comm_errhandler_fn <i>*function</i> , MPI_Errhandler <i>*errhandler</i>)	Creates an error handler that can be attached to communicators.
MPI_Comm_create_keyval (MPI_Comm_copy_attr_function <i>*comm_copy_attr_fn</i> , MPI_Comm_delete_attr_function <i>*comm_delete_attr_fn</i> , int <i>*comm_keyval</i> , void <i>*extra_state</i>)	Generates a new attribute key.
MPI_Comm_delete_attr (MPI_Comm <i>comm</i> , int <i>comm_keyval</i>)	Deletes attribute value associated with a key.
MPI_Comm_disconnect (MPI_Comm <i>*comm</i>)	Deallocates communicator object and sets handle to MPI_COMM_NULL (collective).
MPI_Comm_dup (MPI_Comm <i>comm</i> , MPI_Comm <i>*newcomm</i>)	Duplicates an existing communicator with all its cached information.
MPI_Comm_f2c (MPI_Fint <i>comm</i>)	Translates a Fortran handle into a C handle.

TABLE A-1 Sun MPI Routines (*Continued*)

Routine and C Syntax	Description
MPI_Comm_free (MPI_Comm *comm)	Marks the communicator object for deallocation.
MPI_Comm_free_keyval (int *comm_keyval)	Frees attribute key for communicator cache attribute.
MPI_Comm_get_attr (MPI_Comm comm, int comm_keyval, void *attribute_val, int *flag)	Retrieves attribute value by key.
MPI_Comm_get_errhandler (MPI_Comm comm, MPI_Errhandler *errhandler)	Retrieves error handler associated with a communicator.
MPI_Comm_get_name (MPI_Comm comm, char *comm_name, int *resultlen)	Returns the name that was most recently associated with a communicator.
MPI_Comm_get_parent (MPI_Comm *parent)	Returns the parent intercommunicator of current spawned process.
MPI_Comm_group (MPI_Comm comm, MPI_Group *group)	Accesses the group associated with a communicator.
MPI_Comm_rank (MPI_Comm comm, int *rank)	Determines the rank of the calling process in a communicator.
MPI_Comm_remote_group (MPI_Comm comm, MPI_Group *group)	Accesses the remote group associated with an intercommunicator.
MPI_Comm_remote_size (MPI_Comm comm, int size)	Determines the size of the remote group associated with an intercommunicator.
MPI_Comm_set_attr (MPI_Comm comm, int comm_keyval, void *attribute_val)	Stores attribute value associated with a key.
MPI_Comm_set_errhandler (MPI_Comm comm, MPI_Errhandler *errhandler)	Attaches a new error handler to a communicator.
MPI_Comm_set_name (MPI_Comm comm, char *comm_name)	Associates a name with a communicator.
MPI_Comm_size (MPI_Comm comm, int *size)	Determines the size of the group associated with a communicator.
MPI_Comm_spawn (char *command, char *argv[], int maxprocs, MPI_Info info, int root, MPI_Comm comm, MPI_Comm *intercomm, int array_of_errcodes[])	Spawns a number of identical binaries.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Comm_spawn_multiple (int <i>count</i> , char * <i>array_of_commands</i> [], char ** <i>array_of_argv</i> [], int <i>array_of_maxprocs</i> [], MPI_Info <i>array_of_info</i> [], int <i>root</i> , MPI_Comm <i>comm</i> , MPI_Comm * <i>intercomm</i> , int <i>array_of_errcodes</i> [])	Spawns multiple binaries, or the same binary with multiple sets of arguments.
MPI_Comm_split (MPI_Comm <i>comm</i> , int <i>color</i> , int <i>key</i> , MPI_Comm * <i>newcomm</i>)	Creates new communicators based on colors and keys.
MPI_Comm_test_inter (MPI_Comm <i>comm</i> , int * <i>flag</i>)	Tests whether a communicator is an intercommunicator.
MPI_Dims_create (int <i>nnodes</i> , int <i>ndims</i> , int * <i>dims</i>)	Creates a division of processors in a Cartesian grid.
MPI_Errhandler_create (MPI_Handler_function * <i>function</i> , MPI_Errhandler * <i>errhandler</i>)	<i>Deprecated:</i> Use instead <code>MPI_Comm_create_errhandler()</code> . Creates an MPI error handler.
MPI_Errhandler_free (MPI_Errhandler * <i>errhandler</i>)	Frees an MPI error handler.
MPI_Errhandler_get (MPI_Comm <i>comm</i> , MPI_Errhandler * <i>errhandler</i>)	<i>Deprecated:</i> Use instead <code>MPI_Comm_get_errhandler()</code> . Gets the error handler for a communicator.
MPI_Errhandler_set (MPI_Comm <i>comm</i> , MPI_Errhandler <i>errhandler</i>)	<i>Deprecated:</i> Use instead <code>MPI_Comm_set_errhandler()</code> . Sets the error handler for a communicator.
MPI_Error_class (int <i>errorcode</i> , int * <i>errorclass</i>)	Converts an error code into an error class.
MPI_Error_string (int <i>errorcode</i> , char * <i>string</i> , int * <i>resultlen</i>)	Returns a string for a given error code.
MPI_Finalize ()	Terminates MPI execution environment.
MPI_Finalized (int * <i>flag</i>)	Checks whether <code>MPI_Finalize()</code> has completed.
MPI_Free_mem (void * <i>base</i>)	Frees memory that has been allocated using <code>MPI_Alloc_mem</code> .
MPI_Gather (void * <i>sendbuf</i> , int * <i>sendcount</i> , MPI_Datatype <i>sendtype</i> , void * <i>recvbuf</i> , int <i>recvcount</i> , MPI_Datatype <i>recvtype</i> , int <i>root</i> , MPI_Comm <i>comm</i>)	Gathers values from a group of processes.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Gatherv (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcounts, int *displs, MPI_Datatype recvtpe, int root, MPI_Comm comm)	Gathers into specified locations from all processes in a group. Each process may contribute a different amount of data.
MPI_Get (void *origin_addr, int origin_count, MPI_Datatype origin_datatype, int target_rank, MPI_Aint target_disp, int target_count, MPI_Datatype target_datatype, MPI_Win win)	Copies data from the target memory to the origin.
MPI_Get_address (void *location, MPI_Aint *address)	Gets the address of a location in memory.
MPI_Get_count (MPI_Status *status, MPI_Datatype datatype, int *count)	Gets the number of top-level elements received.
MPI_Get_elements (MPI_Status *status, MPI_Datatype datatype, int *count)	Returns the number of basic elements in a data type.
MPI_Get_processor_name (char *name, int *resultlen)	Gets the name of the processor.
MPI_Get_version (int *version, int *subversion)	Returns the version of the standard corresponding to the current implementation.
MPI_Graph_create (MPI_Comm comm_old, int nnodes, int *index, int *edges, int reorder, MPI_Comm *comm_graph)	Makes a new communicator to which graph topology information has been attached.
MPI_Graph_get (MPI_Comm comm, int maxindex, int maxedges, int *index, int *edges)	Retrieves graph topology information associated with a communicator.
MPI_Graph_map (MPI_Comm comm, int nnodes, int *index, int *edges, int *newrank)	Maps process to graph topology information.
MPI_Graph_neighbors (MPI_Comm comm, int rank, int maxneighbors, int *neighbors)	Returns the neighbors of a node associated with a graph topology.
MPI_Graph_neighbors_count (MPI_Comm comm, int rank, int *neighbors)	Returns the number of neighbors of a node associated with a graph topology.
MPI_Graphdims_get (MPI_Comm comm, int *nnodes, int *nedges)	Retrieves graph topology information associated with a communicator.
MPI_Grequest_complete (MPI_Request request)	Reports that a generalized request is complete.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Grequest_start (MPI_Grequest_query_function * <i>query_fn</i> , MPI_Grequest_free_function * <i>free_fn</i> , MPI_Grequest_cancel_function * <i>cancel_fn</i> , void * <i>extra_state</i> , MPI_Request * <i>request</i>)	Starts a generalized request and returns a handle to it.
MPI_Group_c2f (MPI_Group <i>group</i>)	Translates a C handle into a Fortran handle.
MPI_Group_compare (MPI_Group <i>group1</i> , MPI_Group <i>group2</i> , int * <i>result</i>)	Compares two groups.
MPI_Group_difference (MPI_Group <i>group1</i> , MPI_Group <i>group2</i> , MPI_Group * <i>group_out</i>)	Makes a group from the difference of two groups.
MPI_Group_excl (MPI_Group <i>group</i> , int <i>n</i> , int * <i>ranks</i> , MPI_Group * <i>newgroup</i>)	Produces a group by reordering an existing group and taking only unlisted members.
MPI_Group_f2c (MPI_Fint <i>group</i>)	Translates a Fortran handle into a C handle.
MPI_Group_free (MPI_Group * <i>group</i>)	Frees a group.
MPI_Group_incl (MPI_Group <i>group</i> , int <i>n</i> , int * <i>ranks</i> , MPI_Group * <i>group_out</i>)	Produces a group by reordering an existing group and taking only listed members.
MPI_Group_intersection (MPI_Group <i>group1</i> , MPI_Group <i>group2</i> , MPI_Group * <i>group_out</i>)	Produces a group at the intersection of two existing groups.
MPI_Group_range_excl (MPI_Group <i>group</i> , int <i>n</i> , int <i>ranges</i> [][3], MPI_Group * <i>newgroup</i>)	Produces a group by excluding ranges of processes from an existing group.
MPI_Group_range_incl (MPI_Group <i>group</i> , int <i>n</i> , int <i>ranges</i> [][3], MPI_Group * <i>newgroup</i>)	Creates a new group from ranges of ranks in an existing group.
MPI_Group_rank (MPI_Group <i>group</i> , int * <i>rank</i>)	Returns the rank of this process in the given group.
MPI_Group_size (MPI_Group <i>group</i> , int * <i>size</i>)	Returns the size of a group.
MPI_Group_translate_ranks (MPI_Group <i>group1</i> , int <i>n</i> , int * <i>ranks1</i> , MPI_Group <i>group2</i> , int * <i>ranks2</i>)	Translates the ranks of processes in one group to those in another group.
MPI_Group_union (MPI_Group <i>group1</i> , MPI_Group <i>group2</i> , MPI_Group * <i>group_out</i>)	Produces a group by combining two groups.
MPI_IbSEND (void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , int <i>dest</i> , int <i>tag</i> , MPI_Comm <i>comm</i> , MPI_Request * <i>request</i>)	Starts a nonblocking buffered send.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Info_c2f (MPI_Info <i>info</i>)	Translates a C handle into a Fortran handle.
MPI_Info_create (MPI_Info * <i>info</i>)	Creates a new info object.
MPI_Info_delete (MPI_Info * <i>info</i> , char * <i>key</i> , char * <i>value</i>)	Deletes a key/value pair from info.
MPI_Info_dup (MPI_Info <i>info</i> , MPI_Info * <i>newinfo</i>)	Duplicates an info object.
MPI_Info_f2c (MPI_Fint <i>info</i>)	Translates a Fortran handle into a C handle.
MPI_Info_free (MPI_Info * <i>info</i>)	Frees info and sets it to MPI_INFO_NULL.
MPI_Info_get (MPI_Info * <i>info</i> , char * <i>key</i> , char * <i>value</i>)	Retrieves key value for an info object.
MPI_Info_get_nkeys (MPI_Info <i>info</i> , int * <i>nkeys</i>)	Returns the number of currently defined keys in info.
MPI_Info_get_nthkey (MPI_Info <i>info</i> , int <i>n</i> , char * <i>key</i>)	Returns the <i>n</i> th defined key in info.
MPI_Info_get_valuelen (MPI_Info <i>info</i> , char * <i>key</i> , int * <i>valuelen</i> , int * <i>flag</i>)	Retrieves the length of the key value associated with an info object.
MPI_Info_set (MPI_Info * <i>info</i> , char * <i>key</i> , char * <i>value</i>)	Adds a key/value pair to info.
MPI_Init (int * <i>argc</i> , char *** <i>argv</i>)	Initializes the MPI execution environment.
MPI_Initialized (int * <i>flag</i>)	Indicates whether MPI_Init() has been called.
MPI_Intercomm_create (MPI_Comm <i>local_comm</i> , int <i>local_leader</i> , MPI_Comm <i>peer_comm</i> , int <i>remote_leader</i> , int <i>tag</i> , MPI_Comm * <i>newintercomm</i>)	Creates an intercommunicator.
MPI_Intercomm_merge (MPI_Comm <i>intercomm</i> , int <i>high</i> , MPI_Comm * <i>newintracomm</i>)	Creates an intracommunicator from an intercommunicator.
MPI_Iprobe (int <i>source</i> , int <i>tag</i> , MPI_Comm <i>comm</i> , int * <i>flag</i> , MPI_Status * <i>status</i>)	Nonblocking test for a message.
MPI_Irecv (void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , int <i>source</i> , int <i>tag</i> , MPI_Comm <i>comm</i> , MPI_Request * <i>request</i>)	Begins a nonblocking receive.
MPI_Irsend (void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , int <i>dest</i> , int <i>tag</i> , MPI_Comm <i>comm</i> , MPI_Request * <i>request</i>)	Begins a nonblocking ready send.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Isend (void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm, MPI_Request *request)	Begins a nonblocking send.
MPI_Issend (void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm, MPI_Request *request)	Begins a nonblocking synchronous send.
MPI_Keyval_create (MPI_Copy_function *copy_fn, MPI_Delete_function *delete_fn, int *keyval, void *extra_state)	<i>Deprecated:</i> Use instead <code>MPI_Comm_create_keyval()</code> . Generates a new attribute key.
MPI_Keyval_free (int *keyval)	<i>Deprecated:</i> Use instead <code>MPI_Comm_free_keyval()</code> . Frees attribute key for communicator cache attribute.
MPI_Op_c2f (MPI_Op op)	Translates a C handle into a Fortran handle.
MPI_Op_create (MPI_User_function *function, int commute, MPI_Op *op)	Creates a user-defined combination function handle.
MPI_Op_f2c (MPI_Fint op)	Translates a Fortran handle into a C handle.
MPI_Op_free (MPI_Op *op)	Frees a user-defined combination function handle.
MPI_Open_port (MPI_Info info, char *port_name)	Establishes a network address for a server to accept connections from clients.
MPI_Pack (void *inbuf, int incount, MPI_Datatype datatype, void *outbuf, int outsize, int *position, MPI_Comm comm)	Packs data of a given data type into contiguous memory.
MPI_Pack_size (int incount, MPI_Datatype datatype, MPI_Comm comm, int *size)	Returns the upper bound on the amount of space needed to pack a message.
MPI_Pcontrol (int level, ...)	Controls profiling.
MPI_Probe (int source, int tag, MPI_Comm comm, MPI_Status *status)	Blocking test for a message.
MPI_Put (void *origin_addr, int origin_count, MPI_Datatype origin_datatype, int target_rank, MPI_Aint target_disp, int target_count, MPI_Datatype target_datatype, MPI_Win win)	Copies data from the origin memory to the target.
MPI_Query_thread (int *provided)	Returns the current level of thread support.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Recv (void *buf, int count, MPI_Datatype datatype, int source, int tag, MPI_Comm comm, MPI_Status *status)	Performs a standard receive.
MPI_Recv_init (void *buf, int count, MPI_Datatype datatype, int source, int tag, MPI_Comm comm, MPI_Request *request)	Builds a persistent receive request handle.
MPI_Reduce (void *sendbuf, void *recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)	Reduces values on all processes to a single value.
MPI_Reduce_scatter (void *sendbuf, void *recvbuf, int *recvcounts, MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)	Combines values and scatters the results.
MPI_Request_c2f (MPI_Request request)	Translates a C handle into a Fortran handle.
MPI_Request_f2c (MPI_Fint request)	Translates a Fortran handle into a C handle.
MPI_Request_free (MPI_Request *request)	Frees a communication request object.
MPI_Request_get_status (MPI_Request request, int *flag, MPI_Status *status)	Accesses information associated with a request without freeing the request.
MPI_Rsend (void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm)	Performs a ready send.
MPI_Rsend_init (void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm, MPI_Request *request)	Builds a persistent ready send request handle.
MPI_Scan (void *sendbuf, void *recvbuf, int count, MPI_Datatype datatype, MPI_Op op, MPI_Comm comm)	Computes the scan (partial reductions) of data on a collection of processes.
MPI_Scatter (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm)	Sends data from one job to all other processes in a group.
MPI_Scatterv (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm)	Scatters a buffer in parts to all processes in a group.
MPI_Send (int *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm)	Performs a standard send.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Send_init (void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm, MPI_Request *request)	Builds a persistent send request handle.
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)	Sends and receives two messages at the same time.
MPI_Sendrecv_replace (void *buf, int count, MPI_Datatype datatype, int dest, int sendtag, int source, int recvtag, MPI_Comm comm, MPI_Status *status)	Sends and receives using a single buffer.
MPI_Ssend (void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm)	Performs a synchronous send.
MPI_Ssend_init (void *buf, int count, MPI_Datatype datatype, int dest, int tag, MPI_Comm comm, MPI_Request *request)	Builds a persistent synchronous send request handle.
MPI_Start (MPI_Request *request)	Initiates a communication using a persistent request handle.
MPI_Startall (int count, MPI_Request array_of_requests[])	Starts a collection of requests.
MPI_Status_c2f (MPI_Status *c_status, MPI_Fint *f_status)	Translates a C status into a Fortran status.
MPI_Status_f2c (MPI_Fint *f_status, MPI_Status *c_status)	Translates a Fortran status into a C status.
MPI_Status_set_cancelled (MPI_Status *status, int flag)	Sets <i>status</i> to indicate a request has been cancelled.
MPI_Status_set_elements (MPI_Status *status, MPI_Datatype datatype, int count)	Modifies opaque part of <i>status</i> to allow <code>MPI_Get_elements()</code> to return <i>count</i> .
MPI_Test (MPI_Request *request, int *flag, MPI_Status *status)	Tests for the completion of a send or receive.
MPI_Test_cancelled (MPI_Status *status, int *flag)	Tests whether a request was canceled.
MPI_Testall (int count, MPI_Request array_of_requests, int *flag, MPI_Status *array_of_statuses)	Tests for the completion of all of the given communications.
MPI_Testany (int count, MPI_Request array_of_requests[], int *index, int *flag, MPI_Status status)	Tests for completion of any of the given communications.

TABLE A-1 Sun MPI Routines (*Continued*)

Routine and C Syntax	Description
MPI_Testsome (int <i>incount</i> , MPI_Request <i>array_of_requests</i> [], int <i>*outcount</i> , int <i>*array_of_indices</i> , MPI_Status <i>*array_of_statuses</i>)	Tests for some given communications to complete.
MPI_Topo_test (MPI_Comm <i>comm</i> , int <i>*top_type</i>)	Determines the type of topology (if any) associated with a communicator.
MPI_Type_c2f (MPI_Datatype <i>datatype</i>)	Translates a C handle into a Fortran handle.
MPI_Type_commit (MPI_Datatype <i>*datatype</i>)	Commits a data type.
MPI_Type_contiguous (int <i>count</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates a contiguous data type.
MPI_Type_create_darray (int <i>size</i> , int <i>rank</i> , int <i>ndims</i> , int <i>array_of_gsizes</i> [], int <i>array_of_distrib</i> [], int <i>array_of_dargs</i> [], int <i>array_of_psizes</i> [], int <i>order</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates an array of data types.
MPI_Type_create_hindexed (int <i>count</i> , int <i>array_of_blocklengths</i> , MPI_Aint <i>array_of_displacements</i> [], MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates an indexed data type with offsets in bytes.
MPI_Type_create_hvector (int <i>count</i> , int <i>blocklength</i> , MPI_Aint <i>stride</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates a vector (strided) data type with offset in bytes.
MPI_Type_create_indexed_block (int <i>count</i> , int <i>blocklength</i> , int <i>array_of_displacements</i> [], MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates an indexed block.
MPI_Type_create_keyval (MPI_Type_copy_attr_function <i>*type_copy_attr_fn</i> , MPI_Type_delete_attr_function <i>*type_delete_attr_fn</i> , int <i>*type_keyval</i> , void <i>*extra_state</i>)	Generates a new attribute key.
MPI_Type_create_resized (MPI_Datatype <i>oldtype</i> , MPI_Aint <i>lb</i> , MPI_Aint <i>extent</i> , MPI_Datatype <i>*newtype</i>)	Returns a new data type with new extent and upper and lower bounds.
MPI_Type_create_struct (int <i>count</i> , int <i>array_of_blocklengths</i> [], MPI_Aint <i>array_of_displacements</i> [], MPI_Datatype <i>array_of_types</i> [], MPI_Datatype <i>*newtype</i>)	Creates a struct data type.
MPI_Type_create_subarray (int <i>ndims</i> , int <i>array_of_sizes</i> [], int <i>array_of_subsizes</i> [], int <i>array_of_starts</i> [], int <i>order</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates a data type describing a subarray of an array.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Type_delete_attr (MPI_Datatype <i>type</i> , int <i>type_keyval</i>)	Deletes attribute value associated with a key.
MPI_Type_dup (MPI_Datatype <i>type</i> , MPI_Datatype <i>*newtype</i>)	Duplicates a data type with associated key values.
MPI_Type_extent (MPI_Datatype <i>datatype</i> , MPI_Aint <i>*extent</i>)	<i>Deprecated:</i> Use instead <code>MPI_Type_get_extent()</code> . Returns the extent of a data type, the difference between the upper and lower bounds of the data type.
MPI_Type_f2c (MPI_Fint <i>datatype</i>)	Translates a Fortran handle into a C handle.
MPI_Type_free (MPI_Datatype <i>*datatype</i>)	Frees a data type.
MPI_Type_free_keyval (int <i>*type_keyval</i>)	
MPI_Type_get_attr (MPI_Datatype <i>type</i> , int <i>type_keyval</i> , void <i>*attribute_val</i> , int <i>*flag</i>)	Returns the attribute associated with a data type.
MPI_Type_get_contents (MPI_Datatype <i>datatype</i> , int <i>max_integers</i> , int <i>max_addresses</i> , int <i>max_datatypes</i> , int <i>array_of_integers[]</i> , MPI_Aint <i>array_of_addresses[]</i> , MPI_Datatype <i>array_of_datatypes[]</i>)	Returns information about arguments used in creation of a data type.
MPI_Type_get_envelope (MPI_Datatype <i>datatype</i> , int <i>*num_integers</i> , int <i>*num_addresses</i> , int <i>*num_datatypes</i> , int <i>*combiner</i>)	Returns information about input arguments associated with a data type.
MPI_Type_get_extent (MPI_Datatype <i>datatype</i> , MPI_Aint <i>*lb</i> , MPI_Aint <i>*extent</i>)	Returns the lower bound and extent of a data type.
MPI_Type_get_name (MPI_Datatype <i>type</i> , char <i>*type_name</i> , int <i>*resultlen</i>)	Gets the name of a data type.
MPI_Type_get_true_extent (MPI_Datatype <i>datatype</i> , MPI_Aint <i>*true_lb</i> , MPI_Aint <i>*true_extent</i>)	Returns the true lower bound and extent of a data type's corresponding typemap, ignoring <code>MPI_UB</code> and <code>MPI_LB</code> markers.
MPI_Type_hindexed (int <i>count</i> , int <i>*array_of_blocklengths</i> , MPI_Aint <i>*array_of_displacements</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	<i>Deprecated:</i> Use instead <code>MPI_Type_create_hindexed()</code> . Creates an indexed data type with offsets in bytes.
MPI_Type_hvector (int <i>count</i> , int <i>blocklength</i> , MPI_Aint <i>stride</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	<i>Deprecated:</i> Use instead <code>MPI_Type_create_hvector()</code> . Creates a vector (strided) data type with offset in bytes.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Type_indexed (int <i>count</i> , int <i>*array_of_blocklengths</i> , int <i>*array_of_displacements</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates an indexed data type.
MPI_Type_lb (MPI_Datatype <i>datatype</i> , MPI_Aint <i>*displacement</i>)	<i>Deprecated:</i> Use instead <code>MPI_Type_get_extent()</code> . Returns the lower bound of a data type.
MPI_Type_set_attr (MPI_Datatype <i>type</i> , int <i>type_keyval</i> , void <i>*attribute_val</i>)	Stores attribute value associated with a key.
MPI_Type_set_name (MPI_Comm <i>comm</i> , char <i>*type_name</i>)	Sets the name of a data type.
MPI_Type_size (MPI_Datatype <i>datatype</i> , int <i>*size</i>)	Returns the number of bytes occupied by entries in the data type.
MPI_Type_struct (int <i>count</i> , int <i>*array_of_blocklengths</i> , MPI_Aint <i>*array_of_displacements</i> , MPI_Datatype <i>*array_of_types</i> , MPI_Datatype <i>*newtype</i>)	<i>Deprecated:</i> Use instead <code>MPI_Type_create_struct()</code> . Creates a struct data type.
MPI_Type_ub (MPI_Datatype <i>datatype</i> , MPI_Aint <i>*displacement</i>)	<i>Deprecated:</i> Use instead <code>MPI_Type_get_extent()</code> . Returns the upper bound of a data type.
MPI_Type_vector (int <i>count</i> , int <i>blocklength</i> , int <i>stride</i> , MPI_Datatype <i>oldtype</i> , MPI_Datatype <i>*newtype</i>)	Creates a vector (strided) data type.
MPI_Unpack (void <i>*inbuf</i> , int <i>insize</i> , int <i>*position</i> , void <i>*outbuf</i> , int <i>outcount</i> , MPI_Datatype <i>datatype</i> , MPI_Comm <i>comm</i>)	Unpacks a data type into contiguous memory.
MPI_Wait (MPI_Request <i>*request</i> , MPI_Status <i>*status</i>)	Waits for an MPI send or receive to complete.
MPI_Waitall (int <i>count</i> , MPI_Request <i>array_of_requests</i> [], MPI_Status <i>array_of_statuses</i> [])	Waits for all of the given communications to complete.
MPI_Waitany (int <i>count</i> , MPI_Request <i>array_of_requests</i> [], int <i>*index</i> , MPI_Status <i>*status</i>)	Waits for any of the given communications to complete.
MPI_Waitsome (int <i>incount</i> , MPI_Request <i>array_of_requests</i> [], int <i>*outcount</i> , int <i>array_of_indices</i> [], MPI_Status <i>array_of_statuses</i> [])	Waits for some given communications to complete.
MPI_Win_c2f (MPI_Win <i>win</i>)	Translates a C handle into a Fortran handle.

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
MPI_Win_create (void *base, MPI_Aint size, int disp_unit, MPI_Info info, MPI_Comm comm, MPI_Win *win)	
MPI_Win_create_errhandler (MPI_Win_errhandler_fn *function, MPI_Errhandler *errhandler)	
MPI_Win_create_keyval (MPI_Win_copy_attr_function *win_copy_attr_fn, MPI_Win_delete_attr_function *win_delete_attr_fn, int *win_keyval, void *extra_state)	
MPI_Win_delete_attr (MPI_Win win, int win_keyval)	
MPI_Win_f2c (MPI_Fint win)	Translates a Fortran handle into a C handle.
MPI_Win_fence (int assert, MPI_Win win)	
MPI_Win_free (MPI_Win *win)	Frees the window object and returns a null handle.
MPI_Win_free_keyval (int *win_keyval)	
MPI_Win_get_attr (MPI_Win win, int win_keyval, void *attribute_val, int *flag)	
MPI_Win_get_errhandler (MPI_Win win, MPI_Errhandler *errhandler)	Retrieves the error handler currently associated with a window.
MPI_Win_get_group (MPI_Win win, MPI_Group *group)	Returns a duplicate of the group of the communicator used to create the window.
MPI_Win_get_name (MPI_Win win, char *win_name, int *resultlen)	
MPI_Win_lock (int lock_type, int rank, int assert, MPI_Win win)	
MPI_Win_set_attr (MPI_Win win, int win_keyval, void *attribute_val)	
MPI_Win_set_errhandler (MPI_Win win, MPI_Errhandler errhandler)	Attaches a new error handler to a window.
MPI_Win_set_name (MPI_Win win, char *win_name)	

TABLE A-1 Sun MPI Routines (Continued)

Routine and C Syntax	Description
<code>MPI_Win_unlock(int rank, MPI_Win win)</code>	Completes an RMA access epoch started by a call to <code>MPI_Win_lock()</code> .
<code>double MPI_Wtick()</code>	Returns the resolution of <code>MPI_Wtime()</code> .
<code>double MPI_Wtime()</code>	Returns an elapsed time on the calling processor.

Sun MPI I/O Routines

TABLE A-2 on page 77 lists the Sun MPI I/O routines in alphabetical order. The following sections list the routines by functional category.

File Manipulation

Collective coordination	Noncollective coordination
<code>MPI_File_open()</code>	<code>MPI_File_delete()</code>
<code>MPI_File_close()</code>	<code>MPI_File_get_size()</code>
<code>MPI_File_set_size()</code>	<code>MPI_File_get_group()</code>
<code>MPI_File_preallocate()</code>	<code>MPI_File_get_amode()</code>

File Info

Noncollective coordination	Collective coordination
<code>MPI_File_get_info()</code>	<code>MPI_File_set_info()</code>

Data access

Data Access With Explicit Offsets

Synchronism	Noncollective coordination	Collective coordination
Blocking	MPI_File_read_at()	MPI_File_read_at_all()
	MPI_File_write_at()	MPI_File_write_at_all()
Nonblocking or split collective	MPI_File_iread_at()	MPI_File_read_at_all_begin()
		MPI_File_read_at_all_end()
	MPI_File_iwrite_at()	MPI_File_write_at_all_begin()
		MPI_File_write_at_all_end()

Data Access With Individual File Pointers

Synchronism	Noncollective coordination	Collective coordination
Blocking	MPI_File_read()	MPI_File_read_all()
	MPI_File_write()	MPI_File_write_all()
Nonblocking or split collective	MPI_File_iread()	MPI_File_read_all_begin()
		MPI_File_read_all_end()
	MPI_File_iwrite()	MPI_File_write_all_begin()
		MPI_File_write_all_end()

Data Access With Shared File Pointers

Synchronism	Noncollective coordination	Collective coordination
Blocking	MPI_File_read_shared()	MPI_File_read_ordered()
	MPI_File_write_shared()	MPI_File_write_ordered()
		MPI_File_seek_shared()
		MPI_File_get_position_shared()
Nonblocking or split collective	MPI_File_iread_shared()	MPI_File_read_ordered_begin()
		MPI_File_read_ordered_end()
	MPI_File_iwrite_shared()	MPI_File_write_ordered_begin()
		MPI_File_write_ordered_end()

Pointer Manipulation

```
MPI_File_seek()  
MPI_File_get_position()  
MPI_File_get_byte_offset()
```

File Interoperability

```
MPI_Register_datarep()  
MPI_File_get_type_extent()
```

File Consistency and Semantics

```
MPI_File_set_atomicity()  
MPI_File_get_atomicity()  
MPI_File_sync()
```

Handle Translation

```
MPI_File_f2c()  
MPI_File_c2f()
```


MPI I/O Routines: Alphabetical Listing

TABLE A-2 Sun MPI I/O Routines

Routine and C Syntax	Description
MPI_File_c2f (MPI_File <i>file</i>)	Translates a C handle into a Fortran handle.
MPI_File_close (MPI_File <i>*fh</i>)	Closes a file (collective).
MPI_File_create_errhandler (MPI_File_errhandler_fn <i>*function</i> , MPI_Errhandler <i>*errhandler</i>)	Creates an MPI-style error handler that can be attached to a file.
MPI_File_delete (char <i>*filename</i> , MPI_Info <i>info</i>)	Deletes a file.
MPI_File_f2c (MPI_File <i>file</i>)	Translates a Fortran handle into a C handle.
MPI_File_get_amode (MPI_File <i>fh</i> , int <i>*amode</i>)	Returns mode associated with open file.
MPI_File_get_atomicity (MPI_File <i>fh</i> , int <i>*flag</i>)	Returns current consistency semantics for data-access operations.
MPI_File_get_byte_offset (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , MPI_Offset <i>*disp</i>)	Converts a view-relative offset into an absolute byte position.
MPI_File_get_errhandler (MPI_Comm <i>file</i> , MPI_Errhandler <i>*errhandler</i>)	Gets the error handler for a file.
MPI_File_get_group (MPI_File <i>fh</i> , MPI_Group <i>*group</i>)	Returns the process group of file.
MPI_File_get_info (MPI_File <i>fh</i> , MPI_Info <i>*info_used</i>)	Returns a new info object containing hints.
MPI_File_get_position (MPI_File <i>fh</i> , MPI_Offset <i>*offset</i>)	Returns current position of individual file pointer.
MPI_File_get_position_shared (MPI_File <i>fh</i> , MPI_Offset <i>*offset</i>)	Returns current position of the shared file pointer (collective).
MPI_File_get_size (MPI_File <i>fh</i> , MPI_Offset <i>*size</i>)	Returns current size of file.
MPI_File_get_type_extent (MPI_File <i>fh</i> , MPI_Datatype <i>datatype</i> , MPI_Aint <i>*extent</i>)	Returns the extent of the data type in a file.
MPI_File_get_view (MPI_File <i>fh</i> , MPI_Offset <i>*disp</i> , MPI_Datatype <i>*etype</i> , MPI_Datatype <i>*filetype</i> , char <i>*datarep</i>)	Returns process's view of data in file.

TABLE A-2 Sun MPI I/O Routines (Continued)

Routine and C Syntax	Description
MPI_File_iread (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Request * <i>request</i>)	Reads a file starting at the location specified by the individual file pointer (nonblocking, noncollective).
MPI_File_iread_at (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Request * <i>request</i>)	Reads a file at an explicitly specified offset (nonblocking, noncollective).
MPI_File_iread_shared (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Request * <i>request</i>)	Reads a file using the shared file pointer (nonblocking, noncollective).
MPI_File_iwrite (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Request * <i>request</i>)	Writes a file starting at the location specified by the individual file pointer (nonblocking, noncollective).
MPI_File_iwrite_at (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Request * <i>request</i>)	Writes a file at an explicitly specified offset (nonblocking, noncollective).
MPI_File_iwrite_shared (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Request * <i>request</i>)	Writes a file using the shared file pointer (nonblocking, noncollective).
MPI_File_open (MPI_Comm <i>comm</i> , char * <i>filename</i> , int <i>amode</i> , MPI_Info <i>info</i> , MPI_File * <i>fh</i>)	Opens a file (collective).
MPI_File_preallocate (MPI_File <i>fh</i> , MPI_Offset <i>size</i>)	Preallocates storage space for a portion of a file (collective).
MPI_File_read (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Reads a file starting at the location specified by the individual file pointer.
MPI_File_read_all (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Reads a file starting at the locations specified by individual file pointers (collective).
MPI_File_read_all_begin (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i>)	Reads a file starting at the locations specified by individual file pointers; beginning part of a split collective routine (nonblocking).
MPI_File_read_all_end (MPI_File <i>fh</i> , void * <i>buf</i> , MPI_Status * <i>status</i>)	Reads a file starting at the locations specified by individual file pointers; ending part of a split collective routine (blocking).
MPI_File_read_at (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Reads a file at an explicitly specified offset.

TABLE A-2 Sun MPI I/O Routines (Continued)

Routine and C Syntax	Description
MPI_File_read_at_all (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Reads a file at explicitly specified offsets (collective).
MPI_File_read_at_all_begin (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i>)	Reads a file at explicitly specified offsets; beginning part of a split collective routine (nonblocking).
MPI_File_read_at_all_end (MPI_File <i>fh</i> , void * <i>buf</i> , MPI_Status * <i>status</i>)	Reads a file at explicitly specified offsets; ending part of a split collective routine (blocking).
MPI_File_read_ordered (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Reads a file at a location specified by a shared file pointer (collective).
MPI_File_read_ordered_begin (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i>)	Reads a file at a location specified by a shared file pointer; beginning part of a split collective routine (nonblocking).
MPI_File_read_ordered_end (MPI_File <i>fh</i> , void * <i>buf</i> , MPI_Status * <i>status</i>)	Reads a file at a location specified by a shared file pointer; ending part of a split collective routine (blocking).
MPI_File_read_shared (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Reads a file using the shared file pointer (blocking, noncollective).
MPI_File_seek (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , int <i>whence</i>)	Updates individual file pointers.
MPI_File_seek_shared (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , int <i>whence</i>)	Updates the global shared file pointer (collective).
MPI_File_set_atomicity (MPI_File <i>fh</i> , int <i>flag</i>)	Sets consistency semantics for data-access operations (collective).
MPI_File_set_errhandler (MPI_File <i>file</i> , MPI_Errhandler <i>errhandler</i>)	Sets the error handler for a file.
MPI_File_set_info (MPI_File <i>fh</i> , MPI_Info <i>info</i>)	Sets new values for hints (collective).
MPI_File_set_size (MPI_File <i>fh</i> , MPI_Offset <i>size</i>)	Resizes a file (collective).
MPI_File_set_view (MPI_File <i>fh</i> , MPI_Offset <i>disp</i> , MPI_Datatype <i>etype</i> , MPI_Datatype <i>filetype</i> , char * <i>datarep</i> , MPI_Info <i>info</i>)	Changes process's view of data in file (collective).
MPI_File_sync (MPI_File <i>fh</i>)	Makes semantics consistent for data-access operations (collective).

TABLE A-2 Sun MPI I/O Routines (Continued)

Routine and C Syntax	Description
MPI_File_write (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Writes a file starting at the location specified by the individual file pointer.
MPI_File_write_all (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Writes a file starting at the locations specified by individual file pointers (collective).
MPI_File_write_all_begin (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i>)	Writes a file starting at the locations specified by individual file pointers; beginning part of a split collective routine (nonblocking).
MPI_File_write_all_end (MPI_File <i>fh</i> , void * <i>buf</i> , MPI_Status * <i>status</i>)	Writes a file starting at the locations specified by individual file pointers; ending part of a split collective routine (blocking).
MPI_File_write_at (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Writes a file at an explicitly specified offset.
MPI_File_write_at_all (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Writes a file at explicitly specified offsets (collective).
MPI_File_write_at_all_begin (MPI_File <i>fh</i> , MPI_Offset <i>offset</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i>)	Writes a file at explicitly specified offsets; beginning part of a split collective routine (nonblocking).
MPI_File_write_at_all_end (MPI_File <i>fh</i> , void * <i>buf</i> , MPI_Status * <i>status</i>)	Writes a file at explicitly specified offsets; ending part of a split collective routine (blocking).
MPI_File_write_ordered (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Writes a file at a location specified by a shared file pointer (collective).
MPI_File_write_ordered_begin (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i>)	Writes a file at a location specified by a shared file pointer; beginning part of a split collective routine (nonblocking).

TABLE A-2 Sun MPI I/O Routines (Continued)

Routine and C Syntax	Description
MPI_File_write_ordered_end (MPI_File <i>fh</i> , void * <i>buf</i> , MPI_Status * <i>status</i>)	Writes a file at a location specified by a shared file pointer; ending part of a split collective routine (blocking).
MPI_File_write_shared (MPI_File <i>fh</i> , void * <i>buf</i> , int <i>count</i> , MPI_Datatype <i>datatype</i> , MPI_Status * <i>status</i>)	Writes a file using the shared file pointer (blocking, noncollective).
MPI_Register_datarep (char * <i>datarep</i> , MPI_Datarep_conversion_function * <i>read_conversion_fn</i> , MPI_Datarep_conversion_function * <i>write_conversion_fn</i> , MPI_Datarep_extent_function * <i>dtype_file_extent_fn</i> , void * <i>extra_state</i>)	Defines data representation.

Environment Variables

Many environment variables are available for fine-tuning your Sun MPI environment. All 39 Sun MPI environment variables are listed here with brief descriptions. The same descriptions are also available on the `MPI` man page. If you want to return to the default setting after having set a variable, simply unset it (using `unsetenv`). The effects of some of the variables are explained in more detail in the *Sun HPC ClusterTools Performance Guide*.

The environment variables are listed here in six groups:

- “Informational” on page 83
- “General Performance Tuning” on page 84
- “Tuning Memory for Point-to-Point Performance” on page 85
- “Numerics” on page 87
- “Tuning Rendezvous” on page 88
- “Miscellaneous” on page 89

Informational

`MPI_PRINTENV`

When set to 1, causes the environment variables and `hpc.conf` parameters associated with the MPI job to be printed out. The default is 0.

`MPI_QUIET`

If set to 1, suppresses Sun MPI warning messages. The default value is 0.

MPI_SHOW_ERRORS

If set to 1, the `MPI_ERRORS_RETURN` error handler prints the error message and returns the error. The default value is 0.

MPI_SHOW_INTERFACES

When set to 1, 2 or 3, information regarding which interfaces are being used by an MPI application prints to `stdout`. Set `MPI_SHOW_INTERFACES` to 1 to print the selected internode interface. Set it to 2 to print all the interfaces and their rankings. Set it to 3 for verbose output. The default value, 0, does not print information to `stdout`.

General Performance Tuning

MPI_POLLALL

When set to 1, the default value, all connections are polled for receives, also known as full polling. When set to 0, only those connections are polled where receives are posted. Full polling helps drain system buffers and so lessen the chance of deadlock for “unsafe” codes. Well-written codes should set `MPI_POLLALL` to 0 for best performance.

MPI_PROCBIND

Binds each MPI process to its own processor. By default, `MPI_PROCBIND` is set to 0, which means processor binding is off. To turn processor binding on, set it to 1. The system administrator may allow or disable processor binding by setting the `pbind` parameter in the `hpc.conf` file on or off. If this parameter is set, the `MPI_PROCBIND` environment variable is disabled. Performance can be enhanced with processor binding, but very poor performance will result if processor binding is used for multithreaded jobs or for more than one job at a time.

MPI_SPIN

Sets the spin policy. The default value is 0, which causes MPI processes to spin nonaggressively, allowing best performance when the load is at least as great as the number of CPUs. A value of 1 causes MPI processes to spin aggressively, leading to best performance if extra CPUs are available on each node to handle system daemons and other background activities.

Tuning Memory for Point-to-Point Performance

MPI_RSM_CPOOLSIZE

The requested size, in bytes, to be allocated per stripe for buffers for each remote-shared-memory connection. This value may be overridden when connections are established. The default value is 16384 bytes.

MPI_RSM_NUMPOSTBOX

The number of postboxes per stripe per remote-shared-memory connection. The default is 15 postboxes.

MPI_RSM_PIPESIZE

The limit on the size (in bytes) of a message that can be sent over remote shared memory via the buffer list of one postbox per stripe. The default is 8192 bytes.

MPI_RSM_SBPOOLSIZE

If set, `MPI_RSM_SBPOOLSIZE` is the requested size in bytes of each RSM send buffer pool. An RSM send buffer pool is the pool of buffers on a node that a remote process would use to send to processes on the node. A multiple of 1024 must be used. If unset, then pools of buffers are dedicated to connections rather than to senders.

MPI_RSM_SHORTMSGSIZE

The maximum size, in bytes, of a message that will be sent via remote shared memory without using buffers. The default value is 401 bytes.

MPI_SHM_CPOOLSIZE

The amount of memory, in bytes, that can be allocated to each connection pool. When `MPI_SHM_SBPOOLSIZE` is not set, the default value is 24576 bytes. Otherwise, the default value is `MPI_SHM_SBPOOLSIZE`.

MPI_SHM_CYCLESIZE

The limit, in bytes, on the portion of a shared-memory message that will be sent via the buffer list of a single postbox during a cyclic transfer. The default value is 8192 bytes. A multiple of 1024 that is at most `MPI_SHM_CPOOLSIZE/2` must be used.

MPI_SHM_CYCLESTART

Shared-memory transfers that are larger than `MPI_SHM_CYCLESTART` bytes will be cyclic. The default value is 24576 bytes.

MPI_SHM_NUMPOSTBOX

The number of postboxes dedicated to each shared-memory connection. The default value is 16.

MPI_SHM_PIPESIZE

The limit, in bytes, on the portion of a shared-memory message that will be sent via the buffer list of a single postbox during a pipeline transfer. The default value is 8192 bytes. The value must be a multiple of 1024.

MPI_SHM_PIPESTART

The size, in bytes, at which shared-memory transfers will start to be pipelined. The default value is 2048. Multiples of 1024 must be used.

MPI_SHM_SBPOOLSIZE

If set, `MPI_SHM_SBPOOLSIZE` is the size, in bytes, of the pool of shared-memory buffers dedicated to each sender. A multiple of 1024 must be used. If unset, then pools of shared-memory buffers are dedicated to connections rather than to senders.

MPI_SHM_SHORTMSGSIZE

The size (in bytes) of the section of a postbox that contains either data or a buffer list. The default value is 256 bytes.

Note – If `MPI_SHM_PIPESTART`, `MPI_SHM_PIPESIZE`, or `MPI_SHM_CYCLESIZE` is increased to a size larger than 31744 bytes, then `MPI_SHM_SHORTMSGSIZE` may also have to be increased. See the *Sun HPC ClusterTools Performance Guide* for more information.

Numerics

MPI_CANONREDUCE

Prevents reduction operations from using any optimizations that take advantage of the physical location of processors. This may provide more consistent results in the case of floating-point addition, for example. However, the operation may take longer to complete. The default value is 0, meaning optimizations are allowed. To prevent optimizations, set the value to 1.

Tuning Rendezvous

`MPI_EAGERONLY`

When set to 1, the default, only the eager protocol is used. When set to 0, both eager and rendez-vous protocols are used.

`MPI_RSM_RENDVSIZE`

Messages communicated by remote shared memory that are greater than this size will use the rendezvous protocol unless the environment variable `MPI_EAGERONLY` is set to 1. Default value is 16384 bytes.

`MPI_SHM_RENDVSIZE`

Messages communicated by shared memory that are greater than this size will use the rendezvous protocol unless the environment variable `MPI_EAGERONLY` is set. The default value is 24576 bytes.

`MPI_TCP_RENDVSIZE`

Messages communicated by TCP that contain data of this size and greater will use the rendezvous protocol unless the environment variable `MPI_EAGERONLY` is set. Default value is 49152 bytes.

Miscellaneous

MPI_COSCHED

Specifies the user's preference regarding use of the `spind` daemon for coscheduling. Values can be 0 (prefer no use) or 1 (prefer use). This preference may be overridden by the system administrator's policy. This policy is set in the `hpc.conf` file and can be 0 (forbid use), 1 (require use), or 2 (no policy). If no policy is set and no user preference is specified, coscheduling is not used.

Note – If no user preference is specified, the value 2 will be shown when environment variables are printed with `MPI_PRINTENV`.

MPI_FLOWCONTROL

Limits the number of unexpected messages that can be queued from a particular connection. Once this quantity of unexpected messages has been received, polling the connection for incoming messages stops. The default value, 0, indicates that no limit is set. To limit flow, set the value to some integer greater than zero.

MPI_FULLCONNINIT

Ensures that all connections are established during initialization. By default, connections are established lazily. However, you can override this default by setting the environment variable `MPI_FULLCONNINIT` to 1, forcing full-connection initialization mode. The default value is 0.

MPI_MAXFHANDLES

The maximum number of Fortran handles for objects other than requests. `MPI_MAXFHANDLES` specifies the upper limit on the number of concurrently allocated Fortran handles for MPI objects other than requests. This variable is ignored in the default 32-bit library. The default value is 1024. Users should take care to free MPI objects that are no longer in use. There is no limit on handle allocation for C codes.

MPI_MAXREQHANDLES

The maximum number of Fortran request handles. `MPI_MAXREQHANDLES` specifies the upper limit on the number of concurrently allocated MPI request handles. Users must take care to free up request handles by properly completing requests. The default value is 1024. This variable is ignored in the default 32-bit library.

MPI_OPTCOLL

The MPI collectives are implemented using a variety of optimizations. Some of these optimizations can inhibit performance of point-to-point messages for “unsafe” programs. By default, this variable is 1, and optimized collectives are used. The optimizations can be turned off by setting the value to 0.

MPI_RSM_MAXSTRIPE

Defines the maximum number of interfaces used in a stripe during communication via remote shared memory. The default value is the number of interfaces in the cluster, with a maximum default of 2. Individual connections are load-balanced across multiple stripes in a cluster.

MPI_SHM_BCASTSIZE

On SMPs, the implementation of `MPI_Bcast()` for large messages is done using a double-buffering scheme. The size of each buffer (in bytes) is settable by using this environment variable. The default value is 32768 bytes.

MPI_SHM_GBPOOLSIZE

The amount of memory available, in bytes, to the general buffer pool for use by collective operations. The default value is 20971520 bytes.

MPI_SHM_REDUCE_SIZE

On SMPs, calling `MPI_Reduce()` causes all processors to participate in the reduce. Each processor will work on a piece of data equal to the `MPI_SHM_REDUCE_SIZE` setting. The default value is 256 bytes. Care must be taken when setting this variable because the system reserves `MPI_SHM_REDUCE_SIZE * np * np` memory to execute the reduce.

MPI_SPINDTIMEOUT

When coscheduling is enabled, limits the length of time (in milliseconds) a message will remain in the poll waiting for the `spind` daemon to return. If the timeout occurs before the daemon finds any messages, the process re-enters the polling loop. The default value is 1000 ms. A default can also be set by a system administrator in the `hpc.conf` file.

MPI_TCP_CONNLOOP

Sets the number of times `MPI_TCP_CONNTIMEOUT` occurs before signaling an error. The default value for this variable is 0, meaning that the program will abort on the first occurrence of `MPI_TCP_CONNTIMEOUT`.

MPI_TCP_CONNTIMEOUT

Sets the timeout value in seconds that is used for an `accept()` call. The default value for this variable is 600 seconds (10 minutes). This timeout can be triggered in both full- and lazy-connection initialization. After the timeout is reached, a warning message will be printed. If `MPI_TCP_CONNLOOP` is set to 0, then the first timeout will cause the program to abort.

MPI_TCP_SAFEGATHER

Allows use of a congestion-avoidance algorithm for `MPI_Gather()` and `MPI_Gatherv()` over TCP. By default, `MPI_TCP_SAFEGATHER` is set to 1, which means use of this algorithm is on. If you know that your underlying network can handle gathering large amounts of data on a single node, you may want to override this algorithm by setting `MPI_TCP_SAFEGATHER` to 0.

Troubleshooting

This appendix describes some common problem situations, resulting error messages, and suggestions for fixing the problems. Sun MPI error reporting, including I/O, follows the MPI-2 standard. By default, errors are reported in the form of standard error classes. These classes and their meanings are listed in TABLE C-1 on page 95 (for non-I/O MPI) and TABLE C-2 on page 97 (for MPI I/O) and are also available on the MPI man page.

Three predefined error handlers are available in Sun MPI:

- `MPI_ERRORS_RETURN` – The default, returns an error code if an error occurs.
- `MPI_ERRORS_ARE_FATAL` – I/O errors are fatal, and no error code is returned.
- `MPI_THROW_EXCEPTION` – A special error handler to be used only with C++.

MPI Messages

You can make changes to and get information about the error handler using any of the following routines:

- `MPI_Comm_create_errhandler`
- `MPI_Comm_get_errhandler`
- `MPI_Comm_set_errhandler`

Messages resulting from an MPI program fall into two categories:

- *Error messages* – Error messages stem from within MPI. Usually an error message explains why your program cannot complete, and the program aborts.
- *Warning messages* – Warnings stem from the environment in which you are running your MPI program and are usually sent by `MPI_Init()`. They are not associated with an aborted program, that is, programs continue to run despite warning messages.

Error Messages

Sun MPI error messages use a standard format:

`[xyz] Error in function_name: errclass_string:intern(a) : description: unixerrstring`
where

- `[xyz]` is the *process communication identifier*, and:
 - `x` is the job id (or jid).
 - `y` is the name of the communicator if a name exists; otherwise it is the address of the opaque object.
 - `z` is the rank of the process.

The process communication identifier is present in every error message.

- `function_name` is the name of the associated MPI function. It is present in every error message.
- `errclass_string` is the string associated with the MPI error class. It is present in every error message.
- `intern` is an internal function. It is optional.
- `a` is a system call, if one is the cause of the error. It is optional.
- `description` is a description of the error. It is optional.
- `unixerrstring` is the UNIX error string that describes system call `a`. It is optional.

Warning Messages

Sun MPI warning messages also use a standard format:

`[xyz] Warning message`

where `message` is a description of the error.

Standard Error Classes

Listed below are the error return classes you may encounter in your MPI programs. Error values may also be found in `mpi.h` (for C), `mpif.h` (for Fortran), and `mpi++.h` (for C++).

TABLE C-1 Sun MPI Standard Error Classes

Error Code	Value	Meaning
<code>MPI_SUCCESS</code>	0	Successful return code.
<code>MPI_ERR_BUFFER</code>	1	Invalid buffer pointer.
<code>MPI_ERR_COUNT</code>	2	Invalid count argument.
<code>MPI_ERR_TYPE</code>	3	Invalid datatype argument.
<code>MPI_ERR_TAG</code>	4	Invalid tag argument.
<code>MPI_ERR_COMM</code>	5	Invalid communicator.
<code>MPI_ERR_RANK</code>	6	Invalid rank.
<code>MPI_ERR_ROOT</code>	7	Invalid root.
<code>MPI_ERR_GROUP</code>	8	Null group passed to function.
<code>MPI_ERR_OP</code>	9	Invalid operation.
<code>MPI_ERR_TOPOLOGY</code>	10	Invalid topology.
<code>MPI_ERR_DIMS</code>	11	Illegal dimension argument.
<code>MPI_ERR_ARG</code>	12	Invalid argument.
<code>MPI_ERR_UNKNOWN</code>	13	Unknown error.
<code>MPI_ERR_TRUNCATE</code>	14	Message truncated on receive.
<code>MPI_ERR_OTHER</code>	15	Other error; use <code>Error_string</code> .
<code>MPI_ERR_INTERN</code>	16	Internal error code.
<code>MPI_ERR_IN_STATUS</code>	17	Look in status for error value.
<code>MPI_ERR_PENDING</code>	18	Pending request.
<code>MPI_ERR_REQUEST</code>	19	Illegal <code>MPI_Request()</code> handle.
<code>MPI_ERR_KEYVAL</code>	36	Illegal key value.
<code>MPI_ERR_INFO</code>	37	Invalid info object.
<code>MPI_ERR_INFO_KEY</code>	38	Illegal info key.

TABLE C-1 Sun MPI Standard Error Classes (Continued)

Error Code	Value	Meaning
MPI_ERR_INFO_NOKEY	39	No such key.
MPI_ERR_INFO_VALUE	40	Illegal info value.
MPI_ERR_TIMEOUT	41	Timed out.
MPI_ERR_RESOURCES	42	Out of resources.
MPI_ERR_TRANSPORT	43	Transport layer error.
MPI_ERR_HANDSHAKE	44	Error accepting/connecting.
MPI_ERR_SPAWN	45	Error spawning.
MPI_ERR_WIN	46	Invalid window.
MPI_ERR_BASE	47	Invalid base.
MPI_ERR_SIZE	48	Invalid size.
MPI_ERR_DISP	49	Invalid displacement.
MPI_ERR_LOCKTYPE	50	Invalid locktype.
MPI_ERR_ASSERT	51	Invalid assert.
MPI_ERR_RMA_CONFLICT	52	Conflicting accesses to window.
MPI_ERR_RMA_SYNC	53	Erroneous RMA synchronization.
MPI_ERR_NO_MEM	54	Memory exhausted.
MPI_ERR_LASTCODE	55	Last error code.

MPI I/O message are listed separately, in TABLE C-2 on page 97.

MPI I/O Error Handling

Sun MPI I/O error reporting follows the MPI-2 standard. By default, errors are reported in the form of standard error codes (found in `/opt/SUNWhpc/include/mpi.h`). Error classes and their meanings are listed in TABLE C-2 on page 97. They can also be found in `mpif.h` (for Fortran) and `mpi++.h` (for C++).

You can change the default error handler by specifying `MPI_FILE_NULL` as the file handle with the routine `MPI_File_set_errhandler()`, even no file is currently open. Or, you can use the same routine to change a specific file's error handler.

TABLE C-2 Sun MPI I/O Error Classes

Error Class	Value	Meaning
<code>MPI_ERR_FILE</code>	20	Bad file handle.
<code>MPI_ERR_NOT_SAME</code>	21	Collective argument not identical on all processes.
<code>MPI_ERR_AMODE</code>	22	Unsupported <code>amode</code> passed to open.
<code>MPI_ERR_UNSUPPORTED_DATAREP</code>	23	Unsupported <code>datarep</code> passed to <code>MPI_File_set_view()</code> .
<code>MPI_ERR_UNSUPPORTED_OPERATION</code>	24	Unsupported operation, such as seeking on a file that supports only sequential access.
<code>MPI_ERR_NO_SUCH_FILE</code>	25	File (or directory) does not exist.
<code>MPI_ERR_FILE_EXISTS</code>	26	File exists.
<code>MPI_ERR_BAD_FILE</code>	27	Invalid file name (for example, path name too long).
<code>MPI_ERR_ACCESS</code>	28	Permission denied.
<code>MPI_ERR_NO_SPACE</code>	29	Not enough space.
<code>MPI_ERR_QUOTA</code>	30	Quota exceeded.
<code>MPI_ERR_READ_ONLY</code>	31	Read-only file system.
<code>MPI_ERR_FILE_IN_USE</code>	32	File operation could not be completed, as the file is currently open by some process.
<code>MPI_ERR_DUP_DATAREP</code>	33	Conversion functions could not be registered because a data representation identifier that was already defined was passed to <code>MPI_REGISTER_DATAREP</code> .
<code>MPI_ERR_CONVERSION</code>	34	An error occurred in a user-supplied data-conversion function.
<code>MPI_ERR_IO</code>	35	I/O error.
<code>MPI_ERR_INFO</code>	37	Invalid info object.
<code>MPI_ERR_INFO_KEY</code>	38	Illegal info key.

TABLE C-2 Sun MPI I/O Error Classes *(Continued)*

Error Class	Value	Meaning
MPI_ERR_INFO_NOKEY	39	No such key .
MPI_ERR_INFO_VALUE	40	Illegal info value.
MPI_ERR_LASTCODE	55	Last error code.

TNF Probes

Through Prism, you can use Trace Normal Form (TNF), an extensible system for event-based instrumentation, to aid debugging and to analyze the performance of your Sun MPI programs. TNF is included with the Solaris operating environment. The TNF-instrumented libraries included with Sun MPI (see “Choosing a Library Path” on page 26) include probes for most of the MPI and MPI I/O routines, including some specific arguments. These probes are also categorized into specific groups, so that you can debug and analyze the performance of particular types of routines. For information about using Prism to take advantage of these probes, see the *Prism User’s Guide* and the *Sun HPC ClusterTools Performance Guide*.

This appendix includes all the probes, including their arguments and associated groups, both for MPI (TABLE D-1 on page 104) and for MPI I/O (TABLE D-2 on page 118). The following figure depicts the relationships among the various probe groups. (Some probes fall under both the `mpi_request` and the `mpi_pt2pt` groups.)

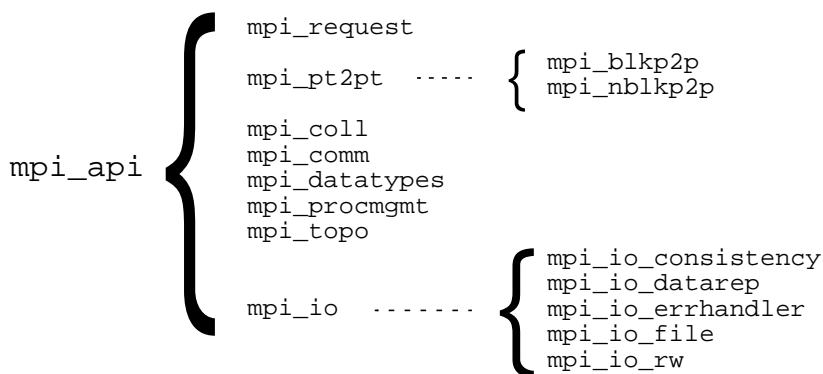


FIGURE D-1 TNF Probe Groups for Sun MPI, Including I/O

TNF Probes for MPI

Each MPI routine is associated with two TNF probes: one ending in `_start`, and one ending in `_end`. Probes are also included for some specific arguments, most of which are defined in the MPI standard and described in the man pages included with Sun MPI. Four of the arguments, however, are not mentioned in the standard or man pages:

- `bytes` – The number of bytes sent or received by an MPI process. See the next section for more information about the `bytes` argument.
- `ctxt` – The *context id* is a number assigned to a particular communicator. The processes in a given communicator may be associated with only one context id. You can determine the context id associated with a communicator using either the `MPI_Comm_set_name_end` or the `MPI_Comm_get_name_end` probe.
- `newctxt` – The context id associated with a communicator that is returned as a `newcomm` or `comm_out` argument.
- `request` – An integer that uniquely refers to a request object. For Fortran calls, this integer is equal to the `request-handle` argument.

The `bytes` Argument

The meaning of the `bytes` argument varies slightly depending on the situation. Here is some general information for different types of sends and receives, followed by some examples of how the `byte` argument works with them.

- *point-to-point blocking sends* – The `start` probes for routines that initiate point-to-point blocking sends report the number of bytes to be sent. These routines are:
 - `MPI_Bsend()`
 - `MPI_Rsend()`
 - `MPI_Send()`
 - `MPI_Ssend()`
 - `MPI_Sendrecv()`
 - `MPI_Sendrecv_replace()`
- *point-to-point nonblocking sends* – The `end` probes for routines that initiate point-to-point nonblocking sends report the number of bytes to be sent. These routines are:
 - `MPI_Bsend_init()`
 - `MPI_Ibsend()`
 - `MPI_Irsend()`
 - `MPI_Isend()`
 - `MPI_Issend()`

- `MPI_Rsend_init()`
- `MPI_Send_init()`
- `MPI_Ssend_init()`
- *point-to-point receives* – The end probes for routines that terminate or could terminate point-to-point nonblocking sends report the number of bytes *actually* received. These routines are:
 - `MPI_Iprobe()`
 - `MPI_Probe()`
 - `MPI_Recv()`
 - `MPI_Test()`
 - `MPI_Testall()`
 - `MPI_Testany()`
 - `MPI_Testsome()`
 - `MPI_Wait()`
 - `MPI_Waitall()`
 - `MPI_Waitany()`
 - `MPI_Waitsome()`
 - `MPI_Sendrecv()`
 - `MPI_Sendrecv_replace()`
- *collectives* – The start probes for collective routines report the number of bytes to be sent from an MPI process and the number to be received at the process. Such byte counts are independent of the algorithm used. For example, the number of bytes sent from the root in a broadcast is given as the number of bytes in the broadcast message, regardless of whether the root sends this message multiple times as part of a binary-tree fan-out. These collective routines are:
 - `MPI_Allgather()`, `MPI_Allgatherv()`
 - `sendbytes` – Number of bytes to be sent from this process.
 - `recvbytes` – Total number of bytes to be received at any process from all processes.
 - `MPI_Allreduce()`, `MPI_Reduce()`, `MPI_Reduce_scatter()`
 - `bytes` – Number of bytes on any process to be reduced.
 - `MPI_Alltoall()`, `MPI_Alltoallv()`
 - `sendbytes` – Total number of bytes to be sent from this process to all processes.
 - `recvbytes` – Total number of bytes to be received at this process from all processes.
 - `MPI_Bcast()`
 - `bytes` – Number of bytes to be broadcast.
 - `MPI_Gather()`, `MPI_Gatherv()`

sendbytes – Root reports total number of bytes to be sent; other processes report 0.

recvbytes – Root reports total number of bytes to be received; other processes report 0.

- MPI_Scan()

bytes – Number of bytes contributed by any process.

- MPI_Scatter(), MPI_Scatterv()

sendbytes – Root reports total number of bytes to be sent; other processes report 0.

recvbytes – Number of bytes to be received at this process from the root.

- *pack and unpack* – The start probes for these routines report the number of bytes packed or unpacked. These routines are:

- MPI_Pack()

- MPI_Unpack()

Examples:

- MPI_Send()

```
call MPI_Send(x,m,MPI_REAL8,...)
```

Probe `mpi_send_start` reports that 8^*m bytes are to be sent.

- MPI_Recv()

```
call MPI_Recv(x,n,MPI_REAL8,...)
```

Probe `mpi_recv_end` reports the number of bytes that were actually received, which must be at most 8^*n .

- MPI_Sendrecv()

```
call MPI_Sendrecv(x,m,MPI_REAL8,...,y,n,MPI_REAL8,...)
```

Probe `mpi_sendrecv_start` reports that 8^*m bytes are to be sent, and probe `mpi_sendrecv_end` reports the number of bytes that were actually received, which must be at most 8^*n .

- MPI_Irecv(), MPI_Wait()

```
integer req
```

```
call MPI_Irecv(x,n,MPI_REAL8,...,req,...)
```

```
call MPI_Wait(req,...)
```

Probe `mpi_wait_end` reports the number of bytes that were actually received, which must be at most 8^*n .

- `MPI_Isend()`, `MPI_Irecv()`, `MPI_Wait()`

```
integer reqs(2)
call MPI_Isend(x,m,MPI_REAL8,...,reqs(1),...)
call MPI_Irecv(Y,N,MPI_REAL8,...,reqs(2),...)
call MPI_Waitany(2,reqs,...)
call MPI_Waitany(2,reqs,...)
```

Probe `mpi_isend_start` reports that 8^m bytes are to be sent. The `MPI_Waitany()` call that completes the receive will show the number of bytes that were actually received, which must be at most 8^n , in its `mpi_waitany_end` probe. The other `MPI_Waitany()` call, which completes the send, will report 0 bytes received.

- `MPI_Waitall()`

```
integer reqs(8)
call MPI_Isend(x1,m,MPI_REAL8,...,reqs(1),...)
call MPI_Isend(x2,m,MPI_REAL8,...,reqs(2),...)
call MPI_Isend(x3,m,MPI_REAL8,...,reqs(3),...)
call MPI_Isend(x4,m,MPI_REAL8,...,reqs(4),...)
call MPI_Irecv(x5,n,MPI_REAL8,...,reqs(5),...)
call MPI_Irecv(x6,n,MPI_REAL8,...,reqs(6),...)
call MPI_Irecv(x7,n,MPI_REAL8,...,reqs(7),...)
call MPI_Irecv(x8,n,MPI_REAL8,...,reqs(8),...)
call MPI_Waitall(8,reqs,...)
```

Probe `mpi_isend_start` reports that 8^m bytes are to be sent in each of the four `MPI_Isend()` cases. Probe `mpi_waitall_end` reports the number of bytes that were actually received, which must be at most $4 \cdot 8^n$.

Groups

Every TNF probe for MPI is associated with the `mpi_api` group, so choosing that group allows Prism to probe all the MPI routines for which probes exist (including the I/O routines). Additional groups exist to probe subsets of the MPI routines, as well. Some routines are associated with more than one group. The ten groups for MPI routine probes are these:

- `mpi_api` – All the TNF probes for MPI routines
- `mpi_blkp2p` – Probes for blocking point-to-point routines
- `mpi_coll` – Probes for collective routines
- `mpi_comm` – Probes for communicator-related routines

- `mpi_datatypes` – Probes for data type-related routines
- `mpi_nblkp2p` – Probes for nonblocking point-to-point routines
- `mpi_procmgmt` – Probes for process-management routines
- `mpi_pt2pt` – Probes for all point-to-point routines (blocking and nonblocking)
- `mpi_request` – Probes for functions producing or acting on request(s)
- `mpi_topo` – Probes for topology-related routines

Probes for MPI (Non-I/O Routines)

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls

Probe	Argument(s)	Group(s), in Addition to <code>mpi_api</code>
<code>MPI_Accumulate_start</code>		<code>mpi_win</code>
<code>MPI_Accumulate_end</code>		<code>mpi_win</code>
<code>MPI_Address_start</code>		
<code>MPI_Address_end</code>		
<code>MPI_Allgather_start</code>	<code>sendbytes</code> <code>recvbytes</code> <code>ctxt</code>	<code>mpi_coll</code>
<code>MPI_Allgather_end</code>		<code>mpi_coll</code>
<code>MPI_Allgatherv_start</code>	<code>sendbytes</code> <code>recvbytes</code> <code>ctxt</code>	<code>mpi_coll</code>
<code>MPI_Allgatherv_end</code>		<code>mpi_coll</code>
<code>MPI_Alloc_mem_start</code>		<code>mpi_comm</code>
<code>MPI_Alloc_mem_end</code>		<code>mpi_comm</code>
<code>MPI_Allreduce_start</code>	<code>bytes</code> <code>ctxt</code>	<code>mpi_coll</code>
<code>MPI_Allreduce_end</code>		<code>mpi_coll</code>
<code>MPI_Alltoall_start</code>	<code>sendbytes</code> <code>recvbytes</code> <code>ctxt</code>	<code>mpi_coll</code>
<code>MPI_Alltoall_end</code>		<code>mpi_coll</code>
<code>MPI_Alltoallv_start</code>	<code>sendbytes</code> <code>recvbytes</code> <code>ctxt</code>	<code>mpi_coll</code>
<code>MPI_Alltoallv_end</code>		<code>mpi_coll</code>
<code>MPI_Attr_delete_start</code>		
<code>MPI_Attr_delete_end</code>		

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls (Continued)

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Attr_get_start		
MPI_Attr_get_end		
MPI_Attr_put_start		
MPI_Attr_put_end		
MPI_Barrier_start	ctxt	mpi_coll
MPI_Barrier_end		mpi_coll
MPI_Bcast_start	bytes root ctxt	mpi_coll
MPI_Bcast_end		mpi_coll
MPI_Bsend_start	bytes dest tag	mpi_pt2pt mpi_blkp2p
MPI_Bsend_end		mpi_pt2pt mpi_blkp2p
MPI_Bsend_init_start		mpi_pt2pt mpi_request
MPI_Bsend_init_end	bytes dest tag request	mpi_pt2pt mpi_request
MPI_Buffer_attach_start	buffer size	
MPI_Buffer_attach_end	buffer size	
MPI_Buffer_detach_start	buffer size	
MPI_Buffer_detach_end		
MPI_Cancel_start	request	mpi_request
MPI_Cancel_end		mpi_request
MPI_Cart_coords_start		mpi_topo
MPI_Cart_coords_end		mpi_topo
MPI_Cart_create_start		mpi_topo
MPI_Cart_create_end		mpi_topo
MPI_Cartdim_get_start		mpi_topo
MPI_Cartdim_get_end		mpi_topo
MPI_Cart_get_start		mpi_topo
MPI_Cart_get_end		mpi_topo
MPI_Cart_map_start		mpi_topo

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls (Continued)

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Cart_map_end		mpi_topo
MPI_Cart_rank_start		mpi_topo
MPI_Cart_rank_end		mpi_topo
MPI_Cart_shift_start		mpi_topo
MPI_Cart_shift_end		mpi_topo
MPI_Cart_sub_start		mpi_topo
MPI_Cart_sub_end		mpi_topo
MPI_Close_port_start	port_name	mpi_procmgmt
MPI_Close_port_end		mpi_procmgmt
MPI_Comm_accept_start	port_name root ctxt	mpi_procmgmt
MPI_Comm_accept_end	port_name root ctxt newctxt	mpi_procmgmt
MPI_Comm_compare_start		mpi_comm
MPI_Comm_compare_end		mpi_comm
MPI_Comm_connect_start	port_name root ctxt	mpi_procmgmt
MPI_Comm_connect_end	port_name root ctxt newctxt	mpi_procmgmt
MPI_Comm_create_start	ctxt group	mpi_comm
MPI_Comm_create_end	ctxt newctxt	mpi_comm
MPI_Comm_create_errhandler_start		
MPI_Comm_create_errhandler_end		
MPI_Comm_create_keyval_start		
MPI_Comm_create_keyval_end		
MPI_Comm_delete_attr_start		
MPI_Comm_delete_attr_end		
MPI_Comm_disconnect_start	ctxt	mpi_procmgmt
MPI_Comm_disconnect_end		mpi_procmgmt
MPI_Comm_dup_start	ctxt	mpi_comm
MPI_Comm_dup_end	ctxt newctxt	mpi_comm

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Comm_free_start	ctxt	mpi_comm
MPI_Comm_free_end		mpi_comm
MPI_Comm_free_keyval_start		
MPI_Comm_free_keyval_end		
MPI_Comm_get_attr_start		
MPI_Comm_get_attr_end		
MPI_Comm_get_errhandler_start		
MPI_Comm_get_errhandler_end		
MPI_Comm_get_name_start		
MPI_Comm_get_name_end	ctxt comm_name	mpi_comm
MPI_Comm_group_start		
MPI_Comm_group_end		
MPI_Comm_remote_group_start		
MPI_Comm_remote_group_end		
MPI_Comm_set_attr_start		
MPI_Comm_set_attr_end		
MPI_Comm_set_errhandler_start		
MPI_Comm_set_errhandler_end		
MPI_Comm_set_name_start		
MPI_Comm_set_name_end	ctxt comm_name	mpi_comm
MPI_Comm_split_start	ctxt	mpi_comm
MPI_Comm_split_end	ctxt newctxt	mpi_comm
MPI_Comm_test_inter_start		
MPI_Comm_test_inter_end		
MPI_Dims_create_start		mpi_topo
MPI_Dims_create_end		mpi_topo
MPI_Errhandler_create_start		
MPI_Errhandler_create_end		
MPI_Errhandler_free_start		
MPI_Errhandler_free_end		

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to <code>mpi_api</code>
<code>MPI_Errhandler_get_start</code>		
<code>MPI_Errhandler_get_end</code>		
<code>MPI_Errhandler_set_start</code>		
<code>MPI_Errhandler_set_end</code>		
<code>MPI_Error_class_start</code>		
<code>MPI_Error_class_end</code>		
<code>MPI_Error_string_start</code>		
<code>MPI_Error_string_end</code>		
<code>MPI_Finalize_start</code>		
<code>MPI_Finalize_end</code>		
<code>MPI_Free_mem_start</code>		
<code>MPI_Free_mem_end</code>		
<code>MPI_Gather_start</code>	<code>sendbytes</code> <code>recvbytes</code> <code>root</code> <code>ctxt</code>	<code>mpi_coll</code>
<code>MPI_Gather_end</code>		<code>mpi_coll</code>
<code>MPI_Gatherv_start</code>	<code>sendbytes</code> <code>recvbytes</code> <code>root</code> <code>ctxt</code>	<code>mpi_coll</code>
<code>MPI_Gatherv_end</code>		<code>mpi_coll</code>
<code>MPI_Get_start</code>		<code>mpi_win</code>
<code>MPI_Get_end</code>		<code>mpi_win</code>
<code>MPI_Get_address_start</code>		
<code>MPI_Get_address_end</code>		
<code>MPI_Get_count_start</code>		
<code>MPI_Get_count_end</code>		
<code>MPI_Get_elements_start</code>		
<code>MPI_Get_elements_end</code>		
<code>MPI_Get_processor_name_start</code>		
<code>MPI_Get_processor_name_end</code>		
<code>MPI_Get_version_start</code>		

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to <code>mpi_api</code>
MPI_Get_version_end		
MPI_Graph_create_start		mpi_topo
MPI_Graph_create_end		mpi_topo
MPI_Graphdims_get_start		mpi_topo
MPI_Graphdims_get_end		mpi_topo
MPI_Graph_get_start		mpi_topo
MPI_Graph_get_end		mpi_topo
MPI_Graph_map_start		mpi_topo
MPI_Graph_map_end		mpi_topo
MPI_Graph_neighbors_start		mpi_topo
MPI_Graph_neighbors_end		mpi_topo
MPI_Graph_neighbors_count_start		mpi_topo
MPI_Graph_neighbors_count_end		mpi_topo
MPI_Grequest_complete_start	request	mpi_request
MPI_Grequest_complete_end		mpi_request
MPI_Grequest_start_start	request	mpi_request
MPI_Grequest_start_end		mpi_request
MPI_Group_compare_start		
MPI_Group_compare_end		
MPI_Group_difference_start		
MPI_Group_difference_end		
MPI_Group_excl_start		
MPI_Group_excl_end		
MPI_Group_free_start		
MPI_Group_free_end		
MPI_Group_incl_start		
MPI_Group_incl_end		
MPI_Group_intersection_start		
MPI_Group_intersection_end		
MPI_Group_range_excl_start		

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to <code>mpi_api</code>
MPI_Group_range_excl_end		
MPI_Group_range_incl_start		
MPI_Group_range_incl_end		
MPI_Group_translate_ranks_start		
MPI_Group_translate_ranks_end		
MPI_Group_union_start		
MPI_Group_union_end		
MPI_Ibsend_start		mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Ibsend_end	bytes dest tag done request	mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Info_create_start		
MPI_Info_create_end		
MPI_Info_delete_start		
MPI_Info_delete_end		
MPI_Info_dup_start		
MPI_Info_dup_end		
MPI_Info_free_start		
MPI_Info_free_end		
MPI_Info_get_start		
MPI_Info_get_end		
MPI_Info_get_nkeys_start		
MPI_Info_get_nkeys_end		
MPI_Info_get_nthkey_start		
MPI_Info_get_nthkey_end		
MPI_Info_get_valuelen_start		
MPI_Info_get_valuelen_end		
MPI_Info_set_start		
MPI_Info_set_end		

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Intercomm_create_start		mpi_comm
MPI_Intercomm_create_end		mpi_comm
MPI_Intercomm_merge_start		mpi_comm
MPI_Intercomm_merge_end		mpi_comm
MPI_Iprobe_start	source tag ctxt	
MPI_Iprobe_end	source tag ctxt flag	
MPI_Irecv_start		mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Irecv_end	done request	mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Irsend_start		mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Irsend_end	bytes dest tag done request	mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Isend_start		mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Isend_end	bytes dest tag done request	mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Issend_start		mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Issend_end	bytes dest tag done request	mpi_pt2pt mpi_nblkp2p mpi_request
MPI_Keyval_create_start		
MPI_Keyval_create_end		
MPI_Keyval_free_start		
MPI_Keyval_free_end		
MPI_Op_create_start		

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls (Continued)

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Op_create_end		
MPI_Open_port_start	port_name	mpi_procmgmt
MPI_Open_port_end	port_name	mpi_procmgmt
MPI_Op_free_start		
MPI_Op_free_end		
MPI_Pack_start	bytes	mpi_datatypes
MPI_Pack_end		mpi_datatypes
MPI_Pack_size_start	count datatype	mpi_datatypes
MPI_Pack_size_end	count datatype size	mpi_datatypes
MPI_Pcontrol_start		
MPI_Pcontrol_end		
MPI_Probe_start	source tag ctxt	
MPI_Probe_end	source tag ctxt	
MPI_Put_start		mpi_win
MPI_Put_end		mpi_win
MPI_Query_thread_start		
MPI_Query_thread_end		
MPI_Recv_start		mpi_pt2pt mpi_blkp2p
MPI_Recv_end	bytes source tag	mpi_pt2pt mpi_blkp2p
MPI_Recv_init_start		mpi_pt2pt mpi_request
MPI_Recv_init_end	request	mpi_pt2pt mpi_request
MPI_Reduce_start	bytes root ctxt	mpi_coll
MPI_Reduce_end		mpi_coll
MPI_Reduce_scatter_start	bytes ctxt	mpi_coll
MPI_Reduce_scatter_end		mpi_coll
MPI_Request_free_start	request	mpi_request
MPI_Request_free_end		mpi_request

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls (Continued)

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Rsend_start	bytes dest tag	mpi_pt2pt mpi_blkp2p
MPI_Rsend_end		mpi_pt2pt mpi_blkp2p
MPI_Rsend_init_start		mpi_pt2pt mpi_request
MPI_Rsend_init_end	bytes dest tag request	mpi_pt2pt mpi_request
MPI_Scan_start	bytes ctxt	mpi_coll
MPI_Scan_end		mpi_coll
MPI_Scatter_start	sendbytes recvbytes root ctxt	mpi_coll
MPI_Scatter_end		mpi_coll
MPI_Scatterv_start	sendbytes recvbytes root ctxt	mpi_coll
MPI_Scatterv_end		mpi_coll
MPI_Send_start	bytes dest tag	mpi_pt2pt mpi_blkp2p
MPI_Send_end		mpi_pt2pt mpi_blkp2p
MPI_Send_init_start		mpi_pt2pt mpi_request
MPI_Send_init_end	bytes dest tag request	mpi_pt2pt mpi_request
MPI_Sendrecv_start	bytes dest sendtag	mpi_pt2pt mpi_blkp2p
MPI_Sendrecv_end	bytes source recvtag	mpi_pt2pt mpi_blkp2p
MPI_Sendrecv_replace_start	bytes dest sendtag	mpi_pt2pt mpi_blkp2p
MPI_Sendrecv_replace_end	bytes source recvtag	mpi_pt2pt mpi_blkp2p
MPI_Ssend_start	bytes dest tag	mpi_pt2pt mpi_blkp2p

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Ssend_end		mpi_pt2pt mpi_blkp2p
MPI_Ssend_init_start		mpi_pt2pt mpi_request
MPI_Ssend_init_end	bytes dest tag request	mpi_pt2pt mpi_request
MPI_Start_start	request	mpi_pt2pt mpi_request
MPI_Start_end		mpi_pt2pt mpi_request
MPI_Startall_start	count	mpi_pt2pt mpi_request
MPI_Startall_end		mpi_pt2pt mpi_request
MPI_Status_set_cancelled_start		
MPI_Status_set_cancelled_end		
MPI_Status_set_elements_start		
MPI_Status_set_elements_end		
MPI_Test_start	request	mpi_request
MPI_Test_end	recvbytes source recvtag flag request	mpi_request
MPI_Testall_start	count	mpi_request
MPI_Testall_end	bytes count flag	mpi_request
MPI_Testany_start	count	mpi_request
MPI_Testany_end	bytes index flag	mpi_request
MPI_Test_cancelled_start		mpi_request
MPI_Test_cancelled_end	flag	mpi_request
MPI_Testsome_start	incount	mpi_request
MPI_Testsome_end	bytes outcount	mpi_request
MPI_Topo_test_start		mpi_topo
MPI_Topo_test_end		mpi_topo

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to <code>mpi_api</code>
<code>MPI_Type_contiguous_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_contiguous_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_hindexed_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_hindexed_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_f90_integer_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_f90_integer_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_keyval_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_keyval_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_struct_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_create_struct_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_delete_attr_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_delete_attr_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_dup_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_dup_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_extent_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_extent_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_free_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_free_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_free_keyval_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_free_keyval_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_attr_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_attr_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_contents_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_contents_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_envelope_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_envelope_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_extent_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_extent_end</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_name_start</code>		<code>mpi_datatypes</code>
<code>MPI_Type_get_name_end</code>		<code>mpi_datatypes</code>

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Type_get_true_extent_start		mpi_datatypes
MPI_Type_get_true_extent_end		mpi_datatypes
MPI_Type_hindexed_start		mpi_datatypes
MPI_Type_hindexed_end		mpi_datatypes
MPI_Type_indexed_start		mpi_datatypes
MPI_Type_indexed_end		mpi_datatypes
MPI_Type_create_indexed_block_start		mpi_datatypes
MPI_Type_create_indexed_block_end		mpi_datatypes
MPI_Type_lb_start		mpi_datatypes
MPI_Type_lb_end		mpi_datatypes
MPI_Type_create_resized_start		mpi_datatypes
MPI_Type_create_resized_end		mpi_datatypes
MPI_Type_set_attr_start		mpi_datatypes
MPI_Type_set_attr_end		mpi_datatypes
MPI_Type_set_name_start		mpi_datatypes
MPI_Type_set_name_end		mpi_datatypes
MPI_Type_size_start		mpi_datatypes
MPI_Type_size_end		mpi_datatypes
MPI_Type_struct_start		mpi_datatypes
MPI_Type_struct_end		mpi_datatypes
MPI_Type_ub_start		mpi_datatypes
MPI_Type_ub_end		mpi_datatypes
MPI_Unpack_start	bytes	mpi_datatypes
MPI_Unpack_end		mpi_datatypes
MPI_Wait_start	request	mpi_request
MPI_Wait_end	recvbytes source recvtag request	mpi_request
MPI_Waitall_start	count	mpi_request
MPI_Waitall_end	bytes count	mpi_request
MPI_Waitany_start	count	mpi_request

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api
MPI_Waitany_end	bytes index	mpi_request
MPI_Waitsome_start	incount	mpi_request
MPI_Waitsome_end	bytes outcount	mpi_request
MPI_Win_create_start		mpi_win
MPI_Win_create_end		mpi_win
MPI_Win_create_errhandler_start		
MPI_Win_create_errhandler_end		
MPI_Win_create_keyval_start		
MPI_Win_create_keyval_end		
MPI_Win_delete_attr_start		
MPI_Win_delete_attr_end		
MPI_Win_fence_start		mpi_win
MPI_Win_fence_end		mpi_win
MPI_Win_free_start		mpi_win
MPI_Win_free_end		mpi_win
MPI_Win_free_keyval_start		
MPI_Win_free_keyval_end		
MPI_Win_get_attr_start		
MPI_Win_get_attr_end		
MPI_Win_get_errhandler_start		
MPI_Win_get_errhandler_end		
MPI_Win_get_group_start		
MPI_Win_get_group_end		
MPI_Win_get_name_start		
MPI_Win_get_name_end		
MPI_Win_lock_start		mpi_win
MPI_Win_lock_end		mpi_win
MPI_Win_set_attr_start		
MPI_Win_set_attr_end		
MPI_Win_set_errhandler_start		

TABLE D-1 TNF Probes, Associated Arguments, & Groups for MPI Calls (Continued)

Probe	Argument(s)	Group(s), in Addition to <code>mpi_api</code>
<code>MPI_Win_set_errhandler_end</code>		
<code>MPI_Win_set_name_start</code>		
<code>MPI_Win_set_name_end</code>		
<code>MPI_Win_unlock_start</code>		<code>mpi_win</code>
<code>MPI_Win_unlock_end</code>		<code>mpi_win</code>

TNF Probes for MPI I/O

Like the MPI routines, each MPI I/O routine is associated with two TNF probes: one ending in `_start`, and one ending in `_end`. Probes are also included for some specific arguments, most of which are defined in the MPI standard and described in the man pages included with Sun MPI. The `ctxt` argument for the context id assigned to a particular communicator, however, is not mentioned in the standard or man pages. It is described in “TNF Probes for MPI” on page 100.

Every TNF probe for MPI I/O is associated with both the `mpi_api` and the `mpi_io` groups. Choosing `mpi_api` allows Prism to probe all the MPI routines for which probes exist, whereas choosing `mpi_io` allows you to focus on the I/O routines. Additional groups exist to probe subsets of the I/O routines, as well. The seven groups for MPI I/O routine probes are these:

- `mpi_api` - All the TNF probes for MPI routines
- `mpi_io` - MPI I/O routines only
- `mpi_io_consistency` - Atomicity and synchronization routines
- `mpi_io_datarep` - Data representation routines
- `mpi_io_errhandler` - Error-handling routines
- `mpi_io_file` - Group(s), in addition to `mpi_api` and `mpi_io`
- `mpi_io_rw` - Read/write routines

TABLE D-2 TNF Probes, Associated Arguments, & Groups for MPI I/O Calls

Probe	Argument(s)	Group(s), in Addition to <code>mpi_api</code> & <code>mpi_io</code>
<code>MPI_File_close_start</code>	<code>filename</code>	<code>mpi_io_file</code>
<code>MPI_File_close_end</code>		<code>mpi_io_file</code>
<code>MPI_File_create_errhandler_start</code>		<code>mpi_io_errhandler</code>

TABLE D-2 TNF Probes, Associated Arguments, & Groups for MPI I/O Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api & mpi_io
MPI_File_create_errhandler_end		mpi_io_errhandler
MPI_File_delete_start	filename	mpi_io_file
MPI_File_delete_end	filename	mpi_io_file
MPI_File_get_amode_start	filename amode	mpi_io_file
MPI_File_get_amode_end	filename amode	mpi_io_file
MPI_File_get_atomicity_start	filename flag	mpi_io_consistency
MPI_File_get_atomicity_end	filename flag	mpi_io_consistency
MPI_File_get_byte_offset_start	filename offset disp	mpi_io_rw
MPI_File_get_byte_offset_end	filename offset disp	mpi_io_rw
MPI_File_get_errhandler_start	filename	mpi_io_errhandler
MPI_File_get_errhandler_end	filename	mpi_io_errhandler
MPI_File_get_group_start	filename	mpi_io_file
MPI_File_get_group_end	filename	mpi_io_file
MPI_File_get_info_start	filename	mpi_io_file
MPI_File_get_info_end	filename	mpi_io_file
MPI_File_get_position_start	filename offset	mpi_io_rw
MPI_File_get_position_end	filename offset	mpi_io_rw
MPI_File_get_position_shared_start	filename offset	mpi_io_rw
MPI_File_get_position_shared_end	filename offset	mpi_io_rw
MPI_File_get_size_start	filename size	mpi_io_file
MPI_File_get_size_end	filename size	mpi_io_file

TABLE D-2 TNF Probes, Associated Arguments, & Groups for MPI I/O Calls (Continued)

Probe	Argument(s)	Group(s), in Addition to mpi_api & mpi_io
MPI_File_get_type_extent_start	filename datatype extent	mpi_io_datarep
MPI_File_get_type_extent_end	filename datatype extent	mpi_io_datarep
MPI_File_get_view_start	filename disp etype filetype datarep_name	mpi_io_file
MPI_File_get_view_end	filename disp etype filetype datarep_name	mpi_io_file
MPI_File_iread_start	filename bytes	mpi_io_rw
MPI_File_iread_end	filename	mpi_io_rw
MPI_File_iread_at_start	filename offset	mpi_io_rw
MPI_File_iread_at_end	filename offset	mpi_io_rw
MPI_File_iread_shared_start	filename	mpi_io_rw
MPI_File_iread_shared_end	filename	mpi_io_rw
MPI_File_iwrite_start	filename	mpi_io_rw
MPI_File_iwrite_end	filename	mpi_io_rw
MPI_File_iwrite_at_start	filename offset	mpi_io_rw
MPI_File_iwrite_at_end	filename offset	mpi_io_rw
MPI_File_iwrite_shared_start	filename	mpi_io_rw
MPI_File_iwrite_shared_end	filename	mpi_io_rw
MPI_File_open_start	filename amode file_handle	mpi_io_file
MPI_File_open_end	filename amode file_handle	mpi_io_file

TABLE D-2 TNF Probes, Associated Arguments, & Groups for MPI I/O Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api & mpi_io
MPI_File_nonblocking_read_actual_end	filename offset count datatype	mpi_io_rw
MPI_File_nonblocking_write_actual_end	filename offset count datatype	mpi_io_rw
MPI_File_preallocate_start	filename size	mpi_io_file
MPI_File_preallocate_end	filename size	mpi_io_file
MPI_File_read_start	filename count datatype	mpi_io_rw
MPI_File_read_end	filename	mpi_io_rw
MPI_File_read_all_start	filename	mpi_io_rw
MPI_File_read_all_end	filename	mpi_io_rw
MPI_File_read_all_begin_start	filename	mpi_io_rw
MPI_File_read_all_begin_end	filename	mpi_io_rw
MPI_File_read_all_end_start	filename	mpi_io_rw
MPI_File_read_all_end_end	filename bytes	mpi_io_rw
MPI_File_read_at_start	filename offset	mpi_io_rw
MPI_File_read_at_end	filename offset	mpi_io_rw
MPI_File_read_at_all_start	filename offset	mpi_io_rw
MPI_File_read_at_all_end	filename offset	mpi_io_rw
MPI_File_read_at_all_begin_start	filename offset	mpi_io_rw
MPI_File_read_at_all_begin_end	filename offset	mpi_io_rw
MPI_File_read_at_all_end_start	filename	mpi_io_rw
MPI_File_read_at_all_end_end	filename bytes	mpi_io_rw

TABLE D-2 TNF Probes, Associated Arguments, & Groups for MPI I/O Calls (Continued)

Probe	Argument(s)	Group(s), in Addition to mpi_api & mpi_io
MPI_File_read_ordered_start	filename	mpi_io_rw
MPI_File_read_ordered_end	filename	mpi_io_rw
MPI_File_read_ordered_begin_start	filename	mpi_io_rw
MPI_File_read_ordered_begin_end	filename	mpi_io_rw
MPI_File_read_ordered_end_start	filename	mpi_io_rw
MPI_File_read_ordered_end_end	filename bytes	mpi_io_rw
MPI_File_read_shared_end	filename	mpi_io_rw
MPI_File_seek_start	filename offset whence	mpi_io_rw
MPI_File_seek_end	filename offset whence	mpi_io_rw
MPI_File_seek_shared_start	filename offset whence	mpi_io_rw
MPI_File_seek_shared_end	filename offset whence	mpi_io_rw
MPI_File_set_atomicsity_start	filename flag	mpi_io_consistency
MPI_File_set_atomicsity_end	filename flag	mpi_io_consistency
MPI_File_set_errhandler_start	filename	mpi_io_errhandler
MPI_File_set_errhandler_end	filename	mpi_io_errhandler
MPI_File_set_info_start	filename	mpi_io_file
MPI_File_set_info_end	filename	mpi_io_file
MPI_File_set_size_start	filename size	mpi_io_file
MPI_File_set_size_end	filename size	mpi_io_file
MPI_File_set_view_start	filename disp etype filetype datarep_name	mpi_io_file

TABLE D-2 TNF Probes, Associated Arguments, & Groups for MPI I/O Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api & mpi_io
MPI_File_set_view_end	filename disp etype filetype datarep_name	mpi_io_file
MPI_File_sync_start	filename	mpi_io_consistency
MPI_File_sync_end	filename	mpi_io_consistency
MPI_File_write_start	filename	mpi_io_rw
MPI_File_write_end	filename	mpi_io_rw
MPI_File_write_all_start	filename	mpi_io_rw
MPI_File_write_all_end	filename	mpi_io_rw
MPI_File_write_all_begin_start	filename	mpi_io_rw
MPI_File_write_all_begin_end	filename	mpi_io_rw
MPI_File_write_all_end_start	filename	mpi_io_rw
MPI_File_write_all_end_end	filename bytes	mpi_io_rw
MPI_File_write_at_start	filename offset	mpi_io_rw
MPI_File_write_at_end	filename offset	mpi_io_rw
MPI_File_write_at_all_start	filename offset	mpi_io_rw
MPI_File_write_at_all_end	filename offset	mpi_io_rw
MPI_File_write_at_all_begin_start	filename offset	mpi_io_rw
MPI_File_write_at_all_begin_end	filename offset	mpi_io_rw
MPI_File_write_at_all_end_start	filename	mpi_io_rw
MPI_File_write_at_all_end_end	filename bytes	mpi_io_rw
MPI_File_write_ordered_start	filename	mpi_io_rw
MPI_File_write_ordered_end	filename	mpi_io_rw
MPI_File_write_ordered_begin_start	filename	mpi_io_rw
MPI_File_write_ordered_begin_end	filename	mpi_io_rw

TABLE D-2 TNF Probes, Associated Arguments, & Groups for MPI I/O Calls *(Continued)*

Probe	Argument(s)	Group(s), in Addition to mpi_api & mpi_io
MPI_File_write_ordered_end_start	filename	mpi_io_rw
MPI_File_write_ordered_end_end	filename bytes	mpi_io_rw
MPI_File_write_shared_start	filename	mpi_io_rw
MPI_File_write_shared_end	filename	mpi_io_rw
MPI_Register_datarep_start	datarep_name	mpi_io_datarep
MPI_Register_datarep_end	datarep_name	mpi_io_datarep

Index

A

Argonne National Laboratory (ANL)
and MPE, 30
array sections, 9
attributes, with communicators, 8

B

blocking routines. *See* routines, blocking.
buffered mode. *See* modes for point-to-point
communication.

C

caching, with communicators, 8
Cartesian topology. *See* topology, Cartesian.
code samples. *See* sample programs.
collective communication. *See* communication,
collective.
communication
buffers, 9
collective, 7, 8, 12
in multithreaded programs, 15
restrictions, 7
“half-channel”, 10
interprocess, 8
persistent request, defined, 10
point-to-point, 8
port, 10
communicator

default, 8
defined, 8
and MPI I/O, 32
and multithreaded programming, 15 to 16
and process topologies, 10
compiling, 24 to 26
with profiling library, 17
See also include syntax.
context, defined, 8

D

-dalignoption, 26
data type
possible values for C, 10
derived (user-defined), 9, 32
possible values for Fortran, 9
primitive, 9
debugging, 29 to 30
with mpe, 30
See also dbx, Prism.
displacement (*disp*), 33, 35

E

elementary data type (*etype*), 32
environmental inquiry functions, 11
error handling, 11
and MPE, 18
and multithreaded programming, 16
error messages

standard error classes (Sun MPI I/O), 96
standard error values (Sun MPI), 95

F

features, 2
file type (`filetype`), 32
Fortran
 compiling with `-dalign` option, 26

G

graph topology. *See* topology, graph.
grid topology. *See* topology, Cartesian.
group, defined, 8

H

header files, 21
“holes” (in an MPI I/O file type), 32, 36

I

I/O. *See* Sun MPI I/O, MPI I/O.
include syntax, 21 to 22
intercommunicator, defined, 8
intracommunicator, 10
 defined, 8

L

libraries
 `libthread.so`, 27
 linking, 24 to 26
linking, 24 to 26
 with profiling library, 17

M

man pages
 Solaris, location, 27
 Sun MPI, location, xi

modes for point-to-point communication, 6
MPI

 Forum, URL, xi
 Mississippi State University URL, xi
 related publications, ix
 Standards, 1
 profiling, 16
 URL, xi
 University of New Mexico URL, xi
 See also Sun MPI.

MPI I/O, 31

 Sun MPI implementation.
 See Sun MPI I/O.

MPI_COMM_GROUP, 8

MPI_COMM_WORLD

 as default communicator, 8, 9

multiprocessing environment (MPE), ?? to 17, 18
 and debugging, 30

multithreaded programming, 13 to 16
 stubbing thread calls, 27

N

nonblocking routines. *See* routines, nonblocking.

O

offset, 33
options
 `-dalign`, 26

P

parallel file system (PFS), 31
persistent communication request. *See*
 communication, persistent request.
PFS. *See* parallel file system.
point-to-point
 communication. *See* communication, point-to-
 point.
 routines. *See* routines, point-to-point.
Prism, 29 to 30
 compilers to use, 25
process

relation to group, 8
process topologies, 10
profiling, 16 to 17

R

rank, of a process, 8, 11
ready mode. *See* modes for point-to-point communication.
receive. *See* routines, receive.
routines
 all-gather, 7
 all-to-all, 7
 barrier, 7
 basic six, 11
 blocking, 6, 7
 broadcast, 7
 collective, 7, 8
 in multithreaded programs, 15
 for constructing communicators, 8
 data access (MPI I/O), 35 to 39
 pointer manipulation, 37 to 38
 with explicit offsets, 36 to 37
 with individual file pointers, 37
 with shared file pointers, 38 to 39
 error-handling, 11
 file consistency (MPI I/O), 40
 file manipulation (MPI I/O), 34
 file views (MPI I/O), 35
 gather, 7
 for constructing groups, 8
 local, 8
 nonblocking, 6
 point-to-point, 6
 receive, 6, 12
 reduction, 7
 scan, 7
 scatter, 7
 semantics (MPI I/O), 40
 send, 6, 12
 Sun MPI
 listed alphabetically, 59 to 74
 listed by functional category, 49 to 58
 Sun MPI I/O
 listed alphabetically, 77 to 80
 listed by functional category, 74 to 76

S

sample programs
 Sun MPI, 22 to 24
 Sun MPI I/O, 41 to 48
send. *See* routines, send.
shutting down, 11
SPMD programs, 29
 defined, 29
standard mode. *See* modes for point-to-point communication.
starting up, 11
static libraries, and relinking, 26
Sun MPI
 contents of library, 2
 definition, 1
 See also MPI.
Sun MPI I/O, 31 to 48
synchronous mode. *See* modes for point-to-point communication.

T

thread safety. *See* multithreaded programming.
timers, 11
topology
 Cartesian, 11
 graph, 11
 virtual, defined, 11
 See also process topologies.

V

view, 33

W

WorkShop Compilers Fortran, 25

