

PARAGON

intel

User's Guide



June 1994

Order Number: 312489-003



**Paragon™**  
**User's Guide**



**Intel® Corporation**





## **WARNING**

Some of the circuitry inside this system operates at hazardous energy and electric shock voltage levels. To avoid the risk of personal injury due to contact with an energy hazard, or risk of electric shock, do not enter any portion of this system unless it is intended to be accessible without the use of a tool. The areas that are considered accessible are the outer enclosure and the area just inside the front door when all of the front panels are installed, and the front of the diagnostic station. There are no user serviceable areas inside the system. Refer any need for such access only to technical personnel that have been qualified by Intel Corporation.

## **CAUTION**

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

## **LIMITED RIGHTS**

The information contained in this document is copyrighted by and shall remain the property of Intel Corporation. Use, duplication or disclosure by the U.S. Government is subject to Limited Rights as set forth in subparagraphs (a)(15) of the Rights in Technical Data and Computer Software clause at 252.227-7013. Intel Corporation, 2200 Mission College Boulevard, Santa Clara, CA 95052. For all Federal use or contracts other than DoD Limited Rights under FAR 52.2272-14, ALT. III shall apply. Unpublished—rights reserved under the copyright laws of the United States.

The handler function you define must be written in C and must have four arguments of type **long**. These arguments are passed the following values when the function is called:

1. Type of the message (as returned by **infotype()**).
2. Length of the message in bytes (as returned by **infocount()**).
3. Node number of the process that sent the message (as returned by **infonode()**).
4. Process type of the process that sent the message (as returned by **infoptype()**).

For example, here's a C code fragment that attaches the functions *funct0()*, *funct1()*, and *funct2()* to message types 0, 1, and 2, respectively. The message types that have handlers are referred to as *handled types*.

```
#include <nx.h>

char buf0[100], buf1[100], buf2[100];
void funct0(), funct1(), funct2();

hrecv(0, buf0, sizeof(buf0), funct0);
hrecv(1, buf1, sizeof(buf1), funct1);
hrecv(2, buf2, sizeof(buf2), funct2);
.
. /* Now perform other work. No blocking happens. */
.
```

The declaration of **funct1()** looks like this (the other functions are similar):

```
void funct1(long type, long count, long node, long ptype)
{
.
.
.
}
```

When a message of type 1 arrives, the message is stored in the buffer specified in the **hrecv()** call (in this case, *buf1*), then **funct1()** is called with the type and length of the message and the node number and process type of the sender as arguments. **funct1()** and the main program then run concurrently until **funct1()** returns. (In previous releases of Paragon OSF/1, the main program was interrupted and did not run at all until **funct1()** returned.)

## CAUTION

The handler runs in the same memory space as the main program (but they have separate stacks).

# Preface

---

This manual tells how to use the Paragon™ OSF/1 operating system on an Intel supercomputer.

This manual assumes that you are an application programmer proficient in the C or Fortran language and the UNIX operating system. The manual provides you with enough detail to begin using your system.

## NOTE

Programming examples in this manual are intended to demonstrate the use of Paragon OSF/1 system calls, not as examples of good programming practice.

For example, in some cases, the return values of functions are not checked for error conditions. This is not recommended, but the error checks have been omitted in order to make the example shorter and easier to read.

## Organization

- |           |  |
|-----------|--|
| Chapter 1 | Provides an overview of the Paragon OSF/1 software and Intel supercomputer hardware.   |
| Chapter 2 | Describes the Paragon OSF/1 commands that you can enter at the shell prompt and the Paragon OSF/1 cross-development commands that run on supported workstations. |
| Chapter 3 | Describes the message-passing system calls available to programs in Paragon OSF/1.   |
| Chapter 4 | Describes the other general-purpose system calls available in Paragon OSF/1.   |
-

- Chapter 5 Describes the parallel I/O calls you can use for parallel access to the Intel supercomputer's file systems.
- Chapter 6 Describes the *pthread*s package, which you can use to create and control multiple threads (also called "lightweight processes") within your programs.
- Chapter 7 Tells how to prepare an application for the Paragon OSF/1 operating system. The steps described are applicable to applications that are written for a parallel computer and applications that are ported from a sequential computer. This chapter discusses three examples: an integration, a matrix\*vector multiplication, and the N-Queens problem.
- Chapter 8 Presents some techniques you can use to improve the performance of your parallel applications.
- Appendix A Summarizes the commands and system calls of Paragon OSF/1. The complete syntax of each command and call is provided, along with a brief description of each.
- Appendix B Describes the level of support offered by Paragon OSF/1 for the commands and system calls of the iPSC® system.

## Notational Conventions

This manual uses the following notational conventions:

- Bold** Identifies command names and switches, system call names, reserved words, and other items that must be used exactly as shown.
- Italic* Identifies variables, filenames, directories, partitions, user names, and writer annotations in examples. Italic type style is also occasionally used to emphasize a word or phrase.
- Plain-Monospace** Identifies computer output (prompts and messages), examples, and values of variables.
- Bold-Italic-Monospace*** Identifies user input (what you enter in response to some prompt).
- Bold-Monospace** Identifies the names of keyboard keys (which are also enclosed in angle brackets). A dash indicates that the key preceding the dash is to be held down *while* the key following the dash is pressed. For example:

<Break>

<s>

<Ctrl-Alt-Del>

- [ ] (Brackets) Surround optional items.
- ... (Ellipsis dots) Indicate that the preceding item may be repeated.
- | (Bar) Separates two or more items of which you may select only one.
- { } (Braces) Surround two or more items of which you must select one.

## Applicable Documents

For more information, refer to the following manuals. See the *Paragon™ System Technical Documentation Guide* for information on the complete Paragon document set and ordering information.

### Paragon™ Manuals

- *Paragon™ Commands Reference Manual*
- *Paragon™ Network Queueing System Manual*
- *Paragon™ C Compiler User's Guide*
- *Paragon™ Fortran Compiler User's Guide*
- *Paragon™ C System Calls Reference Manual*
- *Paragon™ Fortran System Calls Reference Manual*
- *Paragon™ Application Tools User's Guide*
- *Paragon™ Interactive Parallel Debugger Reference Manual*
- *Paragon™ System Administrator's Guide*

For information about limitations and workarounds, see the *Paragon™ System Software Release Notes for the Paragon™ XP/S System*. Release notes are also located in the directory `/vol/share/release_notes` on your Paragon system.

## Other Manuals

- *OSF/1 User's Guide*
- *OSF/1 Programmer's Reference*
- *OSF/1 Command Reference*
- *Effective Fortran 77* - Michael Metcalf
- *C: A Reference Manual* - Harbison and Steele
- *The C Programming Language* - Kernighan and Ritchie
- *CLASSPACK Basic Math Library User's Guide* - Kuck & Associates
- *CLASSPACK Basic Math Library/C User's Guide* - Kuck & Associates

## Comments and Assistance

Intel Supercomputer Systems Division is eager to hear of your experiences with our products. Please call us if you need assistance, have questions, or otherwise want to comment on your Paragon system.

**U.S.A./Canada Intel Corporation**  
**phone: 800-421-2823**  
**Internet: [support@ssd.intel.com](mailto:support@ssd.intel.com)**

---

**Intel Corporation Italia s.p.a.**  
 Milanofiori Palazzo  
 20090 Assago  
 Milano  
 Italy  
 1678 77203 (toll free)

**France Intel Corporation**  
 1 Rue Edison-BP303  
 78054 St. Quentin-en-Yvelines Cedex  
 France  
 0590 8602 (toll free)

**Intel Japan K.K.**  
**Supercomputer Systems Division**  
 5-6 Tokodai, Tsukuba City  
 Ibaraki-Ken 300-26  
 Japan  
 0298-47-8904

**United Kingdom Intel Corporation (UK) Ltd.**  
**Supercomputer System Division**  
 Pipers Way  
 Swindon SN3 IRJ  
 England  
 0800 212665 (toll free)  
 (44) 793 491056 (*answered in French*)  
 (44) 793 431062 (*answered in Italian*)  
 (44) 793 480874 (*answered in German*)  
 (44) 793 495108 (*answered in English*)

**Germany Intel Semiconductor GmbH**  
 Dornacher Strasse 1  
 8016 Feldkirchen bel Muenchen  
 Germany  
 0130 813741 (toll free)

---

**World Headquarters**  
**Intel Corporation**  
**Supercomputer Systems Division**  
 15201 N.W. Greenbrier Parkway  
 Beaverton, Oregon 97006  
 U.S.A.

(503) 629-7600 (Monday through Friday, 8 AM to 5 PM Pacific Time)  
 Fax: (503) 629-9147

If you have comments about our manuals, please fill out and mail the enclosed Comment Card. You can also send your comments electronically to the following address:

**[techpubs@ssd.intel.com](mailto:techpubs@ssd.intel.com)**





# Table of Contents

---

## Chapter 1 Introduction

<b>Introduction</b> .....	1-1
<b>System Hardware</b> .....	1-2
Nodes .....	1-2
Node Interconnect Network .....	1-2
I/O Interfaces .....	1-2
Front Panel LEDs (Paragon™ XP/S Only) .....	1-3
<b>System Software</b> .....	1-4
Paragon™ OSF/1 Operating System .....	1-4
User Model .....	1-5
Programming Model .....	1-6
Cross-Development Facility .....	1-6

## Chapter 2

# Using Paragon™ OSF/1 Commands

<b>Introduction</b> .....	2-1
Terminology .....	2-1
Using Paragon™ OSF/1 Commands on the Intel Supercomputer .....	2-2
Using Paragon™ OSF/1 Commands on Workstations .....	2-2
A Quick Example .....	2-3
Information You Need .....	2-3
Compiling, Linking, and Executing an Application .....	2-3
<b>Compiling and Linking Applications</b> .....	2-5
Configuring Your Environment for Cross-Development .....	2-6
Tips for Compiling and Linking .....	2-8
Using Other Switches .....	2-8
Including <i>nx.h</i> or <i>fnx.h</i> .....	2-8
Specifying Include File and Library Pathnames .....	2-8
Preprocessing a Fortran Program .....	2-9
Order of Switches .....	2-10
<b>Running Applications</b> .....	2-11
I/O Redirection .....	2-12
Controlling the Application's Execution Characteristics .....	2-13
Using the Default Partition .....	2-14
Setting Your Default Partition .....	2-14
Determining the Current Default Partition .....	2-15
Specifying Application Size .....	2-15
Specifying a Rectangle of Nodes .....	2-16
Specifying a Particular Rectangle of Nodes .....	2-16
Using the Default Size .....	2-17
Specifying Application Priority .....	2-17
Specifying Process Type .....	2-18
Running a Program on a Subset of the Nodes .....	2-18

Running Applications Consisting of Multiple Programs .....	2-21
Running an Application in a Particular Partition .....	2-22
<b>Managing Running Applications .....</b>	<b>2-23</b>
<b>Managing Partitions .....</b>	<b>2-25</b>
Special Partitions .....	2-26
The Root Partition .....	2-26
The Service Partition .....	2-27
The Compute Partition .....	2-27
Partition Pathnames .....	2-28
Partition Characteristics .....	2-29
Parent Partition .....	2-29
Partition Name .....	2-30
Nodes Allocated to the Partition .....	2-30
Node Numbers Within a Partition .....	2-30
Unusable Nodes .....	2-31
Owner, Group, and Protection Modes .....	2-32
Scheduling Characteristics .....	2-33
Standard Scheduling .....	2-34
Gang Scheduling .....	2-35
Space Sharing .....	2-37
Summary of Scheduling Types .....	2-38
A Scheduling Example .....	2-39
Making Partitions .....	2-39
Specifying the Nodes Allocated to the Partition .....	2-40
Specifying Protection Modes .....	2-42
Specifying Scheduling Characteristics .....	2-43
Removing Partitions .....	2-45
Removing Partitions Containing Running Applications .....	2-45
Removing Partitions Containing Subpartitions .....	2-45

Showing Partition Characteristics .....2-46

    Showing Free Nodes .....2-48

Listing Subpartitions .....2-49

    Recursively Listing Subpartitions .....2-50

Listing the Applications in a Partition .....2-51

    Applications in Subpartitions .....2-52

    Recursively Listing Applications in Subpartitions .....2-54

Changing Partition Characteristics .....2-54

## **Chapter 3**

# **Using Paragon™ OSF/1 Message-Passing System Calls**

**Introduction** .....3-1

**Process Characteristics** .....3-3

    Node Numbers .....3-3

    Process Types .....3-4

**Message Characteristics** .....3-5

    Message Length .....3-5

    Message Type .....3-6

    Message ID .....3-6

**Message Order** .....3-7

**Names of Send and Receive Calls** .....3-7

**Synchronous Send and Receive** .....3-8

    Synchronous Send to Multiple Nodes .....3-9

**Asynchronous Send and Receive** .....3-10

    Releasing Message IDs .....3-12

    Merging Message IDs .....3-13

**Probing for Pending Messages** .....3-14

<b>Getting Information About Pending or Received Messages .....</b>	<b>3-15</b>
<b>Message Passing with Fortran Commons .....</b>	<b>3-17</b>
<b>Treating a Message as an Interrupt .....</b>	<b>3-18</b>
Passing Information to the Handler .....	3-20
Preventing Interrupts .....	3-22
<b>Extended Receive and Probe .....</b>	<b>3-24</b>
<b>Global Operations .....</b>	<b>3-27</b>

## Chapter 4

# Using Other Paragon™ OSF/1 System Calls

<b>Introduction .....</b>	<b>4-1</b>
<b>Managing Applications .....</b>	<b>4-2</b>
Controlling Application Execution with System Calls .....	4-3
Creating an Application with nx_initve() .....	4-4
Creating a Rectangular Application with nx_initve_rect() .....	4-7
Setting an Application's Priority with nx_pri() .....	4-9
Copying a Process onto the Nodes with nx_nfork() .....	4-10
Loading a Program onto the Nodes with nx_load() .....	4-11
Loading a Program onto the Nodes with nx_loadve() .....	4-13
Waiting for Application Processes with nx_waitall() .....	4-14
Using PIDs .....	4-14
Getting Information About Applications .....	4-16
Finding an Application's Shape with nx_app_rect() .....	4-16
Listing an Application's Nodes with nx_app_nodes() .....	4-17
Listing the Applications in a Partition with nx_pspart() .....	4-19

The Controlling Process .....	4-21
Process Groups .....	4-22
Process Groups in Paragon™ OSF/1 .....	4-23
Killing Application Processes .....	4-23
An Example Controlling Process .....	4-23
Message Passing Between Controlling Process and Application Processes .....	4-25
<b>Managing Partitions .....</b>	<b>4-27</b>
Making Partitions .....	4-28
Removing Partitions .....	4-30
Getting Information About Partitions .....	4-31
Determining a Partition's Attributes with <code>nx_part_attr()</code> .....	4-32
Determining a Partition's Nodes with <code>nx_part_nodes()</code> .....	4-35
Changing Partition Characteristics .....	4-36
<b>Listing Unusable Nodes .....</b>	<b>4-40</b>
<b>Handling Errors .....</b>	<b>4-42</b>
Underscore Calls .....	4-42
Core Dumps .....	4-44
<b>Controlling Floating-Point Behavior .....</b>	<b>4-46</b>
Detecting Not-a-Number .....	4-47
Controlling Floating-Point Behavior .....	4-47
Rounding Mode .....	4-47
Exception Mask and Sticky Flags .....	4-48
Fortran Exception Mask Values .....	4-49
<b>Miscellaneous Calls .....</b>	<b>4-50</b>
Temporarily Releasing Control of the Processor .....	4-50
Timing Execution .....	4-50
<b>iPSC® and Touchstone DELTA Compatibility Calls .....</b>	<b>4-52</b>

## Chapter 5

# Using Parallel File I/O

<b>Introduction</b> .....	5-1
Disks and File Systems .....	5-2
PFS File Systems and PFS Files .....	5-3
PFS Filenames and Pathnames .....	5-3
PFS Limitations .....	5-4
<b>Using PFS Commands</b> .....	5-5
Displaying File System Attributes .....	5-5
Increasing the Size of a File .....	5-7
<b>Using Parallel I/O Calls</b> .....	5-8
<b>Opening Files in Parallel</b> .....	5-9
Using gopen() in C .....	5-10
Using gopen() in Fortran .....	5-10
Opening Files with Standard Operations .....	5-11
Special Considerations for Fortran .....	5-11
Formatted Versus Unformatted I/O .....	5-11
New Files .....	5-12
Unnamed Files .....	5-13
<b>Using I/O Modes</b> .....	5-13
M_UNIX (Mode 0) .....	5-14
M_LOG (Mode 1) .....	5-15
M_SYNC (Mode 2) .....	5-15
M_RECORD (Mode 3) .....	5-16
M_GLOBAL (Mode 4) .....	5-17
An I/O Mode Example .....	5-17
Fortran Example .....	5-18
C Example .....	5-19
Compiling and Running the Example .....	5-20

M_UNIX Output .....	5-21
M_LOG Output .....	5-22
M_SYNC Output .....	5-22
M_RECORD Output .....	5-23
M_GLOBAL Output .....	5-23
<b>Reading and Writing Files in Parallel .....</b>	<b>5-24</b>
Synchronous File I/O .....	5-25
Asynchronous File I/O .....	5-27
<b>Closing Files in Parallel .....</b>	<b>5-28</b>
<b>Detecting End-of-File and Moving the File Pointer .....</b>	<b>5-29</b>
<b>Flushing Fortran Buffered I/O .....</b>	<b>5-30</b>
<b>Using “###” Filenames .....</b>	<b>5-31</b>
<b>Increasing the Size of a File .....</b>	<b>5-32</b>
<b>Using Extended Files .....</b>	<b>5-33</b>
OSF/1 Calls that Do Not Support Extended Files .....	5-34
OSF/1 Commands that Do Not Support Extended Files .....	5-35
Manipulating Extended Files .....	5-36
Performing Extended Arithmetic .....	5-37
<b>Getting Information About PFS File Systems .....</b>	<b>5-39</b>
Getting Information About All Mounted PFS File Systems .....	5-39
Getting PFS Information About a Single File System .....	5-41
<b>Controlling Tape Devices .....</b>	<b>5-44</b>
Naming Tape Devices .....	5-44
Performing Operations on Tape Devices .....	5-45
Getting Status of Tape Devices .....	5-46
<b>Synchronization Summary .....</b>	<b>5-48</b>



## Chapter 6

# Using Pthreads

<b>Introduction</b> .....	6-1
The Pthreads Package .....	6-1
What's In This Chapter .....	6-2
<b>Limitations of Pthreads</b> .....	6-3
<b>Recommended Safe Operating Environment</b> .....	6-4
<b>Compiling and Linking a Pthread Application</b> .....	6-5
<b>Using Reentrant C Library Calls</b> .....	6-6
<b>Using Pthreads Library Calls</b> .....	6-11
Pthreads Library Data Types and Symbols .....	6-11
The Main Thread .....	6-12
Managing Pthread Execution .....	6-13
Managing Pthread Attributes .....	6-15
Managing Mutexes .....	6-16
Managing Mutex Attributes .....	6-17
An Example Pthreads Program .....	6-18
Using Condition Variables to Synchronize Pthreads .....	6-21
Managing Condition Attributes .....	6-23
Examples of Condition Variables .....	6-24
Canceling Pthreads .....	6-28
Cancelability States .....	6-28
Cancellation Examples .....	6-30
Pthreads Cleanup Routines .....	6-32
Managing Pthread Keys .....	6-33
Executing a Routine Once .....	6-34
Managing Signals .....	6-34
<b>Interfacing with Non-Thread-Safe Code</b> .....	6-37
<b>Message Passing and Pthreads Library Calls</b> .....	6-37

<b>File I/O and Pthreads Library Calls .....</b>	<b>6-38</b>
<b>nx_nfork() and nx_initve() and Pthreads Library Calls .....</b>	<b>6-39</b>
<b>Signals and Pthreads Library Calls .....</b>	<b>6-39</b>
Signal Types .....	6-39
Signals are a Per-Process Resource .....	6-40
Dealing with Signals .....	6-41
<b>Handling Errors .....</b>	<b>6-41</b>
errno Confusion .....	6-41
perror() and nx_perror() .....	6-42
Calling exit() .....	6-42
Use of Underscore Versions of Paragon System Calls .....	6-43
Catch Signals Causing Core Dump by Default .....	6-43
When One Pthread Hangs .....	6-43

## Chapter 7

### Designing a Parallel Application

<b>Introduction .....</b>	<b>7-1</b>
<b>The Paragon™ OSF/1 Programming Model .....</b>	<b>7-2</b>
<b>Parallel Programming Techniques .....</b>	<b>7-2</b>
Separating the User Interface from the Computation .....	7-3
Balancing the Load .....	7-3
Domain Decomposition .....	7-3
Control Decomposition .....	7-5
Making the Program Independent of the Number of Nodes .....	7-5
Designing Your Communication Strategy .....	7-6
Using Global Operations .....	7-6
Using Alternate Node Topologies .....	7-6

<b>Example Application: Calculating pi .....</b>	<b>7-7</b>
<b>Example Application: Matrix*Vector Multiplication .....</b>	<b>7-11</b>
<b>Example Application: The N-Queens Problem .....</b>	<b>7-13</b>

## Chapter 8

# Improving Performance

<b>Introduction .....</b>	<b>8-1</b>
<b>Single Node Performance .....</b>	<b>8-2</b>
Use Profiling Tools .....	8-2
Avoid Repeated Use of System Calls .....	8-2
Avoid Virtual Memory Paging .....	8-3
Use Compiler Optimizations .....	8-3
Increase Problem Size .....	8-5
Access Contiguous Memory Locations .....	8-5
Use Caching Wisely .....	8-5
Use Optimized Libraries .....	8-6
Use Assembly Language Subroutines .....	8-7
Avoid Error Checking (C Language Only) .....	8-7
<b>Multi-Node Performance .....</b>	<b>8-7</b>
Use Dynamic Memory Allocation for Large Arrays .....	8-8
Avoid Serializing Calls .....	8-9
Use ParaGraph .....	8-10
Maintain Data Locality .....	8-10
Overlap Computation and Communication .....	8-10
Avoid Message Buffering .....	8-11
Align Application Buffers .....	8-12

Understand Message-Passing Flow Control .....8-13

    Overview of Message-Passing Flow Control .....8-14

    Process Locking .....8-15

    Packetization .....8-16

    System Message Buffers .....8-16

    Message-Passing Configuration Switches .....8-18

        Summary of the Message-Passing Configuration Switches .....8-19

        Default, Maximum, and Minimum Values .....8-20

        Dependencies and Rounding .....8-21

    Recommendations .....8-21

**I/O Performance .....8-23**

    Use PFS File Systems .....8-23

    Use gopen() Instead of open() .....8-23

    Use Parallel I/O Calls .....8-24

    Use Asynchronous Calls .....8-24

    Use the Appropriate I/O Mode .....8-24

    Align I/O Buffers with Virtual Memory Pages .....8-25

    Read or Write Whole File System Blocks .....8-25

    Make Good Use of File Striping .....8-25

## Appendix A

# Summary of Commands and System Calls

**Command Summary ..... A-1**

    Compiling and Linking Applications ..... A-1

    Running Applications ..... A-2

    Managing Partitions ..... A-2

    Parallel File System Commands ..... A-3

    Miscellaneous Commands ..... A-3

<b>C System Call Summary</b> .....	A-4
Process Characteristics .....	A-4
Synchronous Send and Receive .....	A-5
Asynchronous Send and Receive .....	A-6
Probing for Pending Messages .....	A-7
Getting Information About Pending or Received Messages .....	A-7
Treating a Message as an Interrupt .....	A-8
Extended Receive and Probe .....	A-9
Global Operations .....	A-10
Controlling Application Execution .....	A-12
Getting Information About Applications .....	A-13
Partition Management .....	A-14
Finding Unusable Nodes .....	A-15
Handling Errors .....	A-15
Floating-Point Control .....	A-16
Miscellaneous Calls .....	A-16
iPSC® and Touchstone DELTA Compatibility .....	A-17
I/O Modes .....	A-18
Reading and Writing Files in Parallel .....	A-19
Detecting End-of-File and Moving the File Pointer .....	A-20
Increasing the Size of a File .....	A-20
Extended File Manipulation .....	A-21
Performing Extended Arithmetic .....	A-22
Getting Information About PFS File Systems .....	A-23
Managing Pthread Execution .....	A-24
Managing Pthread Attributes .....	A-24
Managing Mutexes .....	A-25
Using Condition Variables to Synchronize Pthreads .....	A-26
Canceling Pthreads .....	A-26
Pthreads Cleanup Routines .....	A-27
Managing Pthread Keys .....	A-27
Miscellaneous Pthread Calls .....	A-27

<b>Fortran System Call Summary</b> .....	A-28
Process Characteristics .....	A-28
Synchronous Send and Receive .....	A-29
Asynchronous Send and Receive .....	A-30
Probing for Pending Messages .....	A-31
Getting Information About Pending or Received Messages .....	A-31
Treating a Message as an Interrupt .....	A-32
Extended Receive and Probe .....	A-33
Global Operations .....	A-35
Controlling Application Execution .....	A-38
Getting Information About Applications .....	A-39
Partition Management .....	A-40
Finding Unusable Nodes .....	A-42
Handling Errors .....	A-42
Floating-Point Control .....	A-42
Miscellaneous Calls .....	A-43
iPSC® and Touchstone DELTA Compatibility .....	A-43
I/O Modes .....	A-45
Reading and Writing Files in Parallel .....	A-45
Detecting End-of-File and Moving the File Pointer .....	A-47
Flushing Fortran Buffered I/O .....	A-47
Increasing the Size of a File .....	A-47
Extended File Manipulation .....	A-48
Performing Extended Arithmetic .....	A-49

## Appendix B

# iPSC® System Compatibility

<b>Introduction</b> .....	B-1
<b>General Compatibility Issues</b> .....	B-1
<b>New Features</b> .....	B-2
<b>Compilers</b> .....	B-4
<b>Commands</b> .....	B-5
Cube Control Commands .....	B-5
CFS Commands .....	B-7
System Administration Commands .....	B-7
Remote Host Commands .....	B-8
Miscellaneous Commands .....	B-8
<b>System Calls</b> .....	B-9
Include Files .....	B-9
Host Calls .....	B-9
Byte-Swapping Calls .....	B-14
Floating-Point Control Calls .....	B-15
CFS Calls .....	B-15
Miscellaneous Calls .....	B-16
<b>Summary</b> .....	B-17

## List of Illustrations

Figure 1-1.	Front Panel LEDs (Paragon™ XP/S Only) .....	1-3
Figure 1-2.	Node Activity LEDs .....	1-4
Figure 2-1.	The Root Partition of a 32-Node System .....	2-27
Figure 2-2.	Node Numbers in Contiguous and Noncontiguous Partitions .....	2-31
Figure 2-3.	Node Numbers in Overlapping Partitions .....	2-32
Figure 4-1.	Sample Partition for nx_part_attr() and nx_part_nodes() .....	4-34
Figure 7-1.	Using Domain Decomposition to Achieve Load Balancing .....	7-4
Figure 7-2.	The Decomposition Used for the pi Example .....	7-9
Figure 7-3.	The N-Queens Solution Tree for a 4 x 4 Board .....	7-15
Figure 8-1.	Two Methods of Improving I/O Performance with M_RECORD .....	8-27



## List of Tables

Table 2-1.	Summary of Scheduling Types .....	2-38
Table 5-1.	File Operations that Accept “###” Filenames .....	5-31
Table 5-2.	OSF/1 Calls Not Supporting Extended Files .....	5-34
Table 5-3.	OSF/1 Commands Not Supporting Extended Files .....	5-35
Table 5-4.	Synchronization in Each I/O Mode .....	5-48
Table 5-5.	File I/O Calls that Synchronize .....	5-48
Table 6-1.	Calls in Reentrant C Library (libc_r.a) .....	6-7
Table 8-1.	Message-Passing Configuration Switches .....	8-20
Table A-1.	Commands for Compiling and Linking Applications .....	A-1
Table A-2.	Commands for Running Applications .....	A-2
Table A-3.	Commands for Managing Partitions .....	A-2
Table A-4.	Parallel File System Commands .....	A-3
Table A-5.	Miscellaneous Commands .....	A-3
Table A-6.	C Calls for Process Characteristics .....	A-4
Table A-7.	C Calls for Synchronous Send and Receive .....	A-5
Table A-8.	C Calls for Asynchronous Send and Receive .....	A-6
Table A-9.	C Calls for Probing for Pending Messages .....	A-7
Table A-10.	C Calls for Getting Information About Pending or Received Messages .....	A-7
Table A-11.	C Calls for Treating a Message as an Interrupt .....	A-8
Table A-12.	C Calls for Extended Receive and Probe .....	A-9
Table A-13.	C Calls for Global Operations .....	A-10
Table A-14.	C Calls for Controlling Application Execution .....	A-12
Table A-15.	C Calls for Getting Information About Applications .....	A-13
Table A-16.	C Calls for Partition Management .....	A-14
Table A-17.	C Calls for Finding Unusable Nodes .....	A-15
Table A-18.	C Calls for Handling Errors .....	A-15

## List of Tables

Table A-19.	C Calls for Floating-Point Control .....	A-16
Table A-20.	Miscellaneous C Calls .....	A-16
Table A-21.	C Calls for iPSC® and Touchstone DELTA Compatibility .....	A-17
Table A-22.	C Calls for I/O Modes .....	A-18
Table A-23.	C Calls for Reading and Writing Files in Parallel .....	A-19
Table A-24.	C Calls for Detecting End-of-File and Moving the File Pointer .....	A-20
Table A-25.	C Calls for Increasing the Size of a File .....	A-20
Table A-26.	C Calls for Extended File Manipulation .....	A-21
Table A-27.	C Calls for Performing Extended Arithmetic .....	A-22
Table A-28.	C Calls for Getting Information About PFST™ File Systems .....	A-23
Table A-29.	C Calls for Managing Pthread Execution .....	A-24
Table A-30.	C Calls for Managing Pthread Attributes .....	A-24
Table A-31.	C Calls for Managing Mutexes .....	A-25
Table A-32.	C Calls for Using Condition Variables to Synchronize Pthreads .....	A-26
Table A-33.	C Calls for Canceling Pthreads .....	A-26
Table A-34.	C Calls for Pthreads Cleanup Routines .....	A-27
Table A-35.	C Calls for Managing Pthread Keys .....	A-27
Table A-36.	Miscellaneous Pthread Calls .....	A-27
Table A-37.	Fortran Calls for Process Characteristics .....	A-28
Table A-38.	Fortran Calls for Synchronous Send and Receive .....	A-29
Table A-39.	Fortran Calls for Asynchronous Send and Receive .....	A-30
Table A-40.	Fortran Calls for Probing for Pending Messages .....	A-31
Table A-41.	Fortran Calls for Getting Information About Pending or Received Messages .....	A-31
Table A-42.	Fortran Calls for Treating a Message as an Interrupt .....	A-32
Table A-43.	Fortran Calls for Extended Receive and Probe .....	A-33
Table A-44.	Fortran Calls for Global Operations .....	A-35
Table A-45.	Fortran Calls for Controlling Application Execution .....	A-38
Table A-46.	Fortran Calls for Getting Information About Applications .....	A-39
Table A-47.	Fortran Calls for Partition Management .....	A-40
Table A-48.	Fortran Calls for Finding Unusable Nodes .....	A-42

## List of Tables

Table A-49.	Fortran Calls for Handling Errors .....	A-42
Table A-50.	Fortran Calls for Floating-Point Control .....	A-42
Table A-51.	Miscellaneous Fortran Calls .....	A-43
Table A-52.	Fortran Calls for iPSC® and Touchstone DELTA Compatibility .....	A-43
Table A-53.	Fortran Calls for I/O Modes .....	A-45
Table A-54.	Fortran Calls for Reading and Writing Files in Parallel .....	A-45
Table A-55.	Fortran Calls for Detecting End-of-File and Moving the File Pointer .....	A-47
Table A-56.	Fortran Calls for Flushing Buffered I/O .....	A-47
Table A-57.	Fortran Calls for Increasing the Size of a File .....	A-47
Table A-58.	Fortran Calls for Extended File Manipulation .....	A-48
Table A-59.	Fortran Calls for Performing Extended Arithmetic .....	A-49
Table B-1.	Unsupported iPSC® System Byte-Swapping Calls .....	B-14
Table B-2.	Summary of Unsupported iPSC® System Commands .....	B-17
Table B-3.	Summary of Unsupported iPSC® System Calls .....	B-19



## Introduction

This chapter introduces the Paragon™ OSF/1 operating system and the hardware it runs on.

In an Intel supercomputer, a large number of processors called *nodes* work concurrently on the parts of a problem. Each node can run multiple processes, and each process can have multiple *threads*. The processes and threads on each node time-share the node's processor, using the standard OSF/1 scheduling mechanisms. Each process can be a stand-alone program (such as a shell, compiler, or editor), or can be part of a *parallel application*.

A parallel application consists of a group of closely related processes that work together on a single problem. They synchronize their actions and share information by passing *messages*, which are created and controlled by special Paragon OSF/1 system calls.

The processes in an application can also share disk files; Paragon OSF/1 *parallel I/O calls* insure that access to these files is efficient and properly synchronized.

## System Hardware

The Paragon OSF/1 operating system runs on several models of Intel supercomputers. These systems all have a large number of *nodes* connected by a high-speed *node interconnect network*, and a number of *I/O interfaces* to communicate with the outside world.

### Nodes

Each node is essentially a separate computer, with one or more i860® processors and 16M bytes or more of memory. Nodes can run distinct programs and have distinct memory spaces. They can team up to work on the same problem and exchange data by passing messages. An Intel supercomputer can have up to 2000 nodes. Each node can run more than one process at the same time; these processes can belong to the same or different applications.

The system administrator can choose to dedicate some nodes to interactive processes, such as shells and editors, and other nodes to compute-intensive applications. The nodes used for interactive processes are called *service nodes*, and the nodes used for compute-intensive applications are called *compute nodes*. However, there are no physical differences between these two types of nodes.

### Node Interconnect Network

The nodes are connected by a high-speed *node interconnect network*. Each node interfaces to this network through special hardware that monitors the network and extracts only those messages addressed to its attached node. Messages addressed to other nodes are passed on without interrupting the node processor. For most applications, you can think of each node as being fully connected to all the other nodes.

### I/O Interfaces

Some nodes are equipped with a SCSI interface, Ethernet interface, or other I/O connection. These nodes manage the system's disk and tape drives, network connections, and other I/O facilities. Nodes with I/O interfaces communicate with the other nodes over the node interconnect network. However, this access is transparent: processes on nodes without I/O hardware access the I/O facilities using standard OSF/1 system calls, just as though they were directly connected. Nodes with I/O interfaces are otherwise identical to nodes without I/O interfaces, and can run user processes.

## Front Panel LEDs (Paragon™ XP/S Only)

On the Paragon XP/S system, each cabinet has a number of Light-Emitting Diodes (LEDs) on its front panel that inform you of the status of the system, the nodes, and messages between nodes. The front panel LEDs are shown in Figure 1-1.

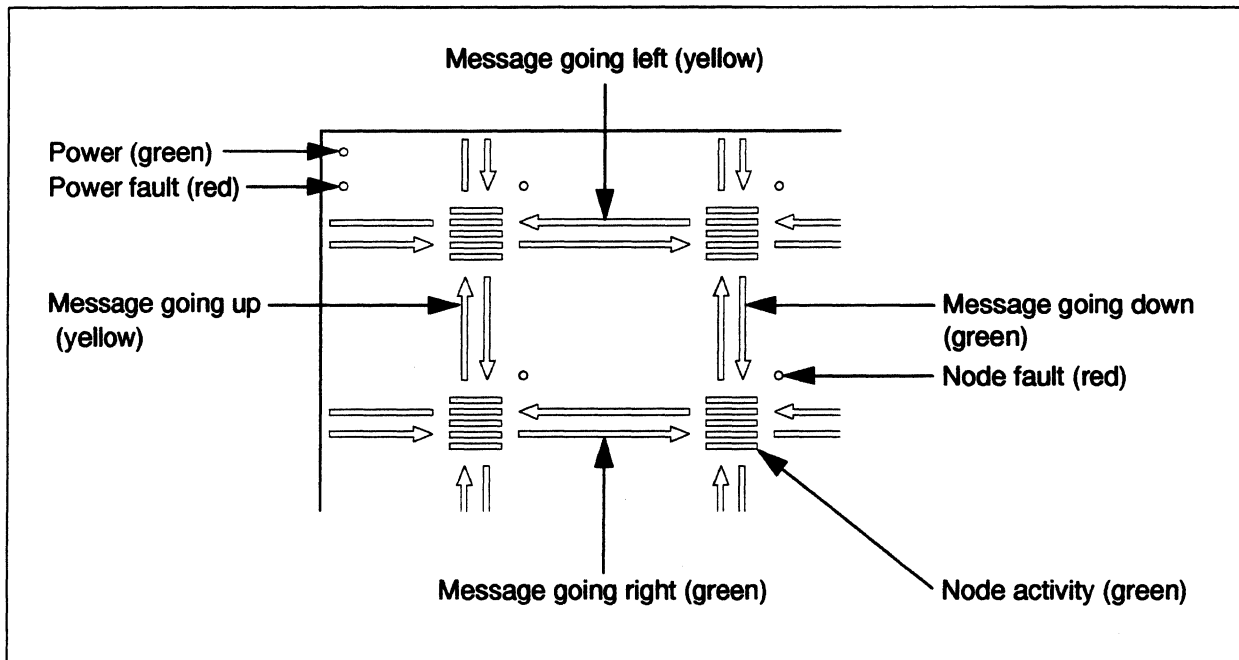



Figure 1-1. Front Panel LEDs (Paragon™ XP/S Only)

Each cabinet has four LED panels, each of which shows the status of 16 nodes in a 4 by 4 grid. Figure 1-1 shows the upper left corner of one LED panel. The meanings of the LEDs are as follows:

- The round green LED in the upper left corner of the top LED panel in each cabinet indicates that power has been supplied to the cabinet. (The corresponding LEDs in the other three panels never illuminate.)
- The round red LED just below the green power LED indicates a fault in the cabinet's power subsystem. If a fault is detected by the cabinet's self-tests, this LED illuminates. (The corresponding LEDs in the other three panels never illuminate.)
- The square groups of horizontal green LED bars show the amount of computational activity on the nodes. Each group represents one node. The more active a node is, the more green LEDs are illuminated, in a bar graph moving out from the center. Figure 1-2 shows the six possible ways these LEDs can be illuminated, showing activity levels from 0% to 100%.



**Figure 1-2. Node Activity LEDs**

- The arrow-shaped yellow and green LED bars indicate messages. When a message is passed from one node to another, all the arrow LEDs along its path illuminate. (Messages always travel first in the X direction (horizontally), then in the Y direction (vertically). Messages never change direction more than once.) Yellow arrows show messages going up or to the left; green arrows show messages going down or to the right. When the arrows are illuminated, a light pattern moves along the arrow to show the direction of motion.
- The round red LED associated with each node indicates a hardware fault on the node. If a fault is detected by the node's self-tests, the red LED illuminates.

## System Software

The nodes run the *Paragon OSF/1* operating system, based on the OSF/1 operating system from the Open Software Foundation. The same operating system runs on every node. OSF/1 is a version of the UNIX operating system that supports most industry standards; Paragon OSF/1 is an extended version of OSF/1 with enhancements to support parallel processing.

The Intel supercomputer also comes with a *cross-development facility*, which you can use to compile and link Paragon OSF/1 programs on supported workstations.

## Paragon™ OSF/1 Operating System

Paragon OSF/1 provides all the standard features of OSF/1, with extensions to provide a *single system image* across multiple nodes. This single system image makes all the nodes appear to be one large system. For example, all the nodes share a single file system, all the nodes have equal access to the system's I/O devices, and process identifiers (PIDs) are unique throughout the system. A process on one node can pipe its output to a process on another node, and the command `kill pid` on any node kills the specified process, no matter which node the process is running on.



The single system image does *not* combine all the nodes' memory into a single address space. Rather, each process has its own address space. The physical memory available to each process is limited to the memory of the node on which it is running. However, because OSF/1 provides *virtual memory*, a process's address space can be up to 2G bytes in size; memory pages that do not fit in physical memory are paged to disk. As in most multi-user systems, the address spaces of the different processes on the system are completely independent, unless two or more processes make special shared virtual memory calls to explicitly share part of their memory.

In addition to the standard facilities of OSF/1, the Paragon OSF/1 operating system provides message passing capability, Parallel File System access, and various other utilities to programs running on the Intel supercomputer. With Paragon OSF/1 calls, your programs can perform the following functions:

- Exchange messages with processes running on other nodes (or the same node).
- Read and write files on the Intel supercomputer's Parallel File System.
- Perform 64-bit integer arithmetic.
- Find out information about the computing environment.
- Perform global operations.
- Create and control parallel applications and partitions.

## User Model

The Paragon OSF/1 operating system is a complete implementation of OSF/1, and provides a full range of services, commands, and system calls. It has its own file system, shells, compilers, editors, network connections, and all the other features needed in a stand-alone computer system. It also supports NFS, the Network File System, so it can share data with other systems on your network. You can edit and compile programs, send and receive mail, read online manual pages, and do all your other daily work on the Intel supercomputer.

You access the Intel supercomputer by logging into a separate computer (typically your UNIX workstation) and then connecting to the Intel supercomputer over a local-area network, using a command such as **rlogin** or **telnet**. The Intel supercomputer does not have any dedicated hardware terminals.

You compile and link your application with the self-hosted Paragon OSF/1 compilers and linker. You then execute your application on the nodes of the Intel supercomputer simply by typing the application's name on the shell command line. Command-line switches, or arguments to system calls in the program, determine the number of nodes on which the application executes.

When you run an application, it runs in a *partition*. A partition is a group of nodes with an associated set of parameters that controls some of the run-time characteristics of the applications within it. You can use commands or system calls to create, modify, and remove partitions. However, the operations you are allowed to perform on your system's partitions may be restricted by the policies of your site.

The Paragon OSF/1 operating system also provides a suite of program development tools, such as a debugger, profiler, and parallel performance analysis tools. These tools are described in the *Paragon™ Application Tools User's Guide*.

## Programming Model

The most common programming model used with Paragon OSF/1 is the "single program, multiple data" (SPMD) model. In this model, the same program runs on each node in the application, but each node works on only part of the data.

- For some problems, called "perfectly parallel" problems, each node can do its work without access to data held by other nodes. In this case, each node operates completely independently.
- For other types of problems, each node needs data from other nodes to do its work. In this case, the nodes can share data by passing messages. Messages can also be used to synchronize node operations.

Because each node is an independent computer, you can also use other programming models. One example is the "manager-worker" model, in which one "manager" program starts up several "worker" programs on other nodes, then gathers and interprets their results.

## Cross-Development Facility

Paragon OSF/1 comes with a complete program development environment, including compilers, linker, libraries, and related tools. You can perform all phases of program development on the Intel supercomputer. In addition, the compilers, linker, and libraries for Paragon OSF/1 are also available on selected UNIX workstations. This *cross-development facility* lets you edit, compile, and link Paragon OSF/1 programs on your own workstation.

Note, though, that the cross-development facility does not include a way to run a Paragon OSF/1 executable that resides on your workstation's disk. You must transfer your executable files to the Intel supercomputer for execution and debugging. You can do this by mounting your workstation's file system onto the Intel supercomputer, or the Intel supercomputer's file system onto your workstation, using the Network File System (NFS). You can also use commands such as **rcp** or **ftp** to copy the executable files to the Intel supercomputer. To execute files on the Intel supercomputer once they are transferred, you can use the standard **rsh** or **rcmd** command.

# Using Paragon™ OSF/1 Commands

2

## Introduction

This chapter tells you how to use Paragon OSF/1 commands to perform the following tasks:

- Compiling and linking applications.
- Running applications.
- Managing running applications.
- Managing partitions.

The commands discussed in this chapter are available to all users. See the *Paragon™ System Administrator's Guide* for information on commands that require root privilege.

This chapter does *not* discuss NQS, the Network Queuing System, which is used at some sites to schedule application execution. See the *Paragon™ Network Queuing System Manual* for information on NQS.

## Terminology

This chapter uses the following terms:

- A *parallel application*, usually just called an *application* in this manual, is a group of cooperating processes that runs on the nodes of the Intel supercomputer.
- A *program* is a file (source or executable). An application consists of one or more programs running on one or more nodes. The term *program* is also used to refer to a non-parallel program (an ordinary program that runs on one node).

- A *partition* is a named group of nodes. When you run a parallel application, you must select a partition to run it in (if you don't, it runs in your *default partition*). The partition places limits on some of the execution characteristics of the application, such as how many nodes it can use and how long it can use them before it is "rolled out" and another application is "rolled in." You can allocate all of the nodes of the partition to the application, or just some of them. This allocation may or may not be exclusive, depending on the characteristics of the partition.

All Intel supercomputers have two special partitions called the *service partition* and the *compute partition*. The service partition is used to run non-parallel programs such as shells and editors, and the compute partition is used to run parallel applications. The other partitions on your system, and what you can do with them, are determined by your system administrator.

## Using Paragon™ OSF/1 Commands on the Intel Supercomputer

The Paragon OSF/1 operating system provides all of the standard commands of OSF/1, such as **cat** and **ls**, which work as specified by the Open Software Foundation. These commands are not described in this chapter; see the *OSF/1 Command Reference* for information on these commands.

Paragon OSF/1 also provides several commands that are not specified by the Open Software Foundation, such as **mkpart** and **rmpart**. These commands are described in this chapter, and manual pages for these commands are provided in the *Paragon™ Commands Reference Manual*.

To use any of these commands, you must first log into an Intel supercomputer. Intel supercomputers have no directly-attached terminals; you must first log into another system (typically a workstation running some variant of the UNIX operating system) and then log into the Intel supercomputer over the network, using a command such as **rlogin** or **telnet**. Once you have logged in, you use these commands in the same way as commands on any other computer running OSF/1.

## Using Paragon™ OSF/1 Commands on Workstations

The Paragon OSF/1 operating system also comes with several commands that run on workstations (for example, the **icc** and **if77** cross-compilers). These commands are described briefly in this chapter; complete descriptions and manual pages for these commands are provided in the *Paragon™ C Compiler User's Guide* and *Paragon™ Fortran Compiler User's Guide*.

To use these commands, you must first log into a workstation on which these commands are supported, then configure your account as described under "Configuring Your Environment for Cross-Development" on page 2-6. Once you have done this, you can use the Paragon OSF/1 cross-development commands in the same way as other commands on the workstation. However, if you compile an application on a workstation you must transfer the executable file to an Intel supercomputer to execute it. Depending on your local configuration, you may be able to use the Network File System (NFS), the **rcp** command, the **ftp** command, or some other technique to do this. Ask your system administrator about how files are shared between the Intel supercomputer and other systems on your network.

## A Quick Example

Here is a quick example that shows you how to compile, link, and execute a simple application on an Intel supercomputer.

### Information You Need

Before you begin, you will need the following information:

- The network name of your Intel supercomputer.
- The command to use to log into the Intel supercomputer, such as **rlogin** or **telnet**.
- Your user name and password on the Intel supercomputer (if necessary).
- The name of the *default partition* you should use to run parallel applications.

This information should be available from your system administrator.

### Compiling, Linking, and Executing an Application

Once you have the necessary information, the procedure to compile, link, and execute an application is as follows:

1. Log into the Intel supercomputer, as instructed by your system administrator.
2. Set the environment variable *NX\_DFLT\_PART* to the name of your default partition:

- If you use the C shell, use the following command:

```
% setenv NX_DFLT_PART partition_name
```

- If you use the Bourne or Korn shell, use the following commands:

```
$ NX_DFLT_PART=partition_name  
$ export NX_DFLT_PART
```

3. Type in a short program:

- If you are a Fortran programmer, type the following program into the file *myapp.f*:

```
program hello
include 'fnx.h'

write(*,100) mynode()
100 format('Hello from node', i4, '!')

end
```

- If you are a C programmer, type the following program into the file *myapp.c*:

```
#include <nx.h>

main()
{
    printf("Hello from node %d!\n", mynode());
}
```

4. Compile the program into an executable file:

- If you are a Fortran programmer, use the following command:

```
% f77 -nx -o myapp myapp.f
```

- If you are a C programmer, use the following command:

```
% cc -nx -o myapp myapp.c
```

5. Execute the resulting file, *myapp*, on four nodes with the following command:

```
% myapp -sz 4
Hello from node 0!
Hello from node 3!
Hello from node 1!
Hello from node 2!
```

The order in which the output lines appear may vary.

That's all there is to it! Of course, Paragon OSF/1 provides many additional commands and switches you can use to control the behavior of the compiler and the resulting application. These commands and switches are described in the rest of this chapter.

## Compiling and Linking Applications

Command Synopsis	Description
<code>cc -nx [ switches ] sourcefile...</code>	Compile a Paragon OSF/1 application written in C on an Intel supercomputer.
<code>f77 -nx [ switches ] sourcefile...</code>	Compile a Paragon OSF/1 application written in Fortran on an Intel supercomputer.
<code>icc -nx [ switches ] sourcefile...</code>	Compile a Paragon OSF/1 application written in C on an Intel supercomputer or cross-development workstation.
<code>if77 -nx [ switches ] sourcefile...</code>	Compile a Paragon OSF/1 application written in Fortran on an Intel supercomputer or cross-development workstation.

You can compile and link applications on the Intel supercomputer itself, or on a workstation that supports the Paragon OSF/1 cross-development environment. On the Intel supercomputer, you can use the “native” commands `cc` and `f77` or the “cross-development” commands `icc` and `if77`. On a workstation, you must use the cross-development commands `icc` and `if77`. The native and cross-development versions of each command take the same switches and work identically.

When compiling and linking an application, you should generally use the switch `-nx` on the command line. The `-nx` switch has three effects:

- If used while linking a C or Fortran program, it links in `libnx.a`, the library that contains all the system calls described in this manual.
- If used while linking a C or Fortran program, it links in a special start-up routine that starts up the program on multiple nodes, as specified by standard command line switches and environment variables.
- If used while compiling a C program, it defines the preprocessor symbol `__NODE`. The program being compiled can use preprocessor statements such as `#ifdef` to control compilation based on whether or not this symbol is defined. (This preprocessor symbol is *not* defined if `-nx` is used while compiling a Fortran program.)

For example, the following command line compiles and links the file `myapp.c` to create an executable file called `myapp` (on the Intel supercomputer):

```
% cc -nx -o myapp myapp.c
```

The following command line has the same effect (on the Intel supercomputer or a cross-development workstation):

```
% icc -nx -o myapp myapp.c
```

## NOTE

Do not use **-nx** if your application calls **nx\_initve()**.

The Paragon OSF/1 operating system provides **nx\_initve()** and related functions to give your application more control over the way it starts up. They let the application perform actions for itself that are normally performed for it by **-nx**. If you link your application with **-nx** and it also calls **nx\_initve()** itself, the application's call to **nx\_initve()** will fail and return -1. See "Managing Applications" on page 4-2 for more information on **nx\_initve()** and related functions.

To link an application that calls **nx\_initve()**, use the switch **-lnx** instead of **-nx**. The **-lnx** switch links in *libnx.a*, but without the special start-up routine supplied by **-nx**. A program linked with **-lnx** can use all the calls described in this manual, but does not automatically start itself on multiple nodes. (Note that the **-lnx** switch must appear on the compiler command line *after* the filenames of any source or object files that use these calls.) Note that the preprocessor symbol **\_\_NODE** is *not* defined by **-lnx**.

A program that is not linked with **-nx** and does not call **nx\_initve()** is not a parallel application. It does not recognize the command-line switches described under "Running Applications" on page 2-11, and it always runs on one node in the service partition. (If it creates additional processes by calling **fork()**, they may run on the same node or a different node, but they will always run in the service partition.)

## Configuring Your Environment for Cross-Development

Before you can use the **icc** and **if77** commands on your workstation, you must configure your environment as follows:

- The environment variable **PARAGON\_XDEV** must be set to the pathname of the directory that contains the Paragon OSF/1 cross-development facility. If you don't know this pathname, ask your system administrator.
- Your execution search path (**PATH** or **path** variable) must include the directory **\$PARAGON\_XDEV/paragon/bin.arch**, where **arch** identifies the architecture of your workstation (such as **sun4** for a Sun-4 workstation).
- If you want to read Paragon OSF/1 online manual pages on your workstation, your online manual page search path (**MANPATH** variable or equivalent facility) must include the directory **\$PARAGON\_XDEV/paragon/man**.



You should put the definitions of these variables into your *.cshrc* or *.login* file (or the equivalent start-up file for your shell). For example, suppose the Paragon OSF/1 cross-development facility is installed in the directory */usr/local/XDEV*. If you use the C shell, you would add these lines to your *.cshrc* file:

```
setenv PARAGON_XDEV /usr/local/XDEV
set path=( $path $PARAGON_XDEV/paragon/bin.`arch` )
setenv MANPATH "${MANPATH} : ${PARAGON_XDEV}/paragon/man"
```

(The curly braces in " *\${MANPATH} : \${PARAGON\_XDEV}/paragon/man*" are necessary because a colon after a variable name is special to the C shell.)

Once your environment is properly configured, you can use the **icc** or **if77** command to compile and link applications on your workstation. For example, the following command line compiles and links the file *myapp.f* to create an executable file called *myapp*:

```
% if77 -nx -o myapp myapp.f
```

The executable file, *myapp*, can only be executed on the Intel supercomputer. You can do this by putting it in a directory that is shared between your workstation and the Intel supercomputer with the Network File System (NFS), or by copying it to the Intel supercomputer with the **ftp** or **rcp** command. If you use the **ftp** command, the resulting file may not have execute permission; if this happens, use the **chmod** command on the Intel supercomputer to give *myapp* execute permission.

## NOTE

The Paragon OSF/1 versions of the compilers are not the same as their iPSC® system equivalents.

If you develop programs for the iPSC series of supercomputers from Intel Corporation as well as for Paragon OSF/1, you must be sure that your execution search path (*PATH* or *path* variable) is set appropriately for your current target system. To compile a program for Paragon OSF/1, the variable *PARAGON\_XDEV* must be set appropriately and your execution search path must include *\$PARAGON\_XDEV/paragon/bin.arch*; to compile a program for the iPSC system, the variable *IPSC\_XDEV* must be set appropriately and your execution search path must include *\$IPSC\_XDEV/i860/bin.arch* instead. Be sure that your execution search path does not include both these directories at the same time.

## Tips for Compiling and Linking

The following sections give you some tips for compiling and linking Paragon OSF/1 applications (on either the Intel supercomputer or a cross-development workstation).

### Using Other Switches

The `cc`, `f77`, `icc`, and `if77` commands have a variety of switches to control their operation. For a description of these switches and other information on these commands, see the online manual pages for the commands or the following printed manuals:

`cc`, `icc`            *Paragon™ C Compiler User's Guide.*

`f77`, `if77`        *Paragon™ Fortran Compiler User's Guide.*

### Including `nx.h` or `fnx.h`

As a general rule, always include the file `nx.h` in all Paragon OSF/1 C programs. This file contains definitions and declarations needed by the Paragon OSF/1 C system calls. Although a specific application may not need the definitions and declarations contained in `nx.h`, the overhead involved in including it in all programs is minor. Include it in your C programs as follows:

```
#include <nx.h>
```

For Fortran programs, the corresponding file is `fnx.h`. Include it in your Fortran programs as follows:

```
include 'fnx.h'
```

### Specifying Include File and Library Pathnames

The standard include and library directories depend on whether you are using the native development commands or the cross-development commands:

- The native development commands search for include files in the directory `/usr/include`, and they search for libraries in the directories `/usr/ccs/lib` (searched first) and `/usr/lib` (searched second).
- The cross-development commands search for include files in the directory `$PARAGON_XDEV/paragon/include`, and they search for all libraries in the directory `$PARAGON_XDEV/paragon/lib-coff`.

Note, though, that on the Intel supercomputer the directories */usr/paragon/XDEV/paragon/lib-coff* and */usr/ccs/lib* are identical, the directories */usr/paragon/XDEV/paragon/include* and */usr/include* are identical, and the default for *\$PARAGON\_XDEV* is */usr/paragon/XDEV*, so this difference may not be significant.

If you need to include a file that is not in the standard include directory or in the same directory as the source file, you must use the **-I** switch on the compiler command line to identify the nonstandard directory. For example, the following command line compiles and links an application that uses include files in the directory */usr/local/include*:

```
% icc -nx myapp.c -I/usr/local/include
```

If you need to link to a library that is not in one of the standard library directories, then you must modify the command line in one of the following ways:

- Use the **-L** switch to provide the pathname of the directory in which the library is located. For example, the following command line compiles and links an application that depends on the library *libfft.a* located in the directory */usr/local/lib*:

```
% icc -nx -L/usr/local/lib myapp.c -lfft
```

- Specify the complete pathname of the appropriate library or libraries on the command line. For example, the following command line also compiles and links an application that depends on the library *libfft.a* located in the directory */usr/local/lib*:

```
% if77 -nx myapp.c /usr/local/lib/libfft.a
```

## Preprocessing a Fortran Program

If your Fortran program is in a file whose filename ends with an uppercase “.F” (rather than the standard lowercase “.f”), the **if77** command runs a preprocessor (like the standard C preprocessor) on the file. This enables you to use lines like the following in a Fortran program:

```
#include <file.h>
```

```
#define MAX 87
```

## Order of Switches

Most **cc**, **f77**, **icc**, and **if77** switches are not order-sensitive. However, order is important for the **-I**, **-L**, and **-l** switches and for listing libraries when linking. When constructing command lines, keep the following guidelines in mind:

- List include directories (**-I** switch) in the order in which they should be searched. The list of include directories you specify with **-I** switches is collected together and used for all source files you specify. For example, the following command looks for include files in the directory *myincludes*, then the directory *./includes*, and finally the standard include directory when compiling *a.c*, *b.c*, and *c.c*:

```
% icc a.c -Imyincludes b.c -I./includes c.c
```

- List libraries in the order in which they should be searched. The Paragon OSF/1 linkers are single-pass linkers; they cannot resolve a backward library reference (i.e., a reference to a library object that was defined in a library that has already been searched). Note that this means that if you use the **-lnx** switch, you should place it *after* any source files that need it, as follows:

```
% if77 -o myapp myapp.f -lnx
```

Backward references between objects (*.o* files), however, are not a problem, as all listed objects are linked unconditionally.

- The **-L** switch affects only the search path of libraries that are listed *after* the **-L** switch. For example, the following command searches only the standard library directories for the library *libnews.a*, but searches the directory *./mylibs* (as well as the standard library directories) for the library *libgx.a*:

```
% icc -nx myprog.c -lnews -L./mylibs -lgx
```

- If you specify more than one **-L** switch, the named directories are searched in reverse order (the directory specified by the first **-L** switch on the command line is searched after the directory specified by the second **-L** switch on the command line). For example:

```
% icc -nx myprog.c -lnews -L../mylibs -lgx -Llocallibs -llocal
```

This command searches for libraries as follows:

- It searches only the standard library directories for the library *libnews.a*.
- It searches the directory *../mylibs* and then the standard library directories for the library *libgx.a*.
- It searches the directory *locallibs*, then *../mylibs*, and then the standard library directories for the library *liblocal.a*.

Note that the **-L** switch also affects system libraries; in fact, directories specified by **-L** are searched for system libraries *before* the standard library directories.

## Running Applications

Once you have compiled your application into a Paragon OSF/1 executable file (and, if necessary, copied the executable to an Intel supercomputer), you run it by typing its name at your Paragon OSF/1 shell command prompt, as you would for any other compiled program.

For example, if *myapp* is a compiled application, you can execute it with the following command:

```
% myapp
```

The way the application runs depends on how you linked it and on what system calls it makes:

- If *myapp* was linked with the **-nx** switch, this command runs *myapp* on your default number of nodes in your default partition. The section "Controlling the Application's Execution Characteristics" on page 2-13 tells you more about the default partition, and about the environment variables and command-line switches you can use to control the execution characteristics of applications linked with the **-nx** switch.
- If *myapp* was linked with the **-lnx** switch, this command runs *myapp* on the nodes and partition specified by system calls within the application. The section "Managing Applications" on page 4-2 tells you how to use these system calls. If *myapp* does not specify the nodes and partition in these calls, it defaults to running on your default number of nodes in your default partition. If *myapp* does not make any of these calls, it runs on one node in the service partition.
- If *myapp* was linked without the **-nx** or **-lnx** switch, it is an ordinary non-parallel program, and it runs on one node in the service partition.

If you see the error message “request overlaps with nodes in use,” it means that your default partition does not allow overlapping applications and someone else is already running an application in that partition. Try again later, or use a different partition (as described under “Running an Application in a Particular Partition” on page 2-22). You can use the **pspart** command to determine which partitions have applications running in them, as described under “Listing the Applications in a Partition” on page 2-51.

If you see the error message “partition permission denied” or “exceeds partition resources,” check to be sure the environment variables *NX\_DFLT\_PART* and *NX\_DFLT\_SIZE* are properly defined. See “Using the Default Partition” on page 2-14 and “Specifying Application Size” on page 2-15 for more information on these variables; see your system administrator for information on the proper settings for these variables at your site.

If you see the error message “error 216 occurred, unknown,” it means that the application was compiled on a previous release of the Paragon OSF/1 operating system and uses an out-of-date version of the libraries. (Error 216 is “parallel application incompatible with OS release”, but the “unknown” message may appear if the application is so out-of-date that it doesn't know about the existence of this error.) If this occurs, recompile the application and try again.

## I/O Redirection

You can redirect the standard input, standard output, and standard error of an application with the usual OSF/1 techniques. For example, the following command redirects the input and output of the application *myapp*:

```
% myapp < myfile.in > myfile.out
```

This command runs the application *myapp* with its standard input redirected from the file *myfile.in* and its standard output redirected to the file *myfile.out*.

Note that, by default, all the nodes read and write their standard input, standard output, and standard error using PFS I/O mode 0. In mode 0, all file access requests are honored on a first-come, first-served basis. You can change this behavior by selecting a different I/O mode; see “Using I/O Modes” on page 5-13 for more information. The standard input, standard output, and standard error are line-buffered by default. This means that if all the nodes write to standard output or standard error, the output from all the nodes is intermixed in the output, line by line; if all the nodes read from standard input, each line of the input goes to an arbitrary node.

## Controlling the Application's Execution Characteristics

Command Synopsis	Description
<pre>application [ -sz size   -sz hXw   -nd hXw:n ]   [ -pri priority ] [ -pt ptype ]   [ -on nodespec ] [ -pn partition ]   [ mp_switches ]   [ \; app2 [ -pt ptype ] [ -on nodespec ] ] ...</pre>	Execute a Paragon OSF/1 application.

When you run an application, you can use command-line switches and environment variables to control the way the application executes. This section discusses all the switches and environment variables except for the *mp\_switches*, which are used for message-passing performance tuning; for information on the *mp\_switches*, see “Message-Passing Configuration Switches” on page 8-18.

Command-line switches can appear in any order on the command line, and may be intermixed with application-specific switches and arguments. If you specify the same command-line switch more than once in a single command, the last occurrence overrides the earlier ones. For example, the following two commands are equivalent:

```
% myapp -sz 4 -sz 50 -pri 8 file.dat
% myapp -pri 8 -sz 4 file.dat -sz 50
```

Each of these commands runs the application *myapp*, with the argument *file.dat*, at priority 8 on 50 nodes of your default partition.

If the application was linked with the *-nx* switch, the command-line switches discussed in this section are interpreted and removed from the command line before the application starts up. In the previous examples, the arguments *-pri 8*, *-sz 4*, and *-sz 50* are interpreted and removed by the *-nx* code; *myapp* sees only the argument *file.dat* (if *myapp* is a C program *argc* is 2, *argv[0]* is “myapp”, and *argv[1]* is “file.dat”).

### NOTE

All the examples in this section assume that *myapp* was linked with the *-nx* switch.

An application that is not linked with *-nx* controls its own execution with system calls, as discussed under “Managing Applications” on page 4-2. Such an application may or may not obey the command-line switches discussed in this section, depending on how it was programmed.

## Using the Default Partition

When you run a parallel application on the Intel supercomputer, it runs in a *partition*. The partition determines the maximum number of nodes used by the application and how the application is scheduled, as described later in this chapter. An application stays in the same partition for its entire run.

If you do not specify otherwise, the application runs in the partition specified by the environment variable *NX\_DFLT\_PART*. If the environment variable *NX\_DFLT\_PART* is not set, the application runs in the *compute partition*, a special partition that is present on all Intel supercomputers. The partition specified by *NX\_DFLT\_PART* (or, if this variable is not set, the compute partition) is called your *default partition*.

For example, to run the application *myapp* in your default partition, use the following command:

```
% myapp
```

This command runs the application *myapp* in the partition specified by the environment variable *NX\_DFLT\_PART*, or in the compute partition if *NX\_DFLT\_PART* is not set.

If you see an error message such as “partition not found” or “partition permission denied,” ask your system administrator what your default partition should be, then use the commands described in the next section to set the variable *NX\_DFLT\_PART* to that value. You can also use the **-pn** switch (described under “Running an Application in a Particular Partition” on page 2-22) to run an application in a different partition.

For more information about partitions, see “Managing Partitions” on page 2-25.

## Setting Your Default Partition

The command you use to set or change your default partition depends on which shell you use.

- If you use the C shell, use the **setenv** command. For example, if you are a C shell user, the following command sets your default partition to *mypart*:

```
% setenv NX_DFLT_PART mypart
```

**setenv** is a built-in command of the shell; see **cs**h in the *OSF/1 Command Reference* for more information.

You can put this command in your *.login* or *.cshrc* file on the Intel supercomputer to have your default partition set to *mypart* each time you log in.



- If you use the Bourne or Korn shell, set the variable and use the **export** command to make its value available to commands other than the shell. For example, if you are a Bourne or Korn shell user, the following commands set your default partition to *mypart*:

```
$ NX_DFLT_PART=mypart
$ export NX_DFLT_PART
```

You do not have to use the **export** command each time you set the variable. You only have to export a variable once in each login session. **export** is a built-in command of the shell; see **sh** or **ksh** in the *OSF/1 Command Reference* for more information.

You can put these commands in your *.profile* file on the Intel supercomputer to have your default partition set to *mypart* each time you log in.

You can use an absolute or relative partition pathname as the value of *NX\_DFLT\_PART*. For example, the following C shell commands are equivalent:

```
% setenv NX_DFLT_PART myorg.mypart
% setenv NX_DFLT_PART .compute.myorg.mypart
```

See “Partition Pathnames” on page 2-28 for more information on partition pathnames.

If you use the C or Korn shell, you can create an alias to change your default partition. For example, the following C shell command creates a “setpart” alias that sets your default partition to its argument:

```
% alias setpart 'setenv NX_DFLT_PART \!*
```

### Determining the Current Default Partition

To find out your default partition once you have set it, use the **echo** command. For example:

```
% echo $NX_DFLT_PART
mypart
```

This command works the same in any shell.

### Specifying Application Size

An application's *size* is the number of nodes allocated to the application from the partition. The processes of the application run only on this set of nodes, and do not exchange messages with processes on nodes outside this set. Depending on the characteristics of the partition, this allocation may or may not be exclusive: some or all of these nodes may also be allocated to other applications and/or other partitions. An application keeps the same size for its entire run.

To set an application's size, use the switch **-sz size**, where *size* is any positive integer less than or equal to the number of nodes in the partition. For example, to run the application *myapp* on 64 nodes of your default partition, use the following command:

```
% myapp -sz 64
```

The **-sz size** switch attempts to allocate a square group of nodes if it can. If this is not possible, it attempts to allocate a rectangular group of nodes that is either twice as wide as it is high or twice as high as it is wide. If this is not possible, it allocates any available nodes; in this case, nodes allocated to the application may not be *contiguous* (that is, they may not all be physically next to each other). If the requested number of nodes is not available, the command fails and the application does not run; an error message is printed to explain why the specified number of nodes is not available.

No matter what the shape of the application, node numbers within the application (as returned by **mynode()**) will always be sequential from 0.

### Specifying a Rectangle of Nodes

To force allocation of a contiguous rectangle of a particular size and shape, use the switch **-sz hXw**, where *h* and *w* are positive integers that specify the height and width of the desired rectangle. (You can use an uppercase or lowercase letter **X** between the integers *h* and *w*.) For example, to run *myapp* on an 8 by 8 node rectangle of your default partition, use the following command:

```
% myapp -sz 8x8
```

If successful, this command runs *myapp* on an 8 by 8 node rectangle of nodes, which could be located anywhere within the partition that it fits. If no 8 by 8 node rectangle is available in the default partition, the command fails immediately and the application does not run, even if there are 64 nodes free in the partition. If this occurs, the command fails with the error message "exceeds partition resources" if no such rectangle can be found that fits within the partition, or "request overlaps with nodes in use" if the rectangle fits within the partition but some of its nodes are busy).

### Specifying a Particular Rectangle of Nodes

To force allocation of a contiguous rectangle of a particular size and shape at a particular location within the partition, use the switch **-nd hXw:n**. (This switch is called **-nd**, rather than **-sz**, because it specifies a particular set of nodes rather than just a size or shape.)

In the **-nd hXw:n** switch, *h* and *w* are positive integers that specify the height and width of the desired rectangle, and *n* is a positive integer that specifies the node number within the partition for the upper left corner of that rectangle. You can use an uppercase or lowercase letter **X** between the integers *h* and *w*. When choosing the value of *n*, remember that in an *m*-node partition the nodes are numbered left to right and top to bottom from 0 to *m*-1.

For example, to run *myapp* on an 8 by 8 node rectangle in the upper left corner of your default partition, use the following command:

```
% myapp -nd 8x8:0
```

In this case, if the specified nodes are not available in the default partition, the application fails immediately (even if there is a different 8 by 8 node rectangle available).

### Using the Default Size

If you don't use the **-sz** or **-nd** switch, the application's size is specified by the environment variable *NX\_DFLT\_SIZE*, whose value must be a single positive integer. You can use the techniques discussed for the *NX\_DFLT\_PART* variable in the previous section to get and set the value of the *NX\_DFLT\_SIZE* variable. If *NX\_DFLT\_SIZE* is not set, the application runs on all nodes of the partition, and its size is set to the size of the partition. The size specified by *NX\_DFLT\_SIZE* (or, if this variable is not set, the size of the partition) is called your *default number of nodes*.

An application can determine its size by calling **numnodes()**, and each process in the application can determine its node number within the application by calling **mynode()**. **mynode()** returns a node number from 0 to one less than the application's size. (See "Process Characteristics" on page 3-3 for more information on these calls.) For example, with **-sz 64**, **-sz 8x8**, or **-nd 8x8:0**, **numnodes()** returns 64 and **mynode()** returns a number from 0 to 63 inclusive. There is no way for an application to change its size.

An application can determine its shape by calling **nx\_app\_rect()**, which returns the height and width of the rectangle of nodes allocated to the application. If the nodes allocated to the application do not form a rectangle, **nx\_app\_rect()** returns a height of 1 and a width equal to **numnodes()**. (**nx\_app\_rect()** can also be called by the name **mypart()** for compatibility with the Touchstone DELTA System.)

### Specifying Application Priority

An application's *priority* is an integer associated with the application that is used in determining how much of a node's processor time the application gets when the node is allocated to more than one application at once. 0 is the lowest priority, and 10 is the highest.

The application's priority is only one of several factors that determine how much processor time it gets. For example, the application's processor time can be affected by the priorities of other applications in the system and by the *effective priority limit* of the partition in which the application runs. See "Scheduling Characteristics" on page 2-33 for more information.

To set the priority of the application, use the switch **-pri priority**, where *priority* is an integer from 0 to 10 inclusive. If you don't use the **-pri** switch, the application's priority is set to 5.

For example, to run the application *myapp* with a priority of 6, use the following command:

```
% myapp -pri 6
```

An application can change its priority by calling `nx_pri()` (see “Setting an Application’s Priority with `nx_pri()`” on page 4-9 for more information).

## Specifying Process Type

A process’s *process type*, or *ptype*, is an integer associated with the process that differentiates it from any other process in the application that is on the same node. The process’s node number and process type together form the process’s “address” for messages within the application.

To set the process type of each process in the application, use the switch `-pt ptype`, where *ptype* is an integer from 0 to 2,147,483,647 ( $2^{31} - 1$ ) inclusive. If you don’t use the `-pt` switch, the process type of each process is 0.

For example, to run the application *myapp* with a process type of 1 for each process, use the following command:

```
% myapp -pt 1
```

A process can find out its current process type by calling `myptype()`. For example, with `-pt 1`, `myptype()` returns 1 on all nodes. Once a process’s process type has been set to a valid value, it cannot change its process type and no other process in the same application on the same node can use that process type for the run of the application. See “Process Characteristics” on page 3-3 for information on process types and the `myptype()` and `setptype()` system calls.

The `-pt` switch is most commonly used when running multiple programs in one application, as discussed under “Running Applications Consisting of Multiple Programs” on page 2-21. In most other circumstances, you can use the default process type of 0.

## Running a Program on a Subset of the Nodes

Usually you run the same program file on all the nodes allocated to the application from the partition. However, you can also run a program on just some of the nodes, leaving the other nodes vacant for other programs. When you do this, the other nodes are allocated to the application, but no processes are started on them.

To run a program on a subset of the nodes of an application, use the switch **-on** *nodespec*, where *nodespec* is one of the following:

<i>x</i>	The node whose node number is <i>x</i> .
<i>x..y</i>	The range of nodes from numbers <i>x</i> to <i>y</i> .
<b>n</b>	The last node of the partition.
<i>nspec[.nspec]...</i>	The specified list of nodes, where each <i>nspec</i> is a node specifier of the form <i>x</i> , <i>x..y</i> , or <b>n</b> (no node may appear more than once in this list). Do not put any spaces in this list.

If you don't use the **-on** switch, the program is run on all nodes allocated to the application.

## NOTE

The numbers you use with **-on** are node numbers within the application (which always range from 0 to one less than the size of the application), not node numbers within the partition.

For example, to run the program *myapp* on the first three nodes of a 20-node application, use the following command:

```
% myapp -sz 20 -on 0,1,2
```

This command creates an application of size 20 in your default partition and runs *myapp* on nodes 0, 1, and 2 of the application. Within this application, the function **numnodes()** returns 20, and the function **mynode()** returns a number from 0 to 19 inclusive. However, no processes are started on nodes 3 through 19.

You can use the letter **n** to represent "the last node in the application." For example, the following command creates an application of your default size in your default partition and runs *myapp* on the first and last nodes of the application:

```
% myapp -on 0,n
```

For example, if your **NX\_DFLT\_SIZE** variable is set to 64 (and there are at least 64 nodes in your default partition), this would run *myapp* on nodes 0 and 63 of the application.

You can also use a pair of numbers separated by two periods ( $x..y$ ) to specify “nodes  $x$  through  $y$  inclusive.” For example, the following command creates an application of size 100 in your default partition and runs the program *myapp* on nodes 10 through 90:

```
% myapp -sz 100 -on 10..90
```

It doesn't matter whether  $y$  is greater than  $x$  or vice versa. For example, the following command *also* creates an application of size 100 in your default partition and runs the program *myapp* on nodes 10 through 90:

```
% myapp -sz 100 -on 90..10
```

These notations can be combined. For example, the following command creates an application of your default size in your default partition and runs *myapp* on all nodes but node 0 of the application:

```
% myapp -on 1..n
```

Another example: the following command creates an application of your default size in your default partition and runs *myapp* on node 1, node 3, nodes 5 through 10 inclusive, and the last node of the application:

```
% myapp -on 1,3,5..10,n
```

## NOTE

Do not use **-on** if you just want to run a single program on a specific number of nodes.

The **-on** switch is designed to be used when running multiple programs as a single application, as discussed in the next section. You can also use the **-on** switch to run a “manager” program on one or a few nodes of an application; the “manager” program can then run “worker” programs on other nodes by calling **nx\_nfork()**, **nx\_load()**, or **nx\_loadve()** (see “Managing Applications” on page 4-2 for information on these functions).

The **-on** switch is *not* designed to run an application on a particular number of nodes or a particular set of nodes. If you want to run an application on a particular number of nodes, use the **-sz** switch. If you want to run an application on a particular set of nodes, allocate a partition containing those nodes and run the application on all nodes of that partition (see “Managing Partitions” on page 2-25 for information on partitions).

If you use **-on** when you should be using **-sz**, the application will be allocated more nodes than it needs. Also, if you use **-on** and do not run a program on every node of the application, global operations will hang. (The global operations described under “Global Operations” on page 3-27, such as **gdsum()**, block until they are called by every node in the application. If you run a program on only a subset of the nodes, these operations will block forever.)

## Running Applications Consisting of Multiple Programs

You can run multiple program files as a single application. For example, you could run two or more separate programs on every node (the resulting processes must have different process types, and the processes time-share the processor while the application is active). You might also run a manager program on one node and worker programs on the other nodes. The programs should be written to work together; you would not usually run two arbitrary programs together in one application.

To run multiple program files as a single application, use the following syntax:

```
% file [ switches ] [ \; file [ -pt ptype ] [ -on nodespec ] ] ...
```

That is, you use two or more complete commands on one line, separated by an escaped semicolon (backslash followed by semicolon).

### NOTE

The escaped semicolon (\;) must be preceded and followed by a space or tab. Otherwise, it will be considered part of the preceding or following argument.

The first *file* must either have been linked with **-nx** or must call **nx\_initve()** without overriding the command line; the second and subsequent *files* may have been linked with or without **-nx**, but must *not* call **nx\_initve()**.

The command-line switches you can use with the *files* are different:

- You can use any application switches (**-sz**, **-pri**, **-pt**, **-on**, **-pn**, and *mp\_switches*) with the first file. The effect of these switches varies according to the switch:
  - The **-sz**, **-pri**, **-pn**, and *mp\_switches* switches you use with the first file affect the entire application.
  - The **-pt** and **-on** switches you use with the first file affect the first file only.
- You can use only the **-pt** and **-on** switches with the second and subsequent files. These switches affect the associated file only.

If you run multiple processes on a single node, you must use the **-pt** switch to specify a unique process type for each process. When two or more processes in an application run on the same node, each must have a different process type. If you don't use the **-pt** switch, each process will have process type 0, and you will receive an error message.

For example, to run the programs *myapp* and *myapp2* as a single application, use the following command:

```
% myapp \; myapp2 -pt 1
```

This command runs the program *myapp* with process type 0 and the program *myapp2* with process type 1 on your default number of nodes in your default partition.

To run the program *manager* on node 0 of a 20-node application and the program *worker* on the remaining nodes, use the following command:

```
% manager -sz 20 -on 0 \; worker -on 1..n
```

This command creates an application of size 20 in your default partition. It then runs the program *manager* on node 0 of the application and the program *worker* on nodes 1 through 19 of the application. All the resulting processes have process type 0, but this does not create a conflict because *manager* and *worker* run on different nodes.

## NOTE

If you forget the backslash before the semicolon, the first program is run as an application by itself and the second program runs after the first program finishes. This usually results in unexpected behavior from the programs.

## Running an Application in a Particular Partition

To run an application in a partition other than your default partition, use the switch **-pn** *partition*. You must have execute permission for the specified partition. The partition specified by **-pn** overrides the value of *NX\_DFLT\_PART*, if any. If you don't use the **-pn** switch, the application runs in your default partition, as described under "Using the Default Partition" on page 2-14.

## NOTE

If your default number of nodes, as specified by the environment variable *NX\_DFLT\_SIZE*, is greater than the number of nodes available in the specified partition, you may get a "partition resources exceeded" or "request overlaps with nodes in use" error.

If you see this error, use the **-sz** switch or change the value of *NX\_DFLT\_SIZE* to specify an application size less than or equal to the size of the specified partition.



For example, to run the application *myapp* on your default number of nodes in the partition *mypart*, use the following command:

```
% myapp -pn mypart
```

You can use an absolute or relative partition pathname with **-pn** (see “Partition Pathnames” on page 2-28 for information on partition pathnames). For example, the following commands are equivalent:

```
% myapp -pn myorg.mypart  
% myapp -pn .compute.myorg.mypart
```

For more information about partitions, see “Managing Partitions” on page 2-25.

## Managing Running Applications

You use the standard OSF/1 techniques to manage running applications. For example, you use your interrupt key (usually **<Del>** or **<Ctrl-c>**) to interrupt a running application. If you use the C shell or Korn shell, you can use your suspend key (usually **<Ctrl-z>**) to suspend an application, and the **fg** or **bg** command to resume it. See **cs**, **sh**, or **ksh** in the *OSF/1 Command Reference* for more information on these techniques.

### NOTE

Interrupting or suspending an application that is “rolled out” will not take effect until the application is “rolled in” again.

Parallel applications can be *gang-scheduled* to make more efficient use of system resources. In gang scheduling, an application is allowed to run for a time period, called the *rollin quantum*, and then is “rolled out” and another application is “rolled in” in its place. If the rollin quantum is long, you may not see any response to a **<Ctrl-c>** or **<Ctrl-z>** for a long time. See “Scheduling Characteristics” on page 2-33 for more information on gang scheduling.

You can also use the **ps** command to determine the status of an application, and the **kill** command to terminate it. For example:

```
% myapp &
[1] 7045
% ps
  PID TT  STAT      TIME COMMAND
 5841 p3  S  +    0:02.50 -csh (csh)
 7045 p3  R           0:00.30 myapp
% kill 7045
% ps
  PID TT  STAT      TIME COMMAND
 5841 p3  S  +    0:02.55 -csh (csh)
[1] + Terminated      myapp
%
```

The **ps** command shows only processes running in the service partition. See **ps** and **kill** in the *OSF/1 Command Reference* for more information on these commands. To show processes running in partitions other than the service partition, use the **pspart** command.

The *myapp* process that you see in the output of **ps** is a special process called the *controlling process* that runs in the service partition; you do not see the other application processes in the output of **ps**. However, sending a signal to the controlling process with **<Del>**, **<Ctrl-c>**, **<Ctrl-z>**, or **kill** signals all the processes in the application. See “Managing Applications” on page 4-2 for more information on the controlling process.

If the application was started from the Bourne shell (**sh**) or from a shell script, you will see *two* processes with the name of the application in the output of **ps**. One of these two processes is the controlling process; the other is another special process, called the *shepherd process*. The shepherd process is necessary for the application; do not kill it. When the application terminates, this process will terminate as well.

To determine which process is which, use the command **ps -f** and examine the **PPID** (parent PID) fields of the two processes. The shepherd process is the parent of the controlling process. For example:

```
$ ps -f
USER      PID    PPID  %CPU  STARTED  TT      TIME COMMAND
chris    131125 131124  0.0   13:55:51 p2      0:00.28 -sh (sh)
chris    131129 131125  0.0   13:56:36 p2      0:00.05 myapp
chris    131130 131129  0.0   13:56:36 p2      0:00.03 myapp
```

In this case the second *myapp* process (PID 131130) is the controlling process. The first *myapp* process, PID 131129, is the parent of the controlling process and is therefore the shepherd process.

You can use the **pspart** command to determine the status of all the applications in a particular partition. See “Listing the Applications in a Partition” on page 2-51 for information on this command.

You can also use the Interactive Parallel Debugger (**ipd**) to control the execution of an application, down to the machine instruction. See the *Paragon™ Interactive Parallel Debugger Reference Manual* for information on **ipd**.

## Managing Partitions

The nodes of the Intel supercomputer are divided into overlapping groups called *partitions*. When you run a parallel application, you must select a partition to run it in. The partition places limits on the execution characteristics of the application, such as which nodes it can use, whether or not it can use nodes that are already in use, and how long it can use them before it is “rolled out” and another application is “rolled in.”

Depending on the policies of your site, you may or may not have to know any more about partitions than what has been discussed in this chapter so far.

- At some sites, the system administrator configures all the partitions; ordinary users can simply set the `NX_DFLT_PART` variable to an appropriate value (or leave it unset and use the compute partition) and then forget all about partitions. If your site is like this, you do not have to read this section. However, you may wish to read it to help you understand how the system works.
- At other sites, users create and configure their own partitions. If your site is like this, you should read this section.

This section includes the following information about partitions:

- Some special partitions that every Intel supercomputer has.
- Specifying partitions with partition pathnames.
- The characteristics of a partition.
- Making partitions with the **mkpart** command.
- Removing partitions with the **rmpart** command.
- Showing the characteristics of a partition with the **showpart** command.
- Listing the subpartitions of a partition with the **lspart** command.
- Listing the applications in a partition with the **pspart** command.
- Changing the characteristics of a partition with the **chpart** command.

## Special Partitions

Every Intel supercomputer has three special partitions:

- The *root partition* directly or indirectly contains all the other partitions in the system. It is the only partition that does not have a parent partition.
- The *service partition* is the partition in which the users' shells and other commands run. Its parent is the root partition.
- The *compute partition* is the partition in which parallel applications run. Its parent is also the root partition.

The characteristics of these partitions are determined by the system administrator. In particular, the system administrator sets the ownership and permissions of these partitions according to local policies. These ownerships and permissions determine whether or not ordinary users can create partitions for their own use, or whether they must run applications in partitions provided for them by the system administrator. If ordinary users are allowed to create partitions, the system administrator can also place restrictions on the characteristics of partitions they create and the use of certain application switches within partitions.

Typically, the service partition and compute partition are the only two children of the root partition and do not overlap. However, the system administrator can choose to configure these partitions differently, and may also create additional child partitions of the root partition.

For example, some systems have an *I/O partition*: a third child of the root partition, which does not overlap with either the service or compute partitions, and which contains the nodes that control disks and other I/O devices. In other systems, the I/O "partition" is not a true partition, but a set of nodes in the root partition that are not part of either the service or the compute partition.


## The Root Partition

The *root partition* is the basis for all other partitions. The name of the root partition is `.` (dot).

The root partition contains every usable node in the system. Depending on the underlying hardware, there may be unusable nodes within the root partition as well. The root partition organizes all the nodes in the system into a two-dimensional grid, or mesh. For example, Figure 2-1 shows the root partition of a 32-node system that is configured as a 4 by 8 node mesh. The nodes are numbered from 0 to 31.

### NOTE

The root partition is always rectangular. (This is *not* true of partitions other than the root partition.)



**Figure 2-1. The Root Partition of a 32-Node System**

For example, a system with 31 nodes would also be a 4-by-8-node rectangle, numbered as shown in Figure 2-1, but one of the nodes would be an *unusable node*, as described under “Unusable Nodes” on page 2-31. You would not be able to start any processes or allocate any subpartitions using this node.

## The Service Partition

The *service partition* is the partition in which the users’ shells, OSF/1 commands, and other non-parallel programs run. The name of the service partition is *service*. The service partition may not contain any subpartitions.

When you log into the Intel supercomputer, a shell is started for you on a node in the service partition; when you execute a command in this shell, the command runs on a node in the service partition. Note that the node the command runs on is not necessarily the same node that the shell runs on; the system starts each new process on the node that is currently the least busy.

## The Compute Partition

The *compute partition* is the partition in which parallel applications run. The name of the compute partition is *compute*.

When you execute a parallel application, one process (called the *controlling process*) runs in the service partition; the other processes of the application run in the compute partition, or in a subpartition of the compute partition. You can specify which partition an application runs in when you execute it.

Your system administrator determines whether or not you can create subpartitions in the compute partition and whether or not you can execute applications in the compute partition itself. There may also be other local policies that affect how you use the compute partition; for example, you may be required to run your applications in certain subpartitions during the day and others at night.

## Partition Pathnames

Since partitions have a hierarchical structure like directories, they also have *pathnames* like directories. Like a file or directory pathname, a *partition pathname* identifies a partition within the hierarchical partition structure by describing the path from a known location to the specified partition.

Unlike file and directory pathnames, however, partition pathnames use a dot (.) instead of a slash (/) to separate the elements of the pathname. This is why the name of the root partition is . (dot). There is also no special partition pathname for “current partition” or “parent of the current partition.” Also, you cannot use wildcards (\* and ?) in partition pathnames.

There are two types of partition pathnames:

- An *absolute partition pathname* specifies the path from the root partition to the specified partition. An absolute partition pathname begins with a dot (.)
- A *relative partition pathname* specifies the path from the compute partition to the specified partition. A relative partition pathname does not begin with a dot.

### NOTE

Relative partition pathnames are always relative to the compute partition (there is no “current partition”).

The absolute partition pathnames of the root partition, service partition, and compute partition are . (dot), *.service*, and *.compute* respectively. Because these partitions are not subpartitions of the compute partition, they do not have relative partition pathnames.

If the partition *mypart* is a subpartition of the compute partition, its absolute partition pathname is *.compute.mypart* and its relative partition pathname is just *mypart*.

If *subpart* is a subpartition of *mypart*, its absolute partition pathname is *.compute.mypart.subpart* and its relative partition pathname is *mypart.subpart*.

## Partition Characteristics

Each partition has the following characteristics:

- A *parent partition* that contains it.
- A *name* that identifies it.
- A set of *nodes* that is allocated to it.
- An *owner* and *group* and a set of *protection modes*, like those of a file or directory, that determine what actions a given user is allowed to perform on it.
- A set of *scheduling characteristics* that determine how applications are scheduled in it.

A partition's characteristics are set when the partition is created. The **mkpart** command, described under "Making Partitions" on page 2-39, lets you specify most of these characteristics on the command line; if you don't specify otherwise, the characteristics of a new partition are set to the same values as those of its parent partition.

You can use the **showpart** command, described under "Showing Partition Characteristics" on page 2-46, to determine a partition's current characteristics.

A partition's parent partition and nodes cannot be changed. You can change the other characteristics with the **chpart** command, described under "Changing Partition Characteristics" on page 2-54.

## Parent Partition

Each partition is contained within another partition. The containing partition is called the *parent* partition, and the contained partition is called a *child* partition or *subpartition* of the parent partition. (There is one exception to this rule: the root partition has no parent.)

You specify a partition's parent when you create it with **mkpart**. The parent partition determines the set of nodes that are available to be allocated to the new partition (a partition cannot include any nodes other than the nodes of its parent). The parent partition also determines the default characteristics of the new partition, as mentioned earlier. A partition's parent does not change for the life of the partition.

## Partition Name

Each partition is identified by a *name*. A partition's name must be unique among all the partitions with the same parent. Partition names can be any length, but must consist of only uppercase letters (A-Z), lowercase letters (a-z), digits (0-9), and underscores (\_).

You specify a partition's name when you create it with **mkpart**, and you can use **chpart** to change an existing partition's name (you must have write permission on the partition's parent partition).

## Nodes Allocated to the Partition

Each partition has a set of *nodes* allocated to it from its parent partition. Depending on the characteristics of the parent partition, this allocation may or may not be exclusive: some or all of these nodes may also be allocated to other partitions and/or applications. The number of nodes in this set is called the partition's *size*.

You can specify the set of nodes allocated to the partition when you create it with **mkpart**. You can specify the partition's size and let the operating system select the nodes, or you can specify certain node numbers from the parent partition. If you don't specify either, the new partition consists of all the nodes of the parent partition.

The set of nodes allocated to a partition does not change for the life of the partition (that is, partitions never move or change their size or shape). Depending on how you allocate the nodes, they may or may not be *contiguous* (all adjacent to each other). Figure 2-2 shows examples of contiguous and noncontiguous partitions.

## Node Numbers Within a Partition

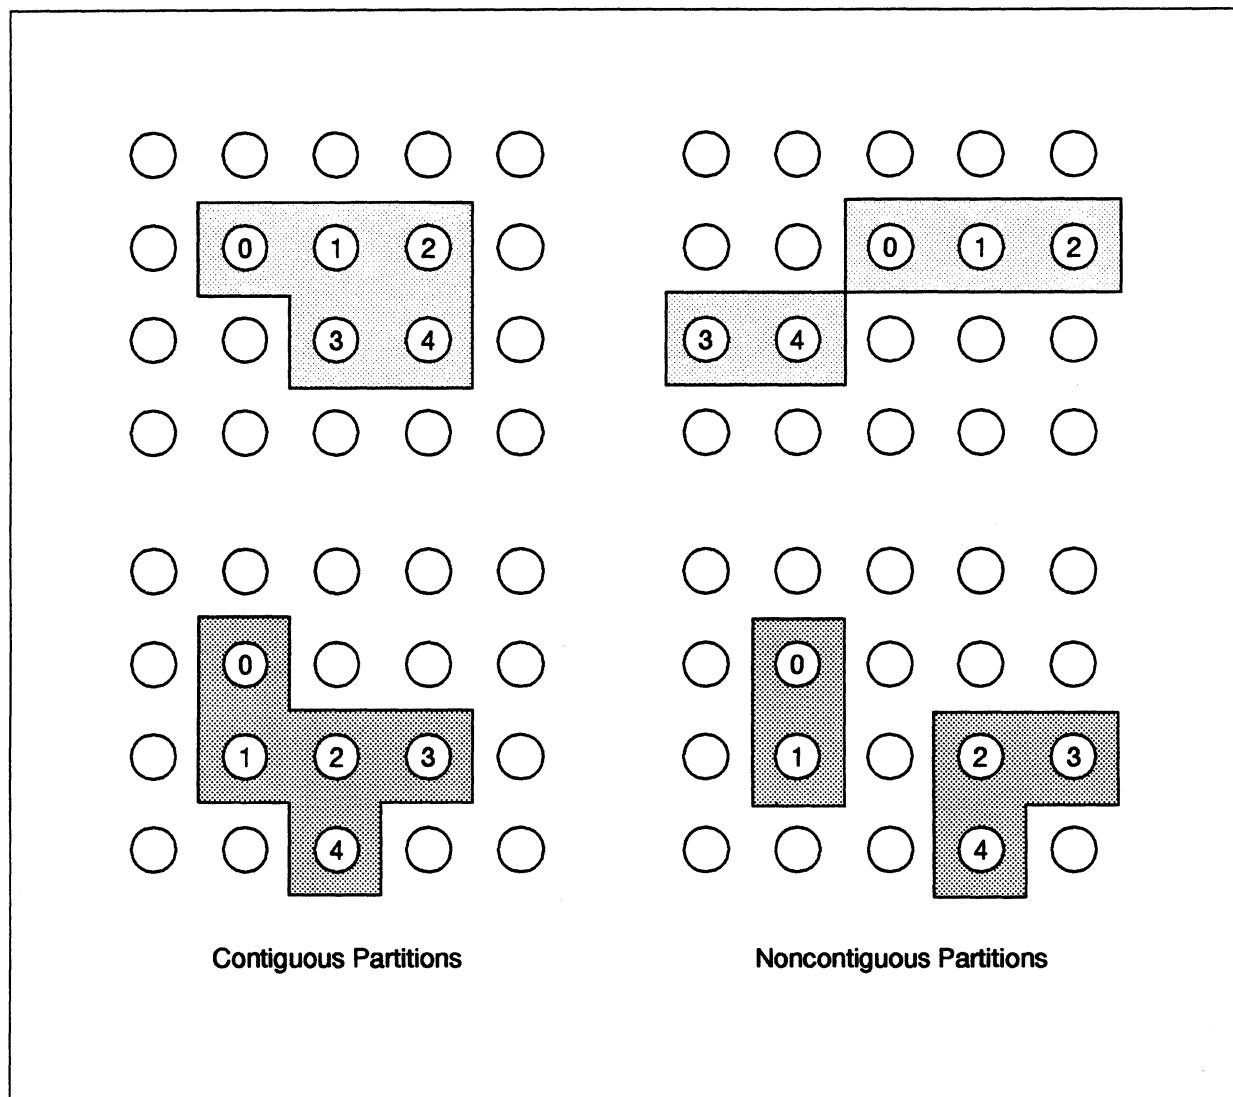
Each node in a partition has a *node number* within the partition: an integer from 0 to one less than the partition's size. The nodes in a partition are typically numbered from left to right and then from top to bottom, as shown in Figure 2-2.

## NOTE

Because partitions can overlap, a single physical node can have many logical node numbers.

For example, Figure 2-3 shows two partitions, called Partition A and Partition B, that have the same parent partition. Partition A consists of nodes 1 through 4 of the parent partition, and Partition B consists of nodes 4 through 8 of the parent partition. In this case, node 4 of the parent partition is also known as node 3 of Partition A and node 0 of Partition B.





**Figure 2-2. Node Numbers in Contiguous and Noncontiguous Partitions**

**Unusable Nodes**

Occasionally a node may become *unusable* because of a hardware or software failure. If this occurs, the node is still allocated to any partitions to which it was allocated before it became unusable, but no applications can be run on that node and no new partitions can include that node until the node becomes usable again. The **showpart** and **lspart** commands indicate if there are any unusable nodes in a partition.

For example, suppose you make a partition containing 20 nodes and later one of those nodes becomes unusable. If you attempt to run an application or make a subpartition with all 20 nodes of this partition while the node is unusable, the attempt will fail.

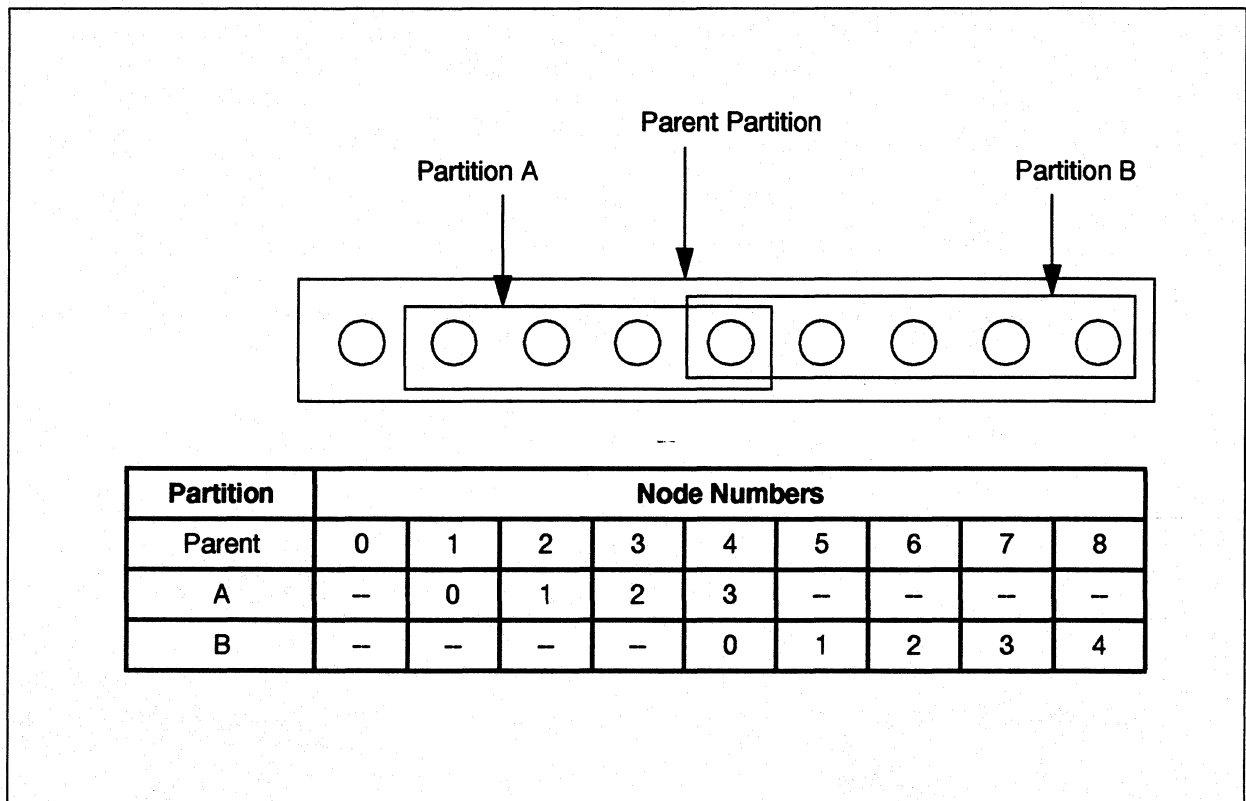


Figure 2-3. Node Numbers in Overlapping Partitions

## Owner, Group, and Protection Modes

Each partition has an *owner*, a *group*, and a set of *protection modes*, like those of a file or directory, that determine who can perform what operations on the partition.

When you create a partition with **mkpart**, you become the new partition's owner; the new partition's group is set to your current group (see **newgrp** in the *OSF/1 Command Reference* for more information on groups). If you are the owner of a partition, you can use **chpart** to change an existing partition's group; only the system administrator can change an existing partition's ownership.

A partition's protection modes consist of three groups of three *permission bits* (read, write, and execute for owner; read, write, and execute for group; and read, write, and execute for "other"), as described for the **chmod** command in the *OSF/1 Command Reference*. The read, write, and execute permission bits have the following meanings for a partition:

- |                    |   |
|--------------------|---|
| <b>r</b> (read)    | Allows listing the subpartitions and characteristics of the partition.                                    |
| <b>w</b> (write)   | Allows creating and removing subpartitions in the partition and changing the partition's characteristics. |
| <b>x</b> (execute) | Allows executing applications in the partition.   |

The system administrator (*root*) is not affected by these permission bits. *root* can do anything to any partition at any time.

The permission bits can be expressed as a three-digit octal number (as for the **chmod** command) or as a string of the form *rwXrwxrwx* (as used by the **ls -l** command, where a letter represents a bit that is "on" and a dash (-) represents a bit that is "off"). For example, the octal number 754 is equivalent to the string *rwXr-xr--*; both grant all permissions to the owner, read and execute permissions to the group, and read permission only to all other users.

When you create a partition with **mkpart**, you can specify its protection modes. If you don't specify a partition's protection modes when you create it, they are set to the same values as those of the parent partition. If you are the owner of a partition or the system administrator, you can use **chpart** to change an existing partition's protection modes.

## Scheduling Characteristics

Each partition has a set of *scheduling characteristics* that determine how the applications running in the partition are scheduled (that is, how the system arbitrates between processes when there are several processes running on a single node).

You can specify a partition's scheduling characteristics when you create it with **mkpart** and change them with **chpart**. If you don't specify a partition's scheduling characteristics when you create it, they are set to the same values as those of the parent partition.

A partition uses one of three different forms of scheduling: *standard scheduling*, *gang scheduling*, or *space sharing*.

- Partitions that use *standard scheduling* use the standard OSF/1 scheduling mechanisms. This gives good response to user input, but may result in poor performance for parallel applications (when one process in the application becomes inactive, other processes that depend on that process for information have to wait until it becomes active again).
- Partitions that use *gang scheduling* use a modified scheduling mechanism that makes all the processes in a parallel application active at the same time. Also, where standard scheduling swaps processes in and out frequently (typically every 100 milliseconds), gang scheduling swaps applications in and out on the basis of the partition's *rollin quantum*: a time period that can be up to 24 hours long. A long rollin quantum gives good performance for parallel applications, because the application can run for a long time without being interrupted, but may result in poor response to user input (when you give input to an application that is rolled out, the application does not respond until it is rolled in again).
- Partitions that use *space sharing* allow only one application per node. When you run an application in a space-shared partition, the partition checks to see if another application or partition is already using the requested nodes. If any of the nodes are in use, your application fails immediately with the error message "request overlaps with nodes in use." However, if all the specified nodes are available, your application begins running immediately and continues running, without interruption, until it completes.

Standard-scheduled partitions should be used to run interactive applications and applications that are being debugged; gang-scheduled and space-shared partitions should be used to run non-interactive (typically either computationally-intensive or I/O-intensive) applications.

The following sections give you more information about these three forms of scheduling.

### Standard Scheduling

*Standard scheduling* is the same as the scheduling technique used on single-processor OSF/1 systems. Standard scheduling is always used in the service partition.

In a partition that uses standard scheduling, each node is scheduled like a separate computer; there is no attempt to coordinate related processes running on separate processors.

### NOTE

A partition that uses standard scheduling may not contain subpartitions, and may not overlap any other partitions that use standard scheduling.

In a partition that uses standard scheduling, each process has a *priority*, a number from -20 (high priority) to 20 (low priority), that is used in determining how much processor time the process gets.

Partitions that use standard scheduling give good interactive performance for each individual process in the partition. However, there is no guarantee that related processes are active at the same time. This means that a process in a parallel application running in such a partition may find itself waiting for a message from a process that is not active, which reduces the performance of the application. To avoid this problem, you can use *gang scheduling*.

### Gang Scheduling

*Gang scheduling* is a special scheduling technique that coordinates the scheduling of related processes running on separate processors. Gang scheduling is typically used only in the compute partition, or is not used at all (this is determined by your system administrator).

In a partition that uses gang scheduling, the nodes are scheduled so that all the processes in an application are active at the same time. If there are multiple processes per node in the active application, standard scheduling is used to schedule these processes against each other while the application is active.

Partitions that use gang scheduling may contain subpartitions, and may overlap other partitions of any type.

In a partition that uses gang scheduling, not only does each process have a priority, but there is a separate priority for the application as a whole. An application's priority is a number from 0 (low priority) to 10 (high priority). A gang-scheduled partition also has a priority of its own, as well as two other quantities called the *effective priority limit* and the *rollin quantum*:

- A partition's *priority* is the lower of the following:
  - The priority of the highest-priority application or subpartition in the partition.
  - The partition's effective priority limit.
- A partition's *effective priority limit* is a number from 0 to 10 that places an upper limit on the partition's priority. It does not affect the priorities of applications or partitions within the partition.
- A partition's *rollin quantum* is the amount of time each application in the partition is allowed to be active before the system considers running another application instead. The term "rollin quantum" comes from the application being "rolled in" when it is made active, and "rolled out" when it is made inactive.

A gang-scheduled partition's effective priority limit and rollin quantum are set when the partition is created, and do not vary unless you change them with the **chpart** command. A gang-scheduled partition's priority may vary over time, depending on the priorities of the applications and subpartitions in the partition.

A partition that uses standard scheduling does not have an effective priority limit or rollin quantum. It also does not have a numeric priority; instead, its priority is “infinite” (that is, higher than the priority of any gang-scheduled partition or application).

Gang scheduling is performed recursively, partition by partition. For each gang-scheduled partition in the system, starting with the root partition, the operating system examines all the entities (applications and partitions) within the partition:

1. Entities that do not overlap other entities (that is, they have no nodes in common with any other entity in the partition) are simply scheduled to run for the partition's rollin quantum.
2. Where two or more entities overlap, the priorities of the overlapping entities are compared, and the highest-priority entity is scheduled to run for the partition's rollin quantum.
3. If two or more entities overlap and are tied for highest priority, they are scheduled in a round-robin fashion (each takes turns running for one full rollin quantum).
4. If an entity that is scheduled to run is a partition, the operating system examines and schedules the entities in the partition as described above. This process continues recursively as necessary.

At the end of each partition's rollin quantum, the operating system examines and schedules the entities in the partition again.

Note that rules 2 and 3 mean that, when applications or partitions overlap, the one with the highest priority gets one rollin quantum after another until it completes. Entities with lower priorities get no processor time at all until the higher-priority entity has completed. If there is a tie for highest priority, the tied high-priority entities take turns running, but entities with lower priority get no processor time until *all* the high-priority entities complete. Partitions that use standard scheduling always have the highest priority, so if a standard-scheduled partition overlaps a gang-scheduled partition or an application, the standard-scheduled partition always wins.

## NOTE

Use of gang scheduling may be limited by the policies of your site.

Your system administrator can require all compute partitions to use space sharing. If gang scheduling is allowed, the administrator can restrict the number of gang-scheduled partitions in the system, can set a minimum rollin quantum, and can restrict the number of applications that can overlap in each gang-scheduled partition. If you try to create a partition that would exceed these restrictions, you see an error message such as “exceeded allocator configuration parameters” or “scheduling parameters conflict with allocator configuration.” See your system administrator for information on the policies in force at your site.

## Space Sharing

*Space sharing*, also referred to as *tiling*, is a scheduling technique that prevents partitions and applications from overlapping. (*Overlapping* means having any physical nodes in common.) Space sharing is typically used in all partitions other than the service and compute partitions. If your system administrator has disallowed gang scheduling, space sharing is used in all partitions other than the service partition. Within a space-shared partition:

- Subpartitions may not overlap other subpartitions.
- Applications may not overlap other applications.
- Active subpartitions may not overlap applications.

An *active* subpartition is a subpartition in which one or more applications is running.

### NOTE

If an application is running *anywhere* in a subpartition or any of its sub-subpartitions—even on a single node—the *entire* subpartition is considered active, and is not allowed to overlap with a running application.

If a subpartition is not active (contains no running applications), it can overlap a running application, but it cannot overlap another partition.

In a space-shared partition, any attempt to create a partition or run an application that would cause an overlap fails immediately. However, once an application is successfully started, it continues running without interruption until it completes. (Exception: if a space-shared partition overlaps with another partition, the entire partition can be interrupted by applications running in that other partition. This can only occur if the space-shared partition's parent is a gang-scheduled partition.)

Space sharing is the opposite of the "time sharing" used in standard scheduling and gang scheduling. In time sharing, multiple applications can use the same nodes at the same time, but each application gets only a fraction of its nodes' processor time. In space sharing, no two applications can use a node at the same time, but each application gets 100% of its nodes' processor time.

Although space sharing allows only one *application* per node, you can have more than one *process* per node within a single application. If there are multiple processes per node within an application, standard scheduling is used to schedule these processes against each other on each node.

Partitions that use space sharing may contain subpartitions, which cannot overlap. The space-shared partition itself can overlap another partition of any type, but the advantages of space sharing may be lost if space-shared partitions overlap with other partitions.

Like gang-scheduled partitions, space-shared partitions have a priority and an effective priority limit. Each application within a space-shared partition has a priority from 0 to 10, and the partition's priority is the lesser of the effective priority limit and the highest application priority in the partition. Since applications in space-scheduled partitions never overlap, their priorities are never compared with each other. However, the priorities of applications in a space-scheduled partition are important because they determine the partition's priority when compared with other partitions at its own hierarchical level.

Unlike gang-scheduled partitions, space-shared partitions do not have a rollin quantum (since applications never overlap, they never have to be rolled in or out). In effect, the rollin quantum of a space-shared partition is "infinite."

### Summary of Scheduling Types

Table 2-1 summarizes the differences between the three scheduling types.

**Table 2-1. Summary of Scheduling Types**

<b>Characteristic</b>	<b>Standard Scheduling</b>	<b>Gang Scheduling</b>	<b>Space Sharing</b>
<b>Scheduling method used within partition</b>	Each process is scheduled by itself using standard UNIX techniques	All processes in an application run at the same time; applications may be rolled in and out	All processes in an application run at the same time; each application runs until it completes
<b>Partitions that typically use this scheduling type</b>	Service partition	Compute partition, or none at all	All other partitions
<b>Restrictions on partition overlap</b>	Partition may not overlap other standard-scheduled partitions	Partition may overlap other partitions	Partition may overlap other partitions (but overlap can lose benefits of space sharing)
<b>Restrictions on subpartition overlap</b>	Subpartitions are not allowed	Subpartitions may overlap; maximum depth of overlap can be restricted by system administrator	Subpartitions may not overlap other subpartitions; active subpartitions may not overlap applications
<b>Restrictions on application overlap</b>	Applications may overlap to any depth	Applications may overlap; maximum depth of overlap can be restricted by system administrator	Applications may not overlap other applications or active subpartitions
<b>Special partition characteristics</b>	Partition priority (always "infinite")	Partition priority, effective priority limit, rollin quantum	Partition priority, effective priority limit



## A Scheduling Example

Suppose that a partition has 10 nodes, and an application is currently running on 5 of those nodes. If you attempt to run a new application on 6 nodes of that partition, the results depend on the partition's scheduling type:

- If the partition uses standard scheduling, both applications run at once. Where the applications overlap, the two applications' processes time-share the node. No attempt is made to coordinate when the processes are active with the rest of the application.
- If the partition uses gang scheduling, the two applications' priorities are compared:
  - If the new application's priority is greater than the old application's, the entire old application is immediately rolled out and the new application starts running. The new application runs until it finishes, then the old application is rolled back in.
  - If the new application's priority is less than the old application's, the entire new application waits until the old application finishes. (During this time it may appear to be "hung.") When the old application finishes, the new application is rolled in and runs until it finishes.
  - If the two applications' priorities are equal, the applications alternate running on each rollin quantum. If one application finishes first, the other runs in every rollin quantum until it finishes.
- If the partition uses space sharing, the new application fails with the error message "request overlaps with nodes in use" and does not run.

You can use the **pspart** command to determine which applications are currently running in a partition and what their priorities are, and you can use the command **showpart -f** to determine which nodes in a partition have applications running on them.

## Making Partitions

Command Synopsis	Description
<b>mkpart</b> [ <i>-sz size</i>   <i>-sz hXw</i>   <i>-nd nodespec</i> ] [ <i>-ss</i>   [ <i>-sps</i>   <i>-rq time</i> ] [ <i>-epl priority</i> ] ] ] [ <i>-mod mode</i> ] <i>name</i>	Create a partition.

To create a partition, use the **mkpart** command. You can specify either a relative or an absolute partition pathname for the new partition. The specified new partition must not exist; the parent partition of the new partition must exist and must grant you write permission.

For example, to create a partition called *mypart* whose parent partition is the compute partition, you can use the following command:

```
% mkpart mypart
```

The following command has the same effect, but uses an absolute partition pathname:

```
% mkpart .compute.mypart
```

## Specifying the Nodes Allocated to the Partition

The **mkpart** command gives you three ways to specify which nodes are allocated to the new partition:

- sz *size*** Creates a partition whose size (number of nodes) is *size*. The **-sz *size*** switch attempts to create a square partition if it can. If this is not possible, it attempts to create a rectangular partition that is either twice as wide as it is high or twice as high as it is wide. If this is not possible, it uses any available nodes. In this case, the nodes allocated to the partition may not be contiguous.
- sz *hXw*** Creates a contiguous rectangular partition that is *h* nodes high and *w* nodes wide. (You can use an uppercase or lowercase letter **X** between the integers *h* and *w*.)
- nd *nodespec*** Creates a partition that consists of exactly the specified nodes, where *nodespec* is one of the following:
  - x* The node whose node number is *x*.
  - x..y* The range of nodes from numbers *x* to *y*.
  - hXw:n* The rectangular group of nodes that is *h* nodes high and *w* nodes wide and whose upper left corner is node number *n*. (You can use an uppercase or lowercase letter **X** between the integers *h* and *w*.)
  - nspec[,nspec]...* The specified list of nodes, where each *nspec* is a node specifier of the form *x*, *x..y*, or *hXw:n* (no node may appear more than once in this list). Do not put any spaces in this list.

The numbers you use with **-nd** are node numbers within the parent partition, which always range from 0 to one less than the size of the partition.

If you don't use the **-sz** or **-nd** switch, all the nodes of the parent partition are allocated to the new partition. You can use at most one **-sz** or **-nd** switch in a single **mkpart** command.

The following examples all create a 50-node partition called *mypart* whose parent partition is the compute partition (that is, the new partition's absolute partition pathname is *.compute.mypart*):

- This command creates a 50-node partition with no specified shape or location:

```
% mkpart -sz 50 mypart
```

The nodes of the new partition are selected from the parent partition by the system, and they may not be contiguous.

- This command creates a partition 10 nodes high and 5 nodes wide:

```
% mkpart -sz 10x5 mypart
```

The position of the new partition within the parent partition is selected by the system, but the new partition is a contiguous rectangle.

- This command creates a partition 10 nodes high and 5 nodes wide located in the upper left corner of the parent partition:

```
% mkpart -nd 10X5:0 mypart
```

The shape and position of the new partition are specified by the user, and the new partition is a contiguous rectangle.

- This command creates a partition that consists of nodes 30 through 79 of the parent partition:

```
% mkpart -nd 30..79 mypart
```

The specific nodes of the partition are specified by the user, and the new partition may or may not be contiguous (its shape depends on the size and shape of the compute partition).

- This command creates a partition that consists of node 0, nodes 3 through 16, and a 5 by 7 node rectangle located at node 21 of the parent partition:

```
% mkpart -nd 0,3..16,5X7:21 mypart
```

The specific nodes of the partition are specified by the user, and the new partition is not contiguous (its shape depends on the size and shape of the compute partition).

No matter how you specify the partition's size, nodes are always numbered from 0 to one less than the partition's size. In most cases, they are numbered from left to right and then top to bottom, as discussed under "Nodes Allocated to the Partition" on page 2-30. However, if you use the **-nd** switch, the nodes in the new partition are numbered in the order you specified them in the **-nd** switch. For example, the following command creates a partition that consists of nodes 30 through 79 of the compute partition:

```
% mkpart -nd 79..30 mypart
```

In this case, node 79 of the parent partition is node 0 of the new partition; node 78 of the parent partition is node 1 of the new partition; and so on to node 30 of the parent partition, which is node 49 of the new partition.

## Specifying Protection Modes

The **mkpart** command gives you two ways to specify the protection modes of the new partition:

- mod *nnn*** Creates a partition whose protection modes are specified by the three-digit octal number *nnn*, as used by the **chmod** command (see **chmod** in the *OSF/1 Command Reference* for more information).
- mod *string*** Creates a partition whose protection modes are specified by the nine-character string *string*. The string must have the form *rwxrwxrwx*, where a letter (r, w, or x) represents a permission granted and a dash (-) represents a permission denied, as displayed by the command **ls -l** (see **ls** in the *OSF/1 Command Reference* for more information).

You can use at most one **-mod** switch in a single **mkpart** command. If you don't use the **-mod** switch, the new partition is given the same protection modes as its parent partition.

For example, the following command creates a partition that is readable, writable, and executable by you; readable and executable by your group, and only readable by others:

```
% mkpart -mod rwxr-xr-- mypart
```

The following command has the same effect, but uses an octal number:

```
% mkpart -mod 754 mypart
```

## Specifying Scheduling Characteristics

The **mkpart** command gives you three switches to specify the scheduling characteristics of the new partition:

<b>-ss</b>	Creates a partition that uses standard scheduling.
	<b>-ss</b> cannot be used with <b>-sps</b> , <b>-rq</b> or <b>-epl</b> .
<b>-rq <i>time</i></b>	Creates a partition that uses gang scheduling with a rollin quantum of <i>time</i> , where <i>time</i> is one of the following:
<i>n</i>	<i>n</i> milliseconds (if <i>n</i> is not a multiple of 100, it is silently rounded up to the next multiple of 100).
<i>ns</i>	<i>n</i> seconds.
<i>nm</i>	<i>n</i> minutes.
<i>nh</i>	<i>n</i> hours.
<b>0</b>	“Infinite” time: once rolled in, an application runs until it exits.

The maximum rollin quantum is 24 hours; the minimum rollin quantum for your system is determined by your system administrator.

**-rq** cannot be used with **-ss** or **-sps**. **-rq** can be used with or without **-epl**; if you use **-rq** without **-epl**, the new partition is a gang-scheduled partition with the same effective priority limit as its parent partition.

If gang-scheduled partitions are not allowed at your site, or creating a gang-scheduled partition would exceed the maximum number of gang-scheduled partitions, any attempt to create a partition with **-rq** fails.

<b>-sps</b>	Creates a partition that uses space sharing.
	<b>-sps</b> cannot be used with <b>-ss</b> or <b>-rq</b> . <b>-sps</b> can be used with or without <b>-epl</b> ; if you use <b>-sps</b> without <b>-epl</b> , the new partition is a space-shared partition with the same effective priority limit as its parent partition.

**-epl *priority*** Creates a partition with an effective priority limit of *priority*, where *priority* is an integer from 0 to 10 inclusive (0 is low priority, 10 is high priority).

**-epl** cannot be used together with **-ss**. If you use **-epl** without either **-sps** or **-rq**, the results depend on the scheduling type of the parent partition:

- If the parent partition is a space-shared partition, the new partition is a space-shared partition with the specified effective priority limit.
- If the parent partition is a gang-scheduled partition, the new partition is a gang-scheduled partition with the specified effective priority limit and the same rollin quantum as its parent. If this would exceed the maximum number of gang-scheduled partitions, the new partition is a space-shared partition instead.

If you don't use the **-ss**, **-rq**, or **-sps** switch, the new partition uses the same scheduling technique, rollin quantum, and effective priority limit as its parent partition.

For example, the following command creates a partition that uses standard scheduling:

```
% mkpart -ss mypart
```

The following command creates a partition that uses gang scheduling with a rollin quantum of 10 seconds and the same effective priority limit as its parent partition:

```
% mkpart -rq 10s mypart
```

The following command creates a partition that uses space sharing with the same effective priority limit as its parent partition:

```
% mkpart -sps mypart
```

The following command creates a partition that uses gang scheduling with a rollin quantum of 5 minutes and an effective priority limit of 6:

```
% mkpart -rq 5m -epl 6 mypart
```

## Removing Partitions

---

Command Synopsis	Description
<b><code>rmpart</code></b> [ <b><code>-f</code></b> ] [ <b><code>-r</code></b> ] <i>partition</i>	Remove a partition.

---

To remove an existing partition, use the **rmpart** command. You must have write permission on the parent partition of the partition to be removed. You can specify the partition to be removed with either a relative or an absolute partition pathname.

For example, to remove the partition called *mypart*, whose parent partition is the compute partition, you can use the following command:

```
% rmpart mypart
```

The following command has the same effect, but uses an absolute partition pathname:

```
% rmpart .compute.mypart
```

## Removing Partitions Containing Running Applications

If you specify a partition that contains any running applications, you see an error message and the partition is not removed. You can force **rmpart** to remove a partition that contains running applications with the **-f** switch. When you use the **-f** switch, **rmpart** terminates all the applications running in the specified partition and then removes it.

For example, if there are applications running in *mypart*, use the following command to terminate the applications and remove the partition:

```
% rmpart -f mypart
```

## Removing Partitions Containing Subpartitions

If you specify a partition that contains any subpartitions, you see an error message and the partition is not removed. You can force **rmpart** to remove a partition that contains subpartitions with the **-r** switch. When you use the **-r** switch, **rmpart** recursively removes all the subpartitions in the specified partition (and their sub-subpartitions, and so on) and then removes the specified partition.

For example, if there are subpartitions in *mypart*, use the following command to remove *mypart* and all its subpartitions:

```
% rmpart -r mypart
```

**rmpart -r** is an “all or nothing” operation. If any subpartitions cannot be removed, the command fails and no subpartitions are removed.

The **-r** switch does not imply **-f**. If *mypart* or any of its subpartitions contains any running applications, you see an error message and none of the partitions are removed. You can force **rmpart** to remove a partition that contains subpartitions and running applications by using the **-r** and **-f** switches together. When you use both these switches, **rmpart** terminates all the applications running in the specified partition and its subpartitions, removes all the subpartitions in the specified partition, and then removes the specified partition.

## Showing Partition Characteristics

Command Synopsis	Description
<b>showpart</b> [ <b>-f</b> ] [ <i>partition</i> ]	Show the characteristics of a partition.

To show the characteristics of a partition, use the **showpart** command. You can specify the partition with either a relative or an absolute partition pathname. If you don't specify a partition, **showpart** shows the characteristics of your default partition (see “Using the Default Partition” on page 2-14). In either case, you must have read permission on the specified partition.

For example, to show the characteristics of the partition called *mypart*, whose parent partition is the compute partition, you can use the following command:

```
% showpart mypart
USER      GROUP    ACCESS  SIZE      FREE  RQ   EPL
smith     eng      777     9         5    15m  5
+-----+
0| . . . . |
4| . * * * |
8| . * * * |
12| . * * * |
+-----+
```

The following command has the same effect, but uses an absolute partition pathname:

```
% showpart .compute.mypart
```



The columns at the top of the **showpart** output have the following meanings:

USER	The owner of the partition, in this case <i>smith</i> .						
GROUP	The group of the partition, in this case <i>eng</i> .						
ACCESS	The access permissions, expressed as an octal number, in this case 777 (which represents the permissions <i>rwxrwxrwx</i> ).						
SIZE	The number of nodes in the partition, in this case 9.						
FREE	The number of free nodes in the partition, in this case 5 (see "Showing Free Nodes" on page 2-48 for more information on free nodes).						
RQ	The rollin quantum or scheduling type of the partition, as follows: <table> <tr> <td>-</td> <td>The partition uses standard scheduling.</td> </tr> <tr> <td>SPS</td> <td>The partition uses space sharing.</td> </tr> <tr> <td><i>time</i></td> <td>The partition uses gang scheduling with a rollin quantum of <i>time</i>. The <i>time</i> is expressed as a number followed by an optional letter: no letter for milliseconds, <i>s</i> for seconds, <i>m</i> for minutes, or <i>h</i> for hours.</td> </tr> </table> <p>In this case, the partition is a gang-scheduled partition with a rollin quantum of 15 minutes.</p>	-	The partition uses standard scheduling.	SPS	The partition uses space sharing.	<i>time</i>	The partition uses gang scheduling with a rollin quantum of <i>time</i> . The <i>time</i> is expressed as a number followed by an optional letter: no letter for milliseconds, <i>s</i> for seconds, <i>m</i> for minutes, or <i>h</i> for hours.
-	The partition uses standard scheduling.						
SPS	The partition uses space sharing.						
<i>time</i>	The partition uses gang scheduling with a rollin quantum of <i>time</i> . The <i>time</i> is expressed as a number followed by an optional letter: no letter for milliseconds, <i>s</i> for seconds, <i>m</i> for minutes, or <i>h</i> for hours.						
EPL	The effective priority limit of the partition, in this case 5, or a dash (-) for a standard-scheduled partition.						

See "Partition Characteristics" on page 2-29 for information on these partition characteristics.

The rectangular picture at the bottom of the **showpart** output shows the size, shape, and position of the specified partition within the system:

- The large rectangle represents the root partition. In this case, the root partition is 4 nodes high and 4 nodes wide.
- The numbers to the left of the rectangle show the node numbers of the nodes in the first column of each row. In this case, the first node in the top row is node 0, the first node in the second row is node 4, the first node in the third row is node 8, and the first node in the bottom row is node 12.

- Asterisks (\*) within the rectangle represent nodes that are allocated to the specified partition; periods (.) represent other nodes. In this case, *mypart* consists of nodes 5–7, 9–11, and 13–15 of the root partition.
- If you see a dash (-) or an X within the rectangle, it represents an unusable node that is allocated to the specified partition. You cannot run any applications or allocate any partitions using this node. See “Unusable Nodes” on page 2-31 for more information.

## Showing Free Nodes

The output of **lspart** or **showpart** includes the number of free nodes in the FREE column. A node is *free* if no application is running on that node and no subpartition in which any applications are running includes that node. (Note that *all* the nodes of a subpartition are considered busy if an application is running *anywhere* in the subpartition, or in any of its sub-subpartitions. This occurs because partitions are scheduled recursively.)

You can use the **-f** switch of **showpart** to see *which* nodes are free. The output of **showpart -f** is the same as the regular **showpart** output, except that free nodes are shown as an F instead of an asterisk.

For example, to show the free nodes in the partition called *mypart*, whose parent partition is the compute partition, you can use the following command:

```
% showpart -f mypart
USER      GROUP    ACCESS  SIZE      FREE  RQ    EPL
smith     eng      777     9         5    15m   5
+-----+
0| . . . . |
4| . * * * |
8| . * F F |
12| . F F F |
+-----+
```

In this case, *mypart* has five free nodes: nodes 4, 5, 6, 7, and 8 of the partition.

## Listing Subpartitions

Command Synopsis	Description
<b>lspart</b> [ -r ] [ <i>partition</i> ]	List the subpartitions of a partition.

To list the subpartitions of a partition with their characteristics, use the **lspart** command. You can specify the partition with either a relative or an absolute partition pathname. If you don't specify a partition, **lspart** lists the subpartitions of your default partition (see "Using the Default Partition" on page 2-14). In either case, you must have read permission on the specified partition.

For example, to list the subpartitions of the partition called *mypart*, whose parent partition is the compute partition, you can use the following command:

```
% lspart mypart
USER      GROUP    ACCESS  SIZE    FREE  RQ    EPL  PARTITION
chris     eng      777     16      4    15m   3    mandelbrot
chris     eng      777     16     16    -     -     debug
pat       mrkt     755     4       0    SPS   10    slalom
*         *        *        *        *     *     *     private
```

The following command has the same effect, but uses an absolute partition pathname:

```
% lspart .compute.mypart
```

The columns in the output of **lspart** are the same as the top part of the output of **showpart** (see "Showing Partition Characteristics" on page 2-46), with the addition of the partition name. In this case, *mypart* has four subpartitions: *mandelbrot*, *debug*, *slalom*, and *private*.

- *mandelbrot* is owned by user *chris* in group *eng*; it has permissions *rwXrwXrwX* and a size of 16 nodes, of which 4 are free (see "Showing Free Nodes" on page 2-48 for more information on free nodes). It is a gang-scheduled partition with a rollin quantum of 15 minutes and an effective priority limit of 3.
- *debug* is also owned by user *chris* in group *eng*; it has permissions *rwXrwXrwX* and a size of 16 nodes, of which all 16 are free. It is a standard-scheduled partition, so it has no rollin quantum or effective priority limit.
- *slalom* is owned by user *pat* in group *mrkt*; it has permissions *rwXr-Xr-X* and a size of 4 nodes, of which none are free. It is a space-shared partition with an effective priority limit of 10.
- *private*'s access permissions do not grant you read permission, so all its characteristics are shown as asterisks (\*).

If you see two numbers separated by a slash in the **SIZE** column, it indicates that one or more of the nodes allocated to the indicated partition is unusable. For example:

```
% lspart mypart
USER      GROUP    ACCESS  SIZE      FREE  RQ    EPL  PARTITION
chris     eng       777     14 / 16   10   15m   3    mandelbrot
```

This indicates that there are 16 nodes allocated to *mandelbrot*, but 2 of them are currently unusable. You cannot run any applications or allocate any partitions using unusable nodes. See “Unusable Nodes” on page 2-31 for more information.

## Recursively Listing Subpartitions

To recursively list all of a partition's subpartitions, sub-subpartitions, and so on, use the **-r** switch. For example:

```
% lspart -r mypart
USER      GROUP    ACCESS  SIZE      FREE  RQ    EPL  PARTITION
.compute.mypart:
chris     eng       777     16         4   15m   3    mandelbrot
chris     eng       777     16        16    -    -    debug
pat       mrkt     755      4         0   SPS   10   slalom
*         *         *         *         *    *    *    private
.compute.mypart.mandelbrot:
chris     eng       777     16        16   15m   10   hi_pri
chris     eng       777     16        16   15m   1    lo_pri
```

The **lspart -r** output reveals that *mypart.mandelbrot* has two subpartitions, *hi\_pri* and *lo\_pri*, neither of which has any sub-subpartitions, and that *slalom* and *debug* have no subpartitions. No information is available on the subpartitions of *private* (if any), because *private* does not grant you read permission.

### NOTE

If you specify a partition that has no subpartitions, **lspart** produces no output.

For example, since *mypart.slalom* has no subpartitions, an **lspart** command on this partition gives no output:

```
% lspart mypart.slalom
%
```

To get information about *mypart.slalom* itself, use the **showpart** command.

## Listing the Applications in a Partition

Command Synopsis	Description
<b>pspart</b> [ -r ] [ <i>partition</i> ]	List the applications in a partition.

To list the applications in a partition, with information about the rollin/rollout status of each, use the **pspart** command. You can specify the partition with either a relative or an absolute partition pathname. If you don't specify a partition, **pspart** lists the applications in your default partition (see "Using the Default Partition" on page 2-14). In either case, you must have read permission on the specified partition.

For example, to list the applications in the partition *mypart*, whose parent partition is the compute partition, you can use the following command:

```
% pspart mypart
  PGID  USER      SIZE  PRI  START          TIME  ACTIVE  TOTAL  TIME  COMMAND
12345  pat        256   5   11:42:20      45.00  75%    0:04:41  mag -sz 256
23456  chris      67    4   Jan 21         -      -      0:12.30  boggle
34567  smith     192  10  02:21:51      0:01:00 100%    2:12:03  myfft
```

The following command has the same effect, but uses an absolute partition pathname:

```
% pspart .compute.mypart
```

The columns in the output of **pspart** have the following meanings:

PGID	The process group ID of the application (see "Process Groups" on page 4-22 for more information).  The process group ID of an application is always the same as the process ID of the application's controlling process. This means that you can use this number with the <b>kill</b> command to kill the application; for example, given the <b>pspart</b> output above, the command <b>kill 34567</b> would kill the application <i>myfft</i> .
USER	The login name of the user who invoked the application.
SIZE	The number of nodes allocated to the application from the partition (see "Specifying Application Size" on page 2-15 for more information).
PRI	The application's priority (see "Specifying Application Priority" on page 2-17 for more information).
START	The time the application was started. If the application was started more than 24 hours ago, the date it was started is shown instead.

**TIME ACTIVE** The amount of time the application has been active (rolled in) in the current rollin quantum (see “Gang Scheduling” on page 2-35 for more information). The time active is shown both as an absolute time (in the format *minutes : seconds . milliseconds* for times less than one minute or *hours : minutes : seconds* for times of one minute or more) and as a percentage of the partition's rollin quantum. If the application is not active in the current rollin quantum, a dash (-) is shown for both quantities. If the partition uses space sharing, the time shown is the total amount of time the application has been running and the percentage is always 100%.

In the example above, the partition *mypart* is a gang-scheduled partition with a rollin quantum of one minute. The application *mag* has been active for 45 seconds, or 75% of the rollin quantum; the application *boggle* is not currently active; and the application *myfft* has been active for one minute, or 100% of the rollin quantum.

**TOTAL TIME** The total amount of time the application has been rolled in since it was started, in the format *minutes : seconds . milliseconds* or *hours : minutes : seconds*. If the partition uses space sharing, the **TOTAL TIME** is always the same as the **TIME ACTIVE**.

In the example above, the application *mag* has been active for a total of 4 minutes and 41 seconds; the application *boggle* has been active for a total of 12.30 seconds; and the application *myfft* has been active for a total of 2 hours, 12 minutes, and 3 seconds.

**COMMAND** The command line by which the application was invoked.

## Applications in Subpartitions

If there are any applications running in subpartitions of the specified partition, the subpartitions appear in the output of **pspart** as follows:

```
% pspart mypart
  PGID  USER      SIZE PRI  START          TIME ACTIVE  TOTAL TIME  COMMAND
  12345  pat         256  5   11:42:20      45.00  75%      0:04:41  mag -sz 256
  23456  chris        67  4    Jan 21         -      -         0:12.30  boggle
  34567  smith       192 10   02:21:51      0:01:00 100%     2:12:03  myfft
Active Partitions
  OWNER  GROUP      SIZE PRI  START          TIME ACTIVE  TOTAL TIME  NAME
smith   eng         64  6   09:16:30         -      -         1:18.10  subpart
```

The columns for the list of active partitions have the following meanings:

<b>OWNER</b>	The owner of the subpartition.
<b>GROUP</b>	The group of the subpartition.
<b>SIZE</b>	The size of the subpartition (note that all nodes of a subpartition containing an active application are considered active, even if not all the nodes in the subpartition are actually in use by applications).
<b>PRI</b>	The current priority of the subpartition (this is the highest priority of all the applications in the subpartition or the subpartition's effective priority limit, whichever is lower).
<b>START</b>	The time or date when the oldest application in the subpartition was started.
<b>TIME ACTIVE</b>	The amount of time the subpartition has been active (rolled in) in the current rollin quantum.
<b>TOTAL TIME</b>	The total amount of time the subpartition has been rolled in since it was started.
<b>NAME</b>	The name of the subpartition.

See "Scheduling Characteristics" on page 2-33 for more information on how subpartitions are scheduled.

## Recursively Listing Applications in Subpartitions

If there are applications running in a subpartition, the output of `pspart` normally shows only that the subpartition is active. To list the applications in subpartitions (and, recursively, in sub-subpartitions and so on), use the `-r` switch. For example:

```
% pspart -r mypart
mypart:
  PGID  USER      SIZE PRI  START      TIME ACTIVE  TOTAL TIME COMMAND
  12345  pat        256  5  11:42:20   45.00 75%    0:04:41 mag -sz 256
  23456  chris      67   4   Jan 21     -    -      0:12.30 boggle
  34567  smith     192  10  02:21:51   0:01:00 100%   2:12:03 myfft
Active Partitions
  OWNER  GROUP      SIZE PRI  START      TIME ACTIVE  TOTAL TIME NAME
smith   eng        64   6  09:16:30   -    -      1:18.10 subpart
mypart.subpart:
  PGID  USER      SIZE PRI  START      TIME ACTIVE  TOTAL TIME COMMAND
  45678  smith     56   7  09:16:30   -    -      1:18.10 span
```

In this case, the `-r` switch shows that the subpartition *subpart* has one application, *span*, which is running on 56 nodes of the subpartition. (Note that even though the application is not running on every node of the subpartition, whenever the application is rolled in the entire subpartition is rolled in. This occurs because subpartitions are scheduled recursively, as discussed under “Gang Scheduling” on page 2-35.)

## Changing Partition Characteristics

Command Synopsis	Description
<pre><b>chpart</b> [ <b>-rq</b> <i>time</i>   <b>-sps</b> ] [ <b>-ep</b> <i>priority</i> ]           [ <b>-nm</b> <i>name</i> ] [ <b>-mod</b> <i>mode</i> ]           [ <b>-g</b> <i>group</i> ] [ <b>-o</b> <i>owner</i> [ <i>.group</i> ] ]           <i>partition</i></pre>	Change certain partition characteristics.

To change the characteristics of a partition, use the `chpart` command. The permissions required depend on the switches you use. You can specify the partition with either a relative or an absolute partition pathname.



**chpart** can change the following partition characteristics:

- Rollin quantum.
- Effective priority limit.
- Partition name.
- Protection modes.
- Owner and group.
- Scheduling type (space-shared to gang-scheduled, or gang-scheduled to space-shared with certain limitations; a partition cannot be changed to or from standard scheduling).

A partition's size and parent partition are determined when the partition is created and cannot be changed.

The switches of **chpart**, which can be used together or separately and in any order (except as noted below), are similar to the corresponding switches of **mkpart**:

<b>-rq</b> <i>time</i>	Changes the partition to a gang-scheduled partition with a rollin quantum of <i>time</i> , where <i>time</i> is one of the following:
<i>n</i>	<i>n</i> milliseconds (if <i>n</i> is not a multiple of 100, it is rounded up to the next multiple of 100).
<i>ns</i>	<i>n</i> seconds.
<i>nm</i>	<i>n</i> minutes.
<i>nh</i>	<i>n</i> hours.
<b>0</b>	“Infinite” time: once rolled in, an application runs until it exits.

The maximum rollin quantum is 24 hours; the minimum rollin quantum for your system is determined by your system administrator.

**-rq** can be used only on a gang-scheduled or space-shared partition, and cannot be used together with **-sps**. To use **-rq**, you must have write permission on the specified partition.

- sps** Changes the partition to a space-shared partition.
- sps** can be used only on a space-shared or gang-scheduled partition, and cannot be used together with **-rq**. If the partition is currently gang-scheduled, it must not contain any overlapping subpartitions or any applications. To use **-sps**, you must have write permission on the specified partition.
- epl *priority*** Changes the partition's effective priority limit to *priority*, where *priority* is an integer from 0 to 10 inclusive.
- epl** can be used only on a gang-scheduled or space-shared partition. To use **-epl**, you must have write permission on the specified partition.
- nm *name*** Changes the partition's name to *name*, where *name* is a valid partition name (a string of any length containing only uppercase letters, lowercase letters, digits, and underscores). To use **-nm**, you must have write permission on the parent partition of the specified partition.
- Note that **-nm** can only change the partition's name "in place;" there is no way to move a partition to a different parent partition.
- mod *nmn*** Changes the partition's protection modes to the value specified by the three-digit octal number *nmn*. To use **-mod**, you must be the owner of the specified partition or the system administrator.
- mod *string*** Changes the partition's protection modes to the value specified by the nine-character string *string*. The string must have the form *rwXrwxrwx*, where a letter (*r*, *w*, or *X*) represents a permission granted and a dash (-) represents a permission denied. To use **-mod**, you must be the owner of the specified partition or the system administrator.
- g *group*** Changes the partition's group to *group*. The *group* can be either a group name or a numeric group ID. To use **-g**, you must be the owner of the specified partition *and* a member of the specified new group, or you must be the system administrator.
- o *owner* [ . *group*]** Changes the partition's owner to *owner*. If *. group* is specified, also changes the partition's group to *group*. The *owner* and *group* can be either user/group names or numeric user/group IDs. To use **-o**, you must be the system administrator.

For example, the following command changes the rollin quantum of *mypart* to 20 minutes:

```
% chpart -rq 20m mypart
```

The following command changes *mypart* to a space-shared partition:

```
% chpart -sps mypart
```

The following command changes the effective priority of *mypart* to 2:

```
% chpart -ep1 2 mypart
```

The following command changes the protection modes of *mypart* so that it is readable, writable, and executable by the owner but not by anyone else:

```
% chpart -mod rwx----- mypart
```

The following command has the same effect as the previous three commands combined, but uses an absolute partition pathname and an octal protection mode specifier:

```
% chpart -ep1 2 -rq 20m -mod 700 .compute.mypart
```

The following command changes the owner of *mypart* to *smith*, but does not affect its group:

```
% chpart -o smith mypart
```

The following command changes the group of *mypart* to *support*, but does not affect its ownership:

```
% chpart -g support mypart
```

The following command changes the owner of *mypart* to *smith* and the group to *support*:

```
% chpart -o smith.support mypart
```

The following command changes the name of *mypart* to *newpart*:

```
% chpart -nm newpart mypart
```

The following command also changes the name of *mypart* to *newpart*, but uses an absolute partition pathname:

```
% chpart -nm newpart .compute.mypart
```

Note that the new name is specified as a name only, not a pathname.



# Using Paragon™ OSF/1 Message-Passing System Calls



3

## Introduction

*Message passing* is the standard means of communication among processes in Paragon OSF/1. As independent processor/memory pairs, the nodes do not share physical memory. If the node processes need to share information, they can do so by passing messages. The calls described in this chapter let your programs send and receive messages.

This chapter introduces the Paragon OSF/1 message-passing system calls. It includes the following sections:

- Process characteristics.
- Message characteristics.
- Names of send and receive calls.
- Synchronous send and receive.
- Asynchronous send and receive.
- Probing for pending messages.
- Getting information about pending or received messages.
- Message passing with Fortran commons.
- Treating a message as an interrupt.
- Extended receive and probe.
- Global operations.

Within each section, the calls are discussed in order of increasing complexity. That is, the “base” calls are discussed first, and the “extended” calls are discussed later.

Each section includes numerous examples in both C and Fortran. A call description at the beginning of each section or subsection gives a language-independent synopsis (call name, parameter names, and brief description) of each call discussed in that section. Differences between C and Fortran are noted where applicable. See Appendix A for information on call and parameter types; see the *Paragon™ C System Calls Reference Manual* or the *Paragon™ Fortran System Calls Reference Manual* for complete information on each call.

This chapter does not describe all the Paragon OSF/1 system calls. For information about system calls that provide general services other than message passing, see Chapter 4. For information about the calls used with the Parallel File System, see Chapter 5. For information about the calls used with graphical interfaces, such as DGL and the X Window System, see the *Paragon™ Graphics Libraries User's Guide*. For information about the system calls that require root privileges, see the *Paragon™ System Administrator's Guide*.

Paragon OSF/1 programs written in C can also issue OSF/1 system calls. The Paragon OSF/1 operating system is a complete OSF/1 system and fully supports all the standard OSF/1 system calls. See the *OSF/1 Programmer's Reference* for information on these calls.

Paragon OSF/1 programs written in Fortran cannot make OSF/1 system calls directly, but the Fortran runtime library includes a number of system interface routines. These routines make a number of OSF/1 system calls available to Fortran programs. See the *Paragon™ Fortran Compiler User's Guide* for information on these routines.

## Process Characteristics

Each process within an application is identified by its *node number* and *process type*. A process must have a valid node number and process type to send and receive messages.

### Node Numbers

Synopsis	Description
<code>mynode()</code>	Obtain the calling process's node number.
<code>numnodes()</code>	Obtain the number of nodes allocated to the current application.

A process's *node number* is an integer that identifies the node on which it is running. Node numbers are assigned by the system, and range from zero to one less than the number of nodes in the application. A process can find out its node number by calling `mynode()`; the node number does not change for the life of the process. A process can also find out the number of nodes in the application by calling `numnodes()`; the maximum node number in the application is `numnodes() - 1`.

When you run an application that was linked with the `-nx` switch, the system creates one process on each node of the default partition (unless you specify otherwise on the application's command line). Each process is the same as the others except for its node number, which is different in each process.

All message-sending system calls have a *node* parameter that specifies the node to which the message is sent. You can use any valid node number, or the special value `-1` to send the message to all nodes in the application except the sending node itself.

Some message-receiving system calls have a *nodesel* parameter that specifies the node from which the message was sent. A *nodesel* parameter can be a valid node number (to receive only messages from that node), or the special value `-1` (to receive messages from any node). Message-receiving system calls that do not have a *nodesel* parameter always receive messages from any node.

The node numbers used in message-passing calls are always node numbers within the application, not physical slot numbers or node numbers within the partition in which the application is running. For example, if you run an application on 30 nodes of a 64-node partition by using the switch `-sz 30`, the node numbers within the application will always be 0 through 29. However, those nodes might not be nodes 0 through 29 of the partition. They might be nodes 0 through 29, or 10 through 39, or a completely arbitrary set of nodes.

## Process Types

Synopsis	Description
<code>myptype()</code>	Obtain the calling process's process type.
<code>setptype(<i>ptype</i>)</code>	Set the calling process's process type (only permitted if the process type is currently <b>INVALID_PTYPE</b> ).

A process's *process type*, or *ptype*, is an integer that distinguishes the process from other processes in the same application running on the same node. Process types are assigned by the user, and can be any integer from 0 to 2,147,483,647 ( $2^{31} - 1$ ) inclusive. A process can find out its process type by calling `myptype()`. A process cannot change its process type once it has been set to a valid value.

When you run an application that was linked with `-nx`, the system sets the process type of all processes in the application to the value you specify with the `-pt` switch on the application's command line (default 0).

All message-sending system calls have a *ptype* parameter that specifies the process type to which the message is sent. You must specify the process type; you cannot use -1.

Some message-receiving system calls have a *ptypesel* parameter that specifies the process type from which the message was sent. A *ptypesel* parameter can be a valid process type (to receive only messages from that process type), or the special value -1 (to receive messages from any process type). Message-receiving system calls that do not have a *ptypesel* parameter always receive messages from any process type.

Certain system calls that involve all the nodes in the application, called *global operations*, require that every node in the application has one process with the same process type. All these processes must call the global operation before the application can proceed.

Within a single application, multiple processes running on the same node must have different process types. However, processes on different nodes may (and usually do) have the same process type. Two processes running on a single node may have the same process type only if they belong to different applications.

### NOTE

The `-pt` switch (or, if not specified, the default process type of 0) applies only to the process type of the initial processes created by running the application.



If an application creates additional processes after it starts up, and no process type is specified for the new process, the new process's process type is set to the special value **INVALID\_PTYPE** (a negative constant defined in the header file *nx.h*). A process whose process type is **INVALID\_PTYPE** cannot send or receive messages. It must use the system call **setptype()** to set its process type to a valid value before it can send or receive any messages. (This is the only valid use of **setptype()**.)

The Paragon OSF/1 system calls that create node processes (**nx\_nfork()**, **nx\_load()**, and **nx\_loadve()**) have a *ptype* parameter that specifies the process type of the newly-created processes. However, the standard OSF/1 system call **fork()**, which creates a new process on the same node as the process that calls it, does not provide any way to specify the new process's process type. This means that the process type of a process created by **fork()** is set to **INVALID\_PTYPE**. The new process must call **setptype()** before it can send or receive messages. The specified process type must be different from the parent's, and different from the process type of any other process in the same application on the same node.

A process's process type is inherited across an **exec()**. This means that if you do a **fork()** followed by an **exec()**, you can call **setptype()** either before or after the **exec()**. However, the **setptype()** must follow the **fork()**.

Once a process has used a process type, that process type is associated with the process for the life of the application. No other process on the same node in the same application can ever use that process type, even if the original process terminates.

If a process has multiple pthreads, all the pthreads in the process have the same process type. See Chapter 6 for information on pthreads.

## Message Characteristics

Messages are characterized by a *length*, a *type*, and sometimes an *ID*. These characteristics are set when the message is sent, and do not change for the life of the message.

### Message Length

The *length* of a message is the number of bytes of information contained in the message. Messages in Paragon OSF/1 can be of any length.

All message-passing system calls have a *count* parameter that specifies the length of the message to be sent or received. The length you specify must be less than or equal to the size in bytes of the buffer used in the call. Message-sending calls read exactly that number of bytes from the buffer and send them as a message; message-receiving calls generate an error if a message is received that is larger than the specified length.

If you program in C, when you send a message you can use the **sizeof** operator to determine the size of your message in bytes. If you program in Fortran, you will need to add up the sizes of all the data elements within the message; see the *Paragon™ Fortran Compiler User's Guide* for information on the default size of each data type. If you pass named common blocks as messages, you may also have to include the space taken up by padding within the common block, as discussed under "Message Passing with Fortran Commons" on page 3-17.

You can also send and receive zero-length messages. This is useful if the message type is sufficient, and there is no need to supply any message content. For example, one process could tell another process to start or stop doing something by sending a zero-length message of type 1 to start, or a zero-length message of type 2 to stop.

## Message Type

The *type* of a message is an integer whose meaning is determined by the programmer.

All message-sending system calls have a *type* parameter that specifies the type of the message sent. You can use any integer from 0 to 999,999,999 (inclusive) as a message type.

All message-receiving system calls have a *typesel* parameter that specifies the type (or types) of messages the call will receive. A *typesel* parameter can be an integer from 0 to 999,999,999 (to receive only messages of the specified type) or the special value -1 (to receive messages of any type).

There are also special message types outside the range 0 to 999,999,999, called *force types* and *typesel masks*, that you can use. Sending with a force type sends a message that uses a limited flow control technique; receiving with a typesel mask receives messages of a selected set of types. See the *Paragon™ Fortran System Calls Reference Manual* or *Paragon™ C System Calls Reference Manual* for information on these special message types. Note, though, that in Paragon OSF/1 regular messages are just as fast as force type messages, so force types are not needed for performance.

## Message ID

The *ID* of a message is an identifier used to check for the completion of asynchronous messages. Synchronous messages do not have IDs.

When you send or receive a message with an *asynchronous* message-passing call (one that returns before the message is completely sent or received), the call returns an ID that you can use to check whether or not the send or receive is complete. See "Asynchronous Send and Receive" on page 3-10 for more information on message IDs.

## Message Order

Paragon OSF/1 guarantees that all messages will arrive in the same order they are sent. That is, if one message is sent from node A to node B, then a second message is sent from node A to node B, the second message will never arrive before the first.

Although the first message always *arrives* at the node first, you can elect to *receive* the second message—that is, to copy its contents into a buffer in user memory—before the first. You do this by specifying different message types in the send calls on node A, and specifying the second message's type in the first receive call on node B.

## Names of Send and Receive Calls

You can tell what each message-passing call does by examining its name.

The first character of the name indicates whether the call is synchronous, asynchronous, or handled:

- c** Synchronous (complete) call. These calls do not return until the message is complete. They are discussed under "Synchronous Send and Receive" on page 3-8.
- i** Asynchronous (incomplete) call. These calls return immediately, so your program can do other work while the message is processed. They are discussed under "Asynchronous Send and Receive" on page 3-10.
- h** Asynchronous with interrupt handler (handled) call. Like the **i...()** calls, the **h...()** calls return immediately. Unlike the **i...()** calls, **h...()** calls indicate that the message is complete by calling a user-supplied interrupt handler. They are discussed under "Treating a Message as an Interrupt" on page 3-18.

The initial **c**, **i**, or **h** is followed by a verb that indicates what the call does:

- send** Send a message.
- recv** Receive a message.
- sendrecv** Send a message and receive the reply.
- probe** Probe for a pending (not yet received) message.

Finally, the verb may be followed by an **x** to indicate that it is an "extended" version of the call (see "Treating a Message as an Interrupt" on page 3-18 and "Extended Receive and Probe" on page 3-24).

The synchronous calls with no additional functionality, such as `csend()`, are the easiest to understand and use. However, the asynchronous calls (such as `isend()`) and the calls with additional functionality (such as `crecvx()`) can offer dramatic improvements in performance when properly used.

## Synchronous Send and Receive

Synopsis	Description
<code>csend(type, buf, count, node, ptype)</code>	Send a message, waiting for completion.
<code>crecv(typesel, buf, count)</code>	Receive a message, waiting for completion.
<code>csendrecv(type, sbuf, scount, node, ptype, typesel, rbuf, rcount)</code>	Send a message and post a receive for the reply. Wait for completion.

The `c...()` message-passing calls perform synchronous sends and receives.

- A synchronous send means that the program executing the send waits until the send is complete. This waiting is referred to as *blocking*. Completing the send, however, does not guarantee that the message has been received. It only means that the message has left the sending process and that the buffer can be reused. You use `csend()` to perform a synchronous send.
- A synchronous receive means that the program executing the receive waits until the message arrives in the specified buffer. You use `crecv()` to perform a synchronous receive.
- A `csendrecv()` is like a `csend()` followed by a `crecv()`. It returns the length of the received message.

Here are two code fragments in C that perform a synchronous send and a synchronous receive.

- Node 1 sends a message of type 0 to the process with the same process type on node 0:

```
#include <nx.h>
#define MSG_TYPE 0
#define DEST_NODE 0
char send_buf[100];
.
.
.
csend(MSG_TYPE, send_buf,
      sizeof(send_buf), DEST_NODE, myptype());
```

- Node 0 receives the message:

```
#include <nx.h>
#define MSG_TYPE 0
char recv_buf[100];
.
.
.
crecv(MSG_TYPE, recv_buf, sizeof(recv_buf));
```

See “Extended Receive and Probe” on page 3-24 for information on a version of the `crecv()` call with additional functionality.

## Synchronous Send to Multiple Nodes

Synopsis	Description
<code>gsendx(type, buf, count, nodes, nodecount)</code>	Send a message to a list of nodes, waiting for completion.

The `gsendx()` call sends a message to multiple nodes. Specifically, it performs a synchronous send of the message specified by the *type*, *buf*, and *count* arguments to the process with the same process type as the caller on the nodes specified by the *nodes* argument. The *nodes* argument is an array of integers; the *nodecount* argument specifies the number of nodes in *nodes*.

For example, the following code fragment in Fortran sends the data in the array *x* to nodes 1 and 3:

```
integer*4 nodenums(2), x(10)
.
.
.
nodenums(1) = 1
nodenums(2) = 3
call gsendx(100, x, 10*4, nodenums, 2)
```

## Asynchronous Send and Receive

Synopsis	Description
<b>isend</b> ( <i>type, buf, count, node, ptype</i> )	Send a message without waiting for completion.
<b>irecv</b> ( <i>typesel, buf, count</i> )	Receive a message without waiting for completion.
<b>isendrecv</b> ( <i>type, sbuf, scount, node, ptype, typesel, rbuf, rcount</i> )	Send a message and post a receive for the reply without waiting for completion.
<b>msgdone</b> ( <i>mid</i> )	Determine whether a send or receive operation has completed.
<b>msgwait</b> ( <i>mid</i> )	Wait for completion of a send or receive operation.
<b>msgignore</b> ( <i>mid</i> )	Release a message ID as soon as a send or receive operation completes.

The **i...()** message-passing calls perform asynchronous sends and receives. The **msgdone()** and **msgwait()** calls are used with the **i...()** calls to determine when the message has completed; the **msgignore()** call is used to discard a message ID as soon as the message has completed.

Unlike a synchronous send or receive, an asynchronous send or receive does not block. It returns a unique message ID, which is not reused until released. You can use this ID to check for completion at a later time.

### NOTE

The number of message IDs is limited, so you *must* release each ID after you use it. See "Releasing Message IDs" on page 3-12 for information on releasing message IDs.

You use **isend()** to perform an asynchronous send, and **irecv()** to perform an asynchronous receive. An **isendrecv()** is like an **isend()** followed by an **irecv()**, except that it returns only one message ID (for the receive). Asynchronous sends can be used together with synchronous receives, and vice versa. For example, a message sent by **isend()** could be received by **crecv()**.

You must make sure that an asynchronous operation has completed before you change the contents of the send buffer or use the contents of the receive buffer. To check if an asynchronous operation has completed, use the `msgdone()` call. It returns 1 if an asynchronous call has completed and 0 otherwise. To block until an asynchronous operation has completed, use the `msgwait()` call. Both `msgdone()` and `msgwait()` take the message ID as an input parameter.

The message ID belonging to an asynchronous receive is distinct from the message ID belonging to any companion asynchronous send. For example, if node 0 sends a message with `isend()` and node 1 receives the message with `irecv()`, the `isend()` has a different message ID from the `irecv()`. When the `isend()` completes, this does *not* indicate that the corresponding `irecv()` has completed.

For example, assume that your application knows that it's going to need a message up ahead. So it posts an asynchronous receive with `irecv()`. It then does work that does not require the message, believing that by the time it needs the message, it will have arrived. When the program comes to where it needs the message, it issues a `msgwait()`. If the message has in fact arrived, the `msgwait()` returns immediately. Otherwise, it blocks until the message arrives. Here is a Fortran code fragment that implements this technique.

Node 1 does an asynchronous send:

```
include 'fnx.h'

integer result, msg_sid
integer MSG_TYPE, DEST_NODE
double precision send_buf(100)
parameter (MSG_TYPE = 1)
parameter (DEST_NODE = 0)
.
.
.
msg_sid = isend(MSG_TYPE, send_buf,
               100*8, DEST_NODE, myptype())
.
.
.
c   Free the asynchronous send ID
   call msgwait(msg_sid)
```

Node 0 does the asynchronous receive:

```

include 'fnx.h'

integer result, msg_rid
integer MSG_TYPE
double precision rec_buffer(100)
parameter (MSG_TYPE = 1)
.
.
.
c Post the receive
msg_rid = irecv(MSG_TYPE, rec_buffer, 100*8)
.
.
.
c Now you need the message.
c
c Free the asynchronous receive ID
call msgwait(msg_rid)

```

When the `msgwait()` returns, the message has been received. You may have blocked on the `msgwait()` if the message had not yet arrived. You may now assign another value to `msg_rid`.

See “Extended Receive and Probe” on page 3-24 for information on a version of the `irecv()` call with additional functionality.

## Releasing Message IDs

Because Paragon OSF/1 has a limited number of message IDs, you must release IDs that are no longer needed. There are four ways to release a message ID:

- You can call `msgwait()`.
- You can keep calling `msgdone()` until it returns 1.
- You can call `msgignore()`.

If you use `msgignore()`, it tells the system to release the message ID as soon as the corresponding send or receive has completed. Note, though, that this leaves you with no way to determine whether or not the message has completed. In this case, your application must have some other means of synchronization to prevent the send or receive buffer from being used before the message is complete.



## NOTE

Re-using a send or receive buffer before the message is complete can result in unexpected behavior. Do not use `msgignore()` unless you are certain this will not occur.

## Merging Message IDs

---

### Synopsis

`msgmerge(mid1, mid2)`

### Description

Merge two message IDs into a single ID that can be used to wait for completion of both operations.

---

The `msgmerge()` call gives you a way to merge two or more message IDs together. It takes two message IDs as parameters, and returns a message ID that does not complete until both the messages identified by the input message IDs have completed.

Once you have merged a message ID with `msgmerge()`, you should not use the input message IDs as arguments to `msgwait()`, `msgdone()`, `msgcancel()`, or `msgignore()`. The input message IDs are automatically released when the merged message IDs are waited for.

For example, the following C code fragment posts two `irecv()`s, one for a message of type 1 and the other for a message of type 2, and then waits until both have completed:

```
#include <nx.h>

int mid1, mid2, midg;
char buf1[10], buf2[10];
.
.
.
mid1 = irecv(1, buf1, 10);
mid2 = irecv(2, buf2, 10);

midg = msgmerge(mid1, mid2);

msgwait(midg);
```

Note that *mid1* and *mid2* are released by the `msgwait()` call on *midg*.

You can use a series of `msgmerge()` calls to merge multiple message IDs together. To help you do this, you can use the value -1 as one of the message IDs; `msgmerge()` returns the other message ID unchanged.

For example, the following Fortran code fragment uses a series of **isend()** calls to send the buffer *buf* as a message of type 1 to the process with the same process type on nodes 1 through 10, then waits for all of the **isend()**s to complete:

```

include 'fnx.h'

integer i, mid
integer buf(100)

mid = -1
i = 1

do while (i .le. 10)
    mid = msgmerge(mid, isend(1, buf, 400, i, myptype()))
    i = i + 1
end do

call msgwait(mid)

```

The message ID returned by each **isend()** call is merged together with the message IDs of the previous **isend()** calls into the merged message ID *mid* (the first message ID is merged with -1, yielding itself). Once all the **isend()**s have been posted, the program uses **msgwait()** on the merged message ID to wait for all of the **isend()**s to complete.

## Probing for Pending Messages

---

### Synopsis

**cprobe**(*typesel*)

**iprobe**(*typesel*)

### Description

Wait for a message of a selected type to arrive.

Determine whether a message of a selected type is pending.

---

When a message arrives for which no receive has been issued, it goes into a system buffer. It is referred to as a *pending message*: a message that is available for receipt, but not yet received. When you issue a receive for that message, the message is moved into the application's buffer (the buffer you specify in the **crecv()** or **irecv()** call). If a receive has already been issued when the message arrives, it goes directly into the application's buffer and bypasses the system buffer.

The **cprobe()** and **iprobe()** calls determine whether there is a message of a given type pending in the system buffer. You can use a message type from 0 to 999,999,999 to probe for a message of a specific type; the special value -1 to probe for a message of any type; or a *typesel mask* to probe for messages of a selected set of types (see the *Paragon™ Fortran System Calls Reference Manual* or *Paragon™ C System Calls Reference Manual* for information on *typesel* masks).

The **cprobe()** call is a blocking call. It takes a type selection parameter as input and returns when a message of the given type has arrived. The **iprobe()** call is similar to **cprobe()**, except that it is nonblocking. **iprobe()** returns 1 if the message is pending and 0 if it is not.

**cprobe()** and **iprobe()** are not the only calls that probe for messages. See "Extended Receive and Probe" on page 3-24 for information on message-probing calls with additional functionality.

## Getting Information About Pending or Received Messages

Synopsis	Description
<b>infocount()</b>	Return size in bytes of a pending or received message.
<b>infonode()</b>	Return node number of the node that sent a pending or received message.
<b>infoftype()</b>	Return process type of the process that sent a pending or received message.
<b>infotype()</b>	Return message type of a pending or received message.

The **info...()** calls return information about received or pending messages. You can obtain the size of the message, its type, and the node number and process type of the sending process.

The return value of the **info...()** calls is defined only in the following cases:

- After a **crecv()**, **cprobe()**, or **msgwait()**.
- After an **iprobe()** or **msgdone()** returns 1.

Note that you must issue the **info...()** call before you perform any other message-passing operation. Otherwise, you will get information about the most recently received or pending message instead.

For example, the following C code receives a message of any type, then uses `infotype()` to determine what type of message was actually received:

```
#include <nx.h>
#define BIGNUM 262144
long buf[BIGNUM], msg_type;
.
.
.
crecv(-1, buf, sizeof(buf));
msg_type = infotype();
```

Another example: the following C code blocks until any message arrives, then allocates a buffer just large enough to hold the message and receives it:

```
#include <nx.h>
char *buf;
long msg_type, msg_len;
.
.
.
cprobe(-1);
msg_type = infotype();
msg_len = infocount();
buf = (char *) calloc(msg_len, 1);
crecv(msg_type, buf, msg_len);
.
.
.
```

Between the `cprobe()` and the `crecv()`, the message is pending; it has arrived, but has not yet been received. Until the message is received, the contents of the message are not accessible to the program.

The `info...()` calls are subject to the following special conditions:

- The return value of the `info...()` calls is undefined after a `msgwait()` or `msgdone()` if the message ID in the `msgwait()` or `msgdone()` call is a “merged” message ID representing more than one message. See “Merging Message IDs” on page 3-13 for more information.
- The return value of the `info...()` calls is undefined after a `crecvx()`, `cprobex()`, or `iprobex()`, except if the last parameter is the special array `msginfo`. See “Extended Receive and Probe” on page 3-24 for more information.
- If you issue an `info...()` call before doing *any* message passing, the call returns -1.

The `info...()` calls are not the only way to get information about a received or pending message. See “Extended Receive and Probe” on page 3-24 for information on message-receiving and message-probing calls that also return information about the received or pending message.

## Message Passing with Fortran Commons

Fortran users often use common blocks to send messages that contain data elements of different types. For example, consider the named common containing a double precision number and an integer. It is good Fortran practice to put the largest data element first in the common list, as follows:

```
integer i
double precision d
common/msg/ d, i
```

To send this common block, specify the name of the first common element as the buffer and the length of the entire common as the length. For example, to send the common block named `msg`, send the variable `d` with a length of 12 bytes (8 for the double precision variable plus 4 for the integer variable). The following `csend()` call sends `msg` to process `ptype` on node `node`.

```
call csend(MSGTYPE, d, 12, node, ptype)
```

If you put smaller data elements before larger data elements in common blocks, the compiler may have to insert padding, or “holes,” between the elements of the common block to preserve data alignment. For example, if you define the common block named `pmsg` as follows, the compiler will place an invisible 4-byte pad between the end of `i` and the beginning of `d` to properly align `d` on an 8-byte boundary:

```
integer i
double precision d
common/pmsg/ i, d
```

This padding has two effects:

- If you send this common block as a message, you must include the padding in the length of the message. For example, even though `pmsg` contains the same two variables as `msg`, `pmsg` is 4 bytes longer than `msg` because of the padding between `i` and `d`. To send `pmsg` to process `ptype` on node `node`, you would use the following call:

```
call csend(MSGTYPE, i, 16, node, ptype)
```

- If another routine uses a different view of the same common block, you may have to add additional variables to the other routine's declaration of the common block to take this padding into account. For example, if another routine wants to view *d* in *pmsg* as an array of two integers, it must declare *pmsg* as follows:

```
integer i, ipad, id(2)
common/pmsg/ i, ipad, id(2)
```

The variable *ipad* corresponds to the 4-byte pad in the original routine's declaration of *pmsg*. Without this variable, the position of *id* would not correspond to the position of *d* in the original common block. This variable is necessary if *pmsg* is shared between these two routines, whether or not the two routines run on different nodes.

When possible, you should define common blocks with the largest data element first, to avoid padding completely. You should also use the `%LOC` function to determine the size of a common block and avoid specifying its size with a hard-coded constant.

## Treating a Message as an Interrupt

Synopsis	Description
<code>hsend(type, buf, count, node, ptype, handler)</code>	Send a message and set up a handler procedure to be called when the send completes.
<code>hrecv(typesel, buf, count, handler)</code>	Receive a message and set up a handler procedure to be called when the receive completes.
<code>hsendrecv(type, sbuf, scount, node, ptype, typesel, rbuf, rcount, handler)</code>	Send a message and post a receive for the reply. Set up a handler procedure to be called when the reply arrives.

The `h...()` message-passing calls perform asynchronous sends and receives. However, unlike the `i...()` calls, the `h...()` calls let you establish a user-provided interrupt handler, which is called when the send or receive is complete.

The `h...()` receive calls let you treat incoming messages as interrupts. For example, consider a program that performs some action based on the type of a received message. One way to implement this program is to block the program at a `crecv()` for messages of all types and then take appropriate action based on the value returned by `infotype()`.

Another way is to issue a number of `hrecv()` calls. Each call attaches a function to a particular message type or set of types. Your program does not block. You can continue with other work; but when the appropriate message comes, the attached function is called to take care of the message. (The message is stored in the receive buffer before the function is called.)

The handler function you define must be written in C and must have four arguments of type **long**. These arguments are passed the following values when the function is called:

1. Type of the message (as returned by **infotype()**).
2. Length of the message in bytes (as returned by **infocount()**).
3. Node number of the process that sent the message (as returned by **infonode()**).
4. Process type of the process that sent the message (as returned by **infoptype()**).

For example, here's a C code fragment that attaches the functions *funct0()*, *funct1()*, and *funct2()* to message types 0, 1, and 2, respectively. The message types that have handlers are referred to as *handled types*.

```
#include <nx.h>

char buf0[100], buf1[100], buf2[100];
void funct0(), funct1(), funct2();

hrecv(0, buf0, sizeof(buf0), funct0);
hrecv(1, buf1, sizeof(buf1), funct1);
hrecv(2, buf2, sizeof(buf2), funct2);
    .
    . /* Now perform other work. No blocking happens. */
    .
```

The declaration of **funct1()** looks like this (the other functions are similar):

```
void funct1(long type, long count, long node, long ptype)
{
    .
    .
    .
}
```

When a message of type 1 arrives, the message is stored in the buffer specified in the **hrecv()** call (in this case, *buf1*), then **funct1()** is called with the type and length of the message and the node number and process type of the sender as arguments. **funct1()** and the main program then run concurrently until **funct1()** returns. (In previous releases of Paragon OSF/1, the main program was interrupted and did not run at all until **funct1()** returned.)

## CAUTION

The handler runs in the same memory space as the main program (but they have separate stacks).

Because of this, parts of the main program may have to be protected from being executed at the same time as the handler; see “Preventing Interrupts” on page 3-22 for information on using **masktrap()** to do this.

## NOTE

Once you have established a handler for a message type, do not attempt to receive a message of that type with a **crecv...()** or **irecv...()** call.

**hsend()** operates the same as **hrecv()**, except that the handler is invoked when the send completes. (Note that completion of the send does not mean that the message has been received, only that the message has been sent and the send buffer can be reused.) **hsendrecv()** is like an **isend()** followed by an **hrecv()**, with the message ID of the **isend()** automatically released by **msgignore()**.

See “Extended Receive and Probe” on page 3-24 for information on a version of the **hrecv()** call with additional functionality.

## Passing Information to the Handler

---

### Synopsis

**hsendx**(*type, buf, count, node, ptype, xhandler, hparam*)

### Description

Send a message and set up an extended handler procedure to be called with the value *hparam* when the send completes. Allows handler sharing.

---

**hsendx()** is identical to **hsend()** except that it has an additional parameter, *hparam*, which is passed to the handler when it is called. The declaration of a handler for **hsendx()** looks like this:

```
void xhandler(long type, long count, long node, long ptype,
              long hparam)
{
    .
    .
    .
}
```



You can use the *hparam* parameter to write handlers that are shared among several `hsendx()` calls, each of which uses a different value of *hparam* to identify itself. For example, here is a C program fragment that sends two messages of type 0 to the process with process type 2 on node 1, then uses an `hsendx()` handler to free each message buffer as soon as the message send completes:

```
#include <nx.h>
#include <malloc.h>

#define NBUFS 2
#define BUFFER_SIZE 10000

char *buf[NBUFS]; /* array of pointers to char */

void freemem(long type, long count, long node, long ptype,
             long hparam)
{
    if( (hparam >= 0) && (hparam < NBUFS) ) {
        free(buf[hparam]);
    } else {
        printf("freemem(): invalid value: %d\n", hparam);
    }
}

main(int argc, char **argv)
{
    /* allocate two buffers with malloc() */
    buf[0] = malloc(BUFFER_SIZE);
    buf[1] = malloc(BUFFER_SIZE);
    .
    .   /* put data into the buffers */
    .
    /* send them */
    hsendx(0, buf[0], BUFFER_SIZE, 1, myptype(), freemem, 0);
    hsendx(0, buf[1], BUFFER_SIZE, 1, myptype(), freemem, 1);
    .
    .   /* Now perform other work */
    .
}
```

Note that you must take care that this handler is not called while the program is in the middle of a call to `malloc()` or `free()`. If the handler attempts to free memory while another part of the program is allocating or freeing memory, `malloc()`'s internal memory structures could become corrupted. You can prevent this by using the `masktrap()` call, described in the following section, to protect each `malloc()` and `free()` call elsewhere in the program that could be interrupted by this handler.

## Preventing Interrupts

---

### Synopsis

**masktrap**(*state*)

### Description

Enable or disable interrupts for message handlers. Required to prevent corruption of global variables.

---

If you have one or more handlers set up and you have some critical code that you do not want interrupted, use the **masktrap**() call. A **masktrap**(1) prevents any handler from running; a **masktrap**(0) re-enables handlers. Any pending interrupts are honored when the mask is removed. **masktrap**() returns the previous masking state (1 or 0). For example:

```
hrecv(6,buf,sizeof(buf),myhandler);
    .
    . /* this code can be interrupted */
    . /* by a message of type 6 */
    .
oldmask = masktrap(1);
    .
    . /* critical code that must not be interrupted */
    .
masktrap(oldmask);
    .
    . /* this code can be interrupted again */
    .
```

Note the use of the variable *oldmask* to save the value of the previous masking state before the call to **masktrap**(1). This means that if the mask were already set before this call (for example, if this code is in a subroutine that could be called when the mask is already set), the following **masktrap**(*oldmask*) would not unset it.

## CAUTION

You must use **masktrap**() around any code in the main program that could interfere with calls in the handler.

For example, if the handler performs any I/O, you must put **masktrap**() calls around any I/O calls (such as **printf**()) in the main program that could be called while the handler is active. If you don't do this, you could find characters from the handler's output interleaved with characters from the main program's output.

Sometimes, though, it's not as obvious which calls could interfere with each other. For example, any two library calls that could allocate or free memory could cause the memory subsystem to become confused if they were called at the same time. To be on the safe side, keep the handler as simple as possible and use **masktrap()** to protect *all* library calls in the rest of the program that could call the same subsystems as the calls in the handler while the handler is active.

These calls to **masktrap()** are necessary because, when the handler is active, the handler and the main program share the same memory space and can change each other's global variables. This could cause any *non-reentrant* function to fail if it is called by both at the same time.

If the handler performs any message passing, any **info...()** call in the main program must be within the same set of **masktrap()** calls as the message-receiving call to which it applies. Otherwise, the **info...()** call could reflect the value of a message received within the handler.

## Extended Receive and Probe

---

### Synopsis

**crecvx**(*typesel, buf, count, nodesel, ptypesel, info*)

**irecvx**(*typesel, buf, count, nodesel, ptypesel, info*)

**hrcvx**(*typesel, buf, count, nodesel, ptypesel, xhandler, hparam*)

**cprobex**(*typesel, nodesel, ptypesel, info*)

**iprobex**(*typesel, nodesel, ptypesel, info*)

### Description

Receive a message of a specified type from a specified sending node and process type, together with information about the message. Wait for completion.

Receive a message of a specified type from a specified sending node and process type, together with information about the message. Do not wait for completion.

Receive a message of a specified type from a specified sending node and process type. Set up an extended handler procedure to be called with information about the message and the value *hparam* when the receive completes.

Wait for a message of a specified type from a specified sending node and process type. Return information about the message.

Determine whether a message of a specified type from a specified sending node and process type is pending. If it is, return information about the message.

---

The extended receive and probe calls, **crecvx()**, **irecvx()**, **hrcvx()**, **cprobex()**, and **iprobex()**, can be used to receive or probe for a message from a particular node or a particular process type, and return information about the message along with the message (instead of using the **info...()** calls).

**crecvx()**, **irecvx()**, **cprobex()**, and **iprobex()** are like **crecv()**, **irecv()**, **cprobe()**, and **iprobe()**, except that they have the following additional parameters:

- |                 |   |
|-----------------|---|
| <i>nodesel</i>  | Specifies the node that sent the message, or -1 for any node.                 |
| <i>ptypesel</i> | Specifies the process type that sent the message, or -1 for any process type. |

*info* An array of eight long integers that receives information about the specified message. The following information is stored into the first four elements of this array:

- Type of the message (as returned by **infoftype()**).
- Length of the message in bytes (as returned by **infocount()**).
- Node number of the process that sent the message (as returned by **infonode()**).
- Process type of the process that sent the message (as returned by **infoftype()**).

The remaining four elements of the array are reserved.

**hrecvx()** is like **hrecv()**, except that it has the same *nodesel* and *ptypesel* parameters as the other *...x()* calls and the same *hparam* parameter as the **hsendx()** call. **hrecvx()** does not have an *info* parameter, because the corresponding information is passed to the handler when it is called.

The *info* parameter of **crecvx()**, **irecvx()**, **cprobex()**, and **iprobex()** must be specified and must not be zero or null. If you do not want this information, or you want it to be available to the **info...()** calls, specify the special array *msginfo*, defined in *nx.h* or *fnx.h*. The array *msginfo* is used by the non-*x* versions of these calls, and the **info...()** calls get their information from *msginfo*. This is why you cannot use the **info...()** calls after **crecvx()**, **cprobex()**, or **iprobex()** unless you specify *msginfo* as the last parameter of the extended receive or probe call.

The *info* parameter of **irecvx()** does not contain valid data until the message is received (as determined by **msgdone()** or **msgwait()**). The *info* parameter of **iprobex()** does not contain valid data unless the **iprobex()** returns 1.



## Global Operations

---

### Synopsis

**gcol**(*x, xlen, y, ylen, ncnt*)  
**gcolx**(*x, xlens, y*)  
**gdhigh**(*x, n, work*)  
**gdlow**(*x, n, work*)  
**gdprod**(*x, n, work*)  
**gdsum**(*x, n, work*)  
**giand**(*x, n, work*)  
**gihigh**(*x, n, work*)  
**gilow**(*x, n, work*)  
**gior**(*x, n, work*)  
**giprod**(*x, n, work*)  
**gisum**(*x, n, work*)  
**gland**(*x, n, work*)  
**glor**(*x, n, work*)  
**gopf**(*x, xlen, work, function*)  
**gshigh**(*x, n, work*)  
**gslow**(*x, n, work*)  
**gsprod**(*x, n, work*)  
**gssum**(*x, n, work*)  
**gsync**()

### Description

Concatenation.  
 Concatenation for contributions of known length.  
 Vector double precision MAX.  
 Vector double precision MIN.  
 Vector double precision MULTIPLY.  
 Vector double precision SUM.  
 Vector integer bitwise AND.  
 Vector integer MAX.  
 Vector integer MIN.  
 Vector integer bitwise OR.  
 Vector integer MULTIPLY.  
 Vector integer SUM.  
 Vector logical AND.  
 Vector logical inclusive OR.  
 Arbitrary commutative function.  
 Vector real MAX.  
 Vector real MIN.  
 Vector real MULTIPLY.  
 Vector real SUM.  
 Global synchronization.

---

The **g...**() calls perform operations that use data from every node in the application. In general, when you make one of these calls, each node contributes a piece of data to the operation, the operation is performed on the whole collection of data, and then the result of the operation is returned to each node.

These operations are *synchronizing calls*: if any node in an application makes one of these calls, it blocks until every node in the application has made the same call. (In the simplest case, **gsync**(), this synchronization is the *only* operation performed by the call.) One process on each node in the application must make the call, and all the processes that make the call must have the same process type.

The global operations are implemented using dynamic algorithm selection for maximum performance. The system considers several ways of exchanging the needed information between the nodes, and selects the one that minimizes the time required to perform the global operation given the size and shape of the application.

**ioctl()** fills in the elements of this structure with information about the device. The value of the *mt\_type* field is always **0x0C** (indicating a generic SCSI device). The values of the *mt\_dsreg* and *mt\_erreg* fields are device-dependent.

For example, the following C program prints the status of the device connected to */dev/io0/rmt6*:

```
#include <fcntl.h>
#include <errno.h>
#include <sys/mtio.h>

main() {
    int fd;
    struct mtget s;

    fd = open("/dev/io0/rmt6", O_RDONLY, 0666);
    if(fd == -1) {
        perror("opening /dev/io0/rmt6");
        exit(1);
    }

    if (ioctl(fd, MTIOCGET, &s) == -1) {
        perror("getting status of tape");
        exit(2);
    }

    printf("mt_type   = 0x%x\n", s.mt_type);
    printf("mt_dsreg  = 0x%x\n", s.mt_dsreg);
    printf("mt_erreg  = 0x%x\n", s.mt_erreg);
    printf("mt_resid  = 0x%x\n", s.mt_resid);
}
```



# Using Other Paragon™ OSF/1 System Calls



4

## Introduction

Paragon OSF/1 system calls are available to all programs running on the Intel supercomputer. These system calls provide a variety of specialized functions that let processes running on different nodes share data and coordinate their activities.

This chapter introduces the Paragon OSF/1 system calls that perform general services other than message passing. It includes the following sections, each of which describes a group of related calls:

- Managing applications.
- Managing partitions.
- Listing unusable nodes.
- Handling errors.
- Controlling floating-point behavior.
- Miscellaneous calls.
- iPSC® system and Touchstone DELTA system compatibility calls.

Within each section, the calls are discussed in order of increasing complexity. That is, the “base” calls are discussed first, and the “extended” calls are discussed later.

Each section includes numerous examples in both C and Fortran. A call description at the beginning of each section or subsection gives a language-independent synopsis (call name, parameter names, and brief description) of each call discussed in that section. Differences between C and Fortran are noted where applicable. See Appendix A for information on call and parameter types; see the *Paragon™ C System Calls Reference Manual* or the *Paragon™ Fortran System Calls Reference Manual* for complete information on each call.

This chapter does not describe all the Paragon OSF/1 system calls. For information about system calls that perform message passing, see Chapter 3. For information about the calls used with the Parallel File System, see Chapter 5. For information about the calls used with graphical interfaces, such as DGL and the X Window System, see the *Paragon™ Graphics Libraries User's Guide*. For information about the system calls that require root privileges, see the *Paragon™ System Administrator's Guide*.

Paragon OSF/1 programs written in C can also issue OSF/1 system calls. The Paragon OSF/1 operating system is a complete OSF/1 system and fully supports all the standard OSF/1 system calls. See the *OSF/1 Programmer's Reference* for information on these calls.

Paragon OSF/1 programs written in Fortran cannot make OSF/1 system calls directly, but the Fortran runtime library includes a number of system interface routines. These routines make a number of OSF/1 system calls available to Fortran programs. See the *Paragon™ Fortran Compiler User's Guide* for information on these routines.

## Managing Applications

Paragon OSF/1 provides system calls that let you create parallel applications, control their execution, and get information about them. See “Running Applications” on page 2-11 and “Managing Running Applications” on page 2-23 for introductory information on applications.

## Controlling Application Execution with System Calls

Synopsis	Description
<code>nx_initve(partition, size, account, argc, argv)</code>	Create a new application.
<code>nx_initve_rect(partition, anchor, rows, cols, account, argc, argv)</code>	Create a new application with a rectangular shape.
<code>nx_pri(pgroup, priority)</code>	Set the priority of an application.
<code>nx_nfork(node_list, numnodes, ptype, pid_list)</code>	Copy the current process onto some or all nodes of an application.
<code>nx_load(node_list, numnodes, ptype, pid_list, pathname)</code>	Execute a stored program on some or all nodes of an application.
<code>nx_loadve(node_list, numnodes, ptype, pid_list, pathname, argv, envp)</code>	Execute a stored program on some or all nodes of an application, with specified argument list and environment.
<code>nx_waitall()</code>	Wait for all application processes.

The simplest way to control the way an application executes is to use the command-line switch **-nx** when you link the application. (See “Compiling and Linking Applications” on page 2-5 for more information on the **-nx** switch.) When you execute a program that was linked with **-nx**, the program is automatically copied onto the specified number of nodes in the specified partition, runs, and then when all the nodes have finished you get your prompt back.

The code linked in by **-nx** reads the command line and environment variables, then performs the following actions for you (see “Controlling the Application’s Execution Characteristics” on page 2-13 for more information):

- Creates a new, empty application in the partition specified by the **-pn** switch and on the nodes of that partition specified by the **-sz** or **-nd** switch. If **-pn** is not used, the partition is specified by `$NX_DFLT_PART`, or `.compute` if `$NX_DFLT_PART` is not set. If neither **-sz** nor **-nd** is used, the number of nodes is specified by `$NX_DFLT_SIZE`, or all nodes of the partition if `$NX_DFLT_SIZE` is not set.
- Sets the application’s priority to the value specified by **-pri** (default 5).
- Loads and starts your program(s) on the nodes specified by **-on** (default all nodes of the application) with the process type specified by **-pt** (default 0).

The `nx_...()` system calls perform the same actions as those of the code linked in by **-nx**, but under program control instead of command-line control. Using these calls is more complicated than using **-nx**, but gives your program more flexibility and control.

## NOTE

If you use `nx_initve()` or `nx_initve_rect()`, do *not* link the program with the `-nx` switch. Use the switch `-lnx` instead.

The switch `-lnx` links in the library `libnx.a`, which contains all the calls discussed in this manual, but does not link in the automatic start-up code linked in by `-nx`.

## Creating an Application with `nx_initve()`

`nx_initve()` creates a new, empty application. The process that calls `nx_initve()` becomes the new application's *controlling process*; see "The Controlling Process" on page 4-21 for information on what this means.

The partition and size of the new application can be specified by parameters or by the command line; the priority and *mp\_switches* are specified by the command line. If command-line switches are not used or the command line is ignored by specifying zero for *argc*, the application's execution characteristics default as discussed under "Controlling the Application's Execution Characteristics" on page 2-13 and the *mp\_switches* default as discussed under "Message-Passing Configuration Switches" on page 8-18.

`nx_initve()` just allocates the specified number of nodes from the partition; it does not start any processes. (This allocation may or may not be exclusive, depending on the characteristics of the partition.) You must call `nx_nfork()`, `nx_load()`, or `nx_loadve()` to start processes in the new application. The nodes allocated to the application are automatically deallocated when all the processes in the application have terminated.

Another effect of `nx_initve()` is to make sure that the calling process is a *process group leader*. If the calling process is not already a process group leader, `nx_initve()` creates a new process group, removes the calling process from its current process group, and makes the calling process the new process group's leader. If you're not familiar with these terms, see "Process Groups" on page 4-22.

Finally, `nx_initve()` also initializes the data structures required by all the other calls described in this manual. In an application linked with `-nx`, the code linked in by `-nx` performs this initialization before the application starts up, so you can use these other calls anywhere in the application. In an application linked with `-lnx`, however, you must call `nx_initve()` before you can use *any* of the other calls described in this manual. If called before `nx_initve()`, these other calls will fail; the way they fail depends on the call (as described under "Handling Errors" on page 4-42). For example, if you call `csend()` before calling `nx_initve()`, the `csend()` prints an error message and terminates the calling process.

The parameters of `nx_initve()` have the following meanings:

*partition*      The relative or absolute partition pathname of the partition to run the application in, or a null string (" " or NULL in C, " " in Fortran) to use the default partition (the partition specified by `$NX_DFLT_PART`, or `.compute` if `$NX_DFLT_PART` is not set). The specified partition must exist and must give execute permission to the calling process.

If the user specifies a partition with the `-pn` switch on the command line, it overrides the value of the *partition* parameter (unless you set the *argc* parameter to 0, as described later in this section).

See "Partition Pathnames" on page 2-28 for more information on partition pathnames; see "Owner, Group, and Protection Modes" on page 2-32 for more information on partition permissions.

*size*            The size of the application (number of nodes to run the application on), or 0 to use the default size (the size specified by `$NX_DFLT_SIZE`, or all nodes of the partition if `$NX_DFLT_SIZE` is not set).

`nx_initve()` attempts to allocate a square group of nodes if it can. If this is not possible, it attempts to allocate a rectangular group of nodes that is either twice as wide as it is high or twice as high as it is wide. If this is not possible, it allocates any available nodes. In this case, nodes allocated to the application may not be contiguous.

If the user specifies the `-sz` or `-nd` switch on the command line, it overrides the value of the *size* parameter (unless you set the *argc* parameter to 0, as described later in this section).

*account*        In the future, this parameter will be used for accounting information. For now, it must be a null string (" " or NULL in C, " " in Fortran).

*argc*            In C, a pointer to an integer whose value is the number of arguments on the command line (counting the application name). If the value of this integer is 0, the command line is ignored. You can use a pointer to the *argc* parameter of `main()`, or you can construct the command line yourself.

In Fortran, this parameter is any nonzero value to search the command line, or 0 to ignore the command line.

**argv**

In C, a pointer to the command-line arguments, which may include arguments that specify application characteristics. You can use the *argv* parameter of **main()**, or you can construct the command line yourself.

In Fortran, **nx\_initve()** gets the command line directly from the system, because Fortran programs don't have an *argv* parameter. This parameter is ignored; it should always be 0.

In either language, if any of the command-line arguments **-sz size**, **-sz hXw**, **-nd hXw:n**, **-pri priority**, **-pn partition**, **-pkt packet\_size**, **-mbf memory\_buffer**, **-mex memory\_export**, **-mea memory\_each**, **-sth send\_threshold**, **-sct send\_count**, **-gth give\_threshold**, or **-plk** is found in the command line:

- The appropriate application characteristic is set as specified by the argument.
- The argument is removed from *argv*.
- The variable pointed to by *argc* is decremented appropriately.

Any remaining arguments are moved to the beginning of *argv* for your program's use.

Note that the arguments **-pt type**, **-on nodelist**, and **\; application** are *not* recognized by **nx\_initve()**. If you want your application to have the same user interface as an application linked with **-nx**, you must examine the argument list for these arguments and pass the appropriate values to **nx\_load()** or **nx\_loadve()** yourself.

**nx\_initve()** returns the number of nodes in the application, or -1 if any error occurs.

For example, the following C call creates an application whose characteristics (partition, number of nodes, and so on) are determined by the user through command-line switches. If the user runs this program with no command-line switches, it runs on the user's default number of nodes in the user's default partition.

```
#include <nx.h>

main(int argc, char *argv[]) {
    int n;

    n = nx_initve("", 0, "", &argc, argv);
}
```

After this call, the variable *n* contains the number of nodes in the new application, or -1 if any error occurred. The variable *argc* contains the count of arguments that were not recognized and removed by **nx\_initve()**, and the array *argv* contains pointers to those arguments.

The following Fortran call creates an application on 50 nodes of the partition *mypart*, ignoring any command-line switches provided by the user:

```
include 'fnx.h'
integer n

n = nx_initve("mypart", 50, "", 0, 0)
```

After this call, the variable *n* contains the number of nodes in the new application, or -1 if any error occurred.

The following restrictions apply to `nx_initve()`:

- A single process cannot call `nx_initve()` more than once.
- An application that calls `nx_initve()` cannot be linked with `-nx`. You must use `-lnx` instead.
- If your application uses any signal handlers, you must set them up *after* the call to `nx_initve()`. See `signal()` in the *OSF/1 Programmer's Reference* for more information on signal handlers.

The reason you cannot use `-nx` when you link an application that calls `nx_initve()` is that the code linked in by `-nx` calls `nx_initve()` itself, and `nx_initve()` can only be called once in an application. If you do use `-nx` when you link, your application's call to `nx_initve()` (actually the second call to `nx_initve()`) fails and returns -1.

## Creating a Rectangular Application with `nx_initve_rect()`

`nx_initve_rect()` works exactly like `nx_initve()` except that it requires that the nodes allocated to the application form a rectangle with a particular height and width. Optionally, it can also specify the rectangle's location within the partition. The parameters of `nx_initve_rect()` are the same as those of `nx_initve()`, except that instead of the *size* parameter it has the following three parameters:

<i>anchor</i>	The node number within the partition for the upper left corner of the rectangle, or -1 to allow the rectangle to be placed anywhere in the partition it will fit.
<i>rows</i>	The height of the rectangle.
<i>cols</i>	The width of the rectangle.

If the specified rectangle of nodes is not available, the `nx_initve_rect()` call fails and returns -1 (even if the equivalent number of nodes is available with a non-rectangular shape).

## NOTE

All the restrictions and cautions in this manual that refer to `nx_initve()` also apply to `nx_initve_rect()`.

If the user specifies a size or shape with the `-sz` or `-nd` switch on the command line, it overrides the values of these three parameters (unless you set the `argc` parameter to 0). `nx_initve_rect()` never uses the environment variable `$NX_DFLT_SIZE`.

For example, the following Fortran call creates an application 8 nodes high and 8 nodes wide (unless otherwise specified by command-line switches) anywhere it will fit in the user's default partition:

```
include 'fnx.h'
integer n

n = nx_initve_rect("", -1, 8, 8, "", 0, 0)
```

The following C call creates an application 10 nodes high and 20 nodes wide whose upper left corner is node 0 (the upper left corner of the partition) in the partition *mypart*, ignoring any command-line switches provided by the user:

```
#include <nx.h>
int n, i;

i = 0;
n = nx_initve_rect("mypart", 0, 10, 20, "", &i, NULL);
```

After either of these calls, the variable *n* contains the number of nodes in the new application, or -1 if any error occurred.

Note that `nx_initve_rect()` will fail if the exact specified rectangle is not available. If you just want to find out the application's size and shape, rather than mandating a particular size and shape, you can use an ordinary `nx_initve()`, followed by a call to `nx_app_rect()` (discussed under "Finding an Application's Shape with `nx_app_rect()`" on page 4-16) to determine the height and width assigned by `nx_initve()`.



## Setting an Application's Priority with `nx_pri()`

`nx_pri()` sets the specified application's priority to the specified value. If you don't call `nx_pri()` and the user doesn't use the `-pri` switch, the default priority is 5. The parameters of `nx_pri()` have the following meanings:

<i>pgroup</i>	The <i>process group ID</i> of the application (see "Process Groups" on page 4-22 for more information), or 0 to specify the application of the calling process. If the specified process group ID is not the process group ID of the calling process, the calling process's user ID must either be <i>root</i> or the same user ID as the specified application.
<i>priority</i>	The new priority, an integer from 0 to 10 inclusive. 0 is the lowest priority, 10 is the highest.

`nx_pri()` returns 0, or -1 if any error occurs.

For example, the following Fortran call sets the priority of the calling application to 7:

```
include 'fnx.h'
integer n

n = nx_pri(0, 7)
```

The following C call sets the priority of the application with process group ID 738423 to 0:

```
#include <nx.h>
int n;

n = nx_pri(738423, 0);
```

In each of these examples, the variable *n* is assigned 0, or -1 if any error occurred.

## Copying a Process onto the Nodes with `nx_nfork()`

`nx_nfork()` copies the process that calls it onto the specified set of nodes with the specified process type. It creates one *child process* on each specified node. `nx_nfork()` is similar to the standard OSF/1 call `fork()` except that it can fork processes onto multiple nodes and specifies the process type for the child processes. The parameters of `nx_nfork()` have the following meanings:

<i>node_list</i>	An array of integers, each of which specifies a node number within the application (no node number may appear more than once in this array). The calling process is copied onto each of the specified nodes.
<i>numnodes</i>	The number of node numbers in <i>node_list</i> , or -1 to use all the nodes in the application (in which case <i>node_list</i> is ignored).
<i>ptype</i>	The process type for each child process.
<i>pid_list</i>	An array of integers, into which are stored the OSF/1 process identifiers (PIDs) of the child processes. See "Using PIDs" on page 4-14 for more information.

`nx_nfork()` returns the number of child processes created to the parent process and 0 to each child process, or -1 if any error occurs.

For example, the following C calls create an application whose characteristics are specified by the user, then copy the calling process onto all nodes of the application. The process type of each child process is set to 0.

```
#include <nx.h>
#include <sys/types.h>

main(int argc, char *argv[]) {
    int n;
    pid_t pids[2000];

    n = nx_initve("", 0, "", &argc, argv);
    n = nx_nfork(NULL, -1, 0, pids);
}
```

Note that the *node\_list* argument is ignored when the *numnodes* argument is -1, so you can specify a **NULL** pointer in this case (in Fortran, you can use the value 0). After the call to `nx_nfork()`, the variable *n* contains the number of child processes created, or -1 if any error occurred; the first *n* elements of the array *pids* contains the PIDs of the child processes. If more than 2000 child processes are created, unexpected results will occur.

The following Fortran calls create an application with 100 nodes and copy the calling process onto the first 50 nodes of the application (nodes 0 through 49). The process type of each child process is set to 0.

```

      include 'fnx.h'
      integer n
      integer nodes(50), pids(50)

      n = nx_initve("mypart", 100, "", 0, 0)

      do 2, i = 1, 50
         nodes(i) = i - 1
2      continue

      n = nx_nfork(nodes, 50, 0, pids)

```

After the call to **nx\_nfork()**, the variable *n* contains 50, or -1 if any error occurred; the array *pids* contains the PIDs of the child processes.

## Loading a Program onto the Nodes with **nx\_load()**

**nx\_load()** executes the specified file on the specified set of nodes with the specified process type. Like **nx\_nfork()**, **nx\_load()** creates one child process on each specified node. The parameters of **nx\_load()** have the following meanings:

<i>node_list</i>	An array of integers, each of which specifies a node number within the application (no node number may appear more than once in this array). The specified file is loaded onto each of the specified nodes.
<i>numnodes</i>	The number of node numbers in <i>node_list</i> , or -1 to use all the nodes in the application (in which case <i>node_list</i> is ignored).
<i>ptype</i>	The process type for each child process.
<i>pid_list</i>	An array of integers, into which are stored the OSF/1 process identifiers (PIDs) of the child processes. See "Using PIDs" on page 4-14 for more information.
<i>pathname</i>	The relative or absolute pathname of the file to load.

**nx\_load()** returns the number of child processes created, or -1 if any error occurs.

For example, the following Fortran calls create an application whose characteristics are specified by the user, then load and start the program *myprog* on all nodes of the application. The process type of each child process is set to 0.

```
include 'fnx.h'
integer n
integer pids(2000)

n = nx_initve("", 0, "", 1, 0)
n = nx_load(0, -1, 0, pids, "myprog")
```

After the call to **nx\_load()**, the variable *n* contains the number of child processes created, or -1 if any error occurred; the first *n* elements of the array *pids* contains the PIDs of the child processes. If more than 2000 child processes are created, unexpected results will occur.

The following C calls create an application with 10 nodes in the partition *mypart*, then load and start the program *../bin/myprog* on nodes 1, 5, and 7 of the application. The process type of each child process is set to 1.

```
#include <nx.h>
#include <sys/types.h>
int n, i;
int nodes[3];
pid_t pids[3];

i = 0;
n = nx_initve("mypart", 10, "", &i, NULL);

nodes[0] = 1;
nodes[1] = 5;
nodes[2] = 7;

n = nx_load(nodes, 3, 1, pids, "../bin/myprog");
```

After the call to **nx\_load()**, the variable *n* contains 3, or -1 if any error occurred; the array *pids* contains the PIDs of the child processes.

## Loading a Program onto the Nodes with `nx_loadve()`

`nx_loadve()` is just like `nx_load()` except that it also lets you specify the argument list and environment variables for the new processes (in C). `nx_loadve()` has the following additional parameters:

<i>argv</i>	In C, this parameter contains the command line for the child process (you can use the <i>argv</i> parameter of <code>main()</code> or construct the command line yourself).
<i>envp</i>	In C, this parameter contains the environment variables for the child process (you can use the <i>envp</i> parameter of <code>main()</code> or construct the environment yourself).

In Fortran, you must specify the value 0 for the *argv* and *envp* parameters (or use `nx_load()` instead). This is necessary because these parameters are pointers to arrays of strings, which have no equivalent in Fortran.

`nx_loadve()` returns the number of child processes created, or -1 if any error occurs. If an error occurs, the value 0 is also stored into the *pid\_list* for each process that was not successfully started.

For example, the following C calls create an application as specified by the user (if not specified, the default number of nodes in the default partition), then set the value of the environment variable *HOME* to */tmp*, then load and start the program *myprog* on all nodes of the application with process type 0:

```
#include <nx.h>
#include <stdlib.h>
#include <sys/types.h>
extern char **environ;

main(int argc, char *argv[]) {
    int n;
    pid_t pids[2000];

    n = nx_initve(NULL, 0, NULL, &argc, argv);
    putenv("HOME=/tmp");
    n = nx_loadve(NULL, -1, 0, pids, "myprog", argv, environ);
}
```

The argument list of *myprog* consists of any command-line arguments to the calling program that were not recognized and removed by `nx_initve()`, and the environment of *myprog* is the same as the user's environment except for the value of *HOME*.

## Waiting for Application Processes with `nx_waitall()`

`nx_nfork()`, `nx_load()`, and `nx_loadve()` return immediately to the calling process. To wait for the processes created by `nx_nfork()`, `nx_load()`, or `nx_loadve()` to complete, call `nx_waitall()`.

`nx_waitall()` simply blocks until all the child processes of the calling process have terminated. It returns 0, or -1 if any error occurs.

For example, the following Fortran calls create a new application as specified by the user, run the program *myprog* on all nodes of the application, and wait until all the node processes have completed:

```
include 'fnx.h'
integer n
integer pids(2000)

n = nx_initve("", 0, "", 1, 0)
n = nx_load(0, -1, 0, pids, "myprog")
n = nx_waitall()
```

## Using PIDs

The *pid\_list* argument of `nx_nfork()`, `nx_load()`, and `nx_loadve()` receives the OSF/1 process identifiers (PIDs) of the child processes created by the call. The specified array must be large enough to hold all the PIDs—that is, it must have at least as many elements as the *maximum* number of processes that could be created by the call. If more child processes are created than the number of elements in the *pid\_list*, unexpected results will occur (the program will probably crash).

In the typical case where you create one process per node of the application, you can use the value returned by `nx_initve()` to determine the number of nodes in the application, then use `malloc()` or an equivalent call to dynamically allocate a *pid\_list* with the same number of elements. For example, the following example allocates the appropriate number of elements to the array *pids* based on the application size specified by the user in *argv*:

```
#include <nx.h>
#include <stdio.h>
#include <malloc.h>

main(int argc, char **argv) {
    int    nnodes;
    long  *pids;

    nnodes = nx_initve(NULL, 0, NULL, &argc, argv);
    pids = (long *)calloc(nnodes, sizeof(long));
    nx_nfork(NULL, -1, 0, pids);
```

If you don't use dynamic allocation, you should give the *pid\_list* as many elements as the number of nodes on the largest system on which the application will be run. For portability to very large Intel supercomputers, this array should have at least 1000 elements (and possibly more in the future).

Each element in the *pid\_list* receives the PID of the process on the node specified by the corresponding element of the *node\_list*. If *numnodes* is -1, the PID of the process on node 0 is stored into the first element of *pid\_list*, the PID of the process on node 1 is stored into the second element of *pid\_list*, and so on. If one or more processes were not successfully started, the value 0 is stored into the corresponding element of the *pid\_list*.

## NOTE

The PIDs stored into the *pid\_list* are OSF/1 PIDs, not Paragon OSF/1 process types.

OSF/1 PIDs are unique throughout the system; they are used with standard OSF/1 system calls such as `kill()`. (Note that `kill()` and other system interface routines are supported by the Fortran runtime library; see the *Paragon™ Fortran Compiler User's Guide* for information on these routines.) Paragon OSF/1 process types are unique only within a single application and a single node; they are used with Paragon OSF/1 message-passing calls such as `csend()`.

For example, the following C calls create an application as specified by the user, run the program *myprog* on all nodes of the application with process type 0, and then send the signal `SIGKILL` to all the node processes:

```
#include <nx.h>
#include <signal.h>
#include <sys/types.h>

main(int argc, char *argv[]) {
    int n, i;
    pid_t pids[2000];

    n = nx_initve(NULL, 0, NULL, &argc, argv);
    n = nx_load(NULL, -1, 0, pids, "myprog");

    for(i=0; i<n; i++) {
        kill(pids[i], SIGKILL);
    }
}
```

## Getting Information About Applications

Synopsis	Description
<b>nx_app_rect</b> ( <i>rows, cols</i> )	Obtain the height and width of the rectangle of nodes allocated to the current application.
<b>nx_app_nodes</b> ( <i>pgroup, node_list, list_size</i> )	List the nodes allocated to an application.
<b>nx_ψpart</b> ( <i>partition, ψpart_list, list_size</i> )	Obtain information about all applications and active subpartitions in a partition (C only).

To get information about applications once they are running, use **nx\_app\_rect()**, **nx\_app\_nodes()**, and **nx\_ψpart()**. **nx\_app\_rect()** returns the application's shape (height and width of the rectangle of nodes allocated to the application); **nx\_app\_nodes()** returns a list of the nodes that are allocated to the application; and **nx\_ψpart()** returns information about all the active applications and subpartitions in a partition (like the **ψpart** command).

### NOTE

Do not call **nx\_app\_nodes()** or **nx\_ψpart()** on more than a few nodes at once.

If many nodes use the application information calls at the same time, the allocator daemon can become overwhelmed with requests, which could slow down your application or reduce system stability. If all the nodes in your application need this information, you should have one node make the call and then distribute the information to the other nodes. Note, though, that **nx\_app\_rect()** is not subject to this restriction.

### Finding an Application's Shape with **nx\_app\_rect()**

Sometimes, in addition to its node number and the number of nodes in the application, a process needs to know the *shape* of the application. For example, an application might use a different message-passing algorithm depending on whether the nodes allocated to the application form a square, a tall skinny rectangle, a short wide rectangle, or something else (such as a group of noncontiguous nodes).

To find out the rectangular dimensions of the nodes allocated to the application, call **nx\_app\_rect()**. **nx\_app\_rect()** stores the height of the application into *rows* and the width of the application into *cols*. If the nodes allocated to the application do not form a rectangle, **nx\_app\_rect()** stores 1 into *rows* and **numnodes()** into *cols*. **nx\_app\_rect()** returns 0 if it is successful, or -1 if any error occurs.



For example, the following code fragment in Fortran stores the height of the application into *y* and the width of the application into *x*:

```
integer*4 x, y, result
      .
      .
      .
result = nx_app_rect(y, x)
```

The following code fragment in C does the same:

```
long x, y, result;
      .
      .
      .
result = nx_app_rect(&y, &x);
```

See “Specifying a Rectangle of Nodes” on page 2-16 for information on how to run your application on a rectangular group of nodes with a specific shape.

## NOTE

**nx\_app\_rect()** can also be called by the name **mypart()** for compatibility with the Touchstone DELTA System.

## Listing an Application's Nodes with nx\_app\_nodes()

Occasionally you want to know an application's physical location within the system. You can use this information to help track down possible hardware problems or make use of nodes with special hardware features (such as extra memory or special I/O interfaces).

To list the nodes allocated to an application, call `nx_app_nodes()`. `nx_app_nodes()` has the following parameters:

<i>pgroup</i>	The process group ID of the application (see "Process Groups" on page 4-22 for more information), or 0 to specify the application of the calling process. If the specified process group ID is not the process group ID of the calling process, the calling process's user ID must either be <i>root</i> or the same user ID as the specified application.
<i>node_list</i>	Pointer variable into which <code>nx_app_nodes()</code> stores the address of the list of nodes. The call allocates the memory for this list; when you are finished using the information, you should release this memory by calling <code>free()</code> .
<i>list_size</i>	Variable into which <code>nx_app_nodes()</code> stores the number of entries in <i>node_list</i> .

The node numbers returned by `nx_app_nodes()` are node numbers from the root partition, which tell you where in the machine the application is located. `nx_app_nodes()` returns 0 for success, or -1 if any error occurs.

For example, the following Fortran program fragment prints the root node numbers of the nodes on which the current application is running:

```

include 'fnx.h'

integer*4  mynodes(1)
pointer    (ptr, mynodes)
integer    nnodes
integer    i, status

status = nx_app_nodes(0, ptr, nnodes)

if(status .ne. 0) then
    call nx_perror("nx_app_nodes()")
    stop
end if

do 2, i = 1, nnodes
    print *, mynodes(i)
2   continue

call free(ptr)

```

The equivalent C code is as follows:

```
#include <nx.h>

nx_nodes_t    mynodes;
unsigned long  nnodes;
int           i, status;

status = nx_app_nodes(0, &mynodes, &nnodes);

if(status != 0) {
    nx_perror("nx_app_nodes()");
    exit(1);
}

for(i = 0; i < nnodes; i++) {
    printf("%d\n", mynodes[i]);
}

free(mynodes);
```

Note the use of the & operator on the variables *mynodes* and *nnodes* in the call to `nx_app_nodes()`.

## Listing the Applications in a Partition with `nx_pspart()`

`nx_pspart()` returns information about each of the applications and subpartitions in a partition, like the `pspart` command. It is callable only from C, not Fortran. It has the following parameters:

<i>partition</i>	The relative or absolute pathname of the partition. The specified partition must exist and must give read permission to the calling process.
<i>pspart_list</i>	Pointer variable into which <code>nx_pspart()</code> stores the address of an array of <code>nx_pspart_t</code> structures. Each structure in the array describes one object (application or subpartition). The <code>nx_pspart_t</code> structure is defined in <code>allocsys.h</code> , which is automatically included by <code>nx.h</code> and <code>fnx.h</code> . It includes the following fields:
<i>object_type</i>	The type of the object described by this structure: <code>NX_APPLICATION</code> or <code>NX_PARTITION</code> . (These are constants defined in <code>nx.h</code> or <code>fnx.h</code> ).
<i>object_id</i>	If the object is an application, this is its process group ID. If the object is a partition, this is an arbitrary value and should be ignored.
<i>uid</i>	The numeric user ID of the object's owner.

<i>gid</i>	The numeric group ID of the object's group.
<i>size</i>	The number of nodes allocated to the object.
<i>priority</i>	The current priority of the object.
<i>rolled_in</i>	The amount of time the object has been rolled in during the current rollin quantum, expressed as an integer number of milliseconds.
<i>rollin_q</i>	The rollin quantum for the object's parent partition (that is, the partition specified in the <b>nx_pspart()</b> call), expressed as an integer number of milliseconds.
<i>elapsed</i>	The total amount of time the object has been rolled in since it was started, expressed as an integer number of milliseconds.
<i>active</i>	Whether or not the object is currently rolled in: 1 if it is, 0 if it is not.
<i>time_started</i>	The time the object was started, as returned by <b>time()</b> . (If the object is a subpartition, the time the oldest application in the subpartition was started.)

**nx\_pspart()** allocates the memory for the *pspart\_list* array; when you are finished using the information, you should release this memory by calling **free()**.

*list\_size* Variable into which **nx\_pspart()** stores the number of *nx\_pspart\_t* structures in *pspart\_list*.

**nx\_pspart()** returns 0 for success, or -1 if any error occurs.

For example, the following program fragment prints the numeric user ID and size for every application and subpartition in the partition *mypart*:

```
#include <nx.h>

nx_pspart_t  *info;
unsigned long nobjs;
int          status, i;

status = nx_pspart("mypart", &info, &nobj);

if(status != 0) {
    nx_perror("nx_pspart()");
    exit(1);
}

for(i = 0; i < nobjs; i++) {
    printf("uid = %d, size = %d\n", info->uid, info->size);
}

free(info);
```

Note the use of the `&` operator on the structure *info* and the variable *nobj*s in the call to `nx_pspart()`.

## The Controlling Process

By calling `nx_initve()`, a process creates a new application. The process that called `nx_initve()` becomes the new application's *controlling process*. Each application has exactly one controlling process, and each process controls at most one application.

The controlling process is a special process that creates and controls the application:

- The controlling process can create new processes in the application, using the Paragon OSF/1 function `nx_nfork()`, `nx_load()`, or `nx_loadve()`.
- The controlling process can wait for an application process to complete, using `nx_waitall()` or the standard OSF/1 function `wait()` or `waitpid()`.
- The controlling process can send a signal to an application process, or terminate it, using the standard OSF/1 function `kill()`. In particular, the controlling process can send a signal to all the processes in the application (including itself) by using `kill(0, signal)`.

You can terminate the entire application by terminating the controlling process, using the `kill` command or your interrupt key (normally `<Ctrl-c>` or `<Del>`). The controlling process always runs in the service partition; the application processes run in the partition specified by `nx_initve()`. If the application processes are running in a gang-scheduled partition, the controlling process is

rolled in and out along with its application (that is, when the application is rolled out, the controlling process gets no processor time; when the application is rolled in, the controlling process gets its normal share of the service partition's processor time).

In OSF/1 terms, the controlling process is a *parent* process and the processes created by `nx_nfork()`, `nx_load()`, or `nx_loadve()` are its *child* processes. (In this respect, `nx_nfork()` is similar to `fork()`, `nx_load()` is similar to a `fork()` followed by an `execv()` with a null argument list, and `nx_loadve()` is similar to a `fork()` followed by an `execve()`). The controlling process and the application processes all belong to the same *process group*, and the controlling process is the *process group leader* of the group. No process outside the application belongs to this process group.

The controlling process does not usually do heavy computational work, because it runs in the service partition along with users' shells and other interactive processes. Since it is scheduled interactively, the controlling process will not give as much throughput as application processes running in gang-scheduled compute partitions.

See the *OSF/1 Programmer's Reference* for information on `wait()`, `waitpid()`, `kill()`, `fork()`, and `execve()`.

## Process Groups

*Process groups* are a standard OSF/1 concept, not unique to Paragon OSF/1. A process group is a set of related processes. You can send a signal to all the processes in a group at once with `kill()`, and you can wait for any process in a group with `waitpid()`. The processes in a process group also share access to a terminal, called the *controlling terminal* of the group. Each process belongs to exactly one process group.

The processes in a process group are all children (or grandchildren, and so on) of the oldest process in the group, called the *process group leader*. The process group leader's process ID (PID) is used to identify the group, and is also called the *process group ID* of the whole group. (Note that this is the process group leader's OSF/1 PID, not its process type.) A process can determine its process group ID by calling `getpgrp()`.

Normally, a process belongs to the same process group as its parent process. However, a process can leave its parent's process group and start a new process group of its own by making such calls as `setpgid()`, `setpgrp()`, or `setsid()`. These calls create a new process group, then remove the calling process from its current group and place it in the new group. The calling process becomes the new group's process group leader, and the caller's PID becomes the new group's process group ID. After that, any processes created by the process group leader belong to the new process group. See the *OSF/1 Programmer's Reference* for information on `setpgid()` and `getpgrp()`.

## Process Groups in Paragon™ OSF/1

In Paragon OSF/1, process groups work the same as they do in standard OSF/1. In addition, `nx_initve()` makes sure that the calling process is a process group leader. If the calling process is not already a process group leader, `nx_initve()` has the same effect as `setpgid()`: it creates a new process group and makes the calling process the new group's process group leader. Because all the processes in the application are created by the controlling process, all the processes in an application are members of the same process group, and no other process in the system is a member of that process group. This means that the application's process group ID uniquely identifies the application, which is why you use a process group ID to identify the application in `nx_pri()`.

This also means that if a process in an application leaves the application's process group by calling `nx_initve()` (or `setpgid()`, `setpgrp()`, or `setsid()`), it leaves the application. If a process leaves its application's process group, it is no longer considered part of the application and can no longer exchange messages with the other processes in the application. You shouldn't do this unless you know exactly what you are doing.

## Killing Application Processes

You can take advantage of the fact that all the processes in the application are members of the same process group by using OSF/1 system calls that affect process groups. For example, specifying a process ID of 0 (zero) to `kill()` sends the specified signal to all the members of the calling process's process group, so the following call kills the entire application (including the calling process):

```
kill(0, SIGKILL);
```

This call differs from the example discussed under "Using PIDs" on page 4-14 in that it also kills the calling process.

## An Example Controlling Process

The following C program (which must be linked with `-lnx`, not `-nx`) copies itself onto eight nodes of the partition `mypart` with a process type of 0 and a priority of 7. The eight application processes print "Hello from node *n*" and then exit. The controlling process waits for the application processes to finish, then prints "Hello from controlling process" before exiting itself. Note that this program is executed by both the controlling process and the application processes.

```
#include <nx.h>
#include <sys/types.h>
#include <stdio.h>
#define NUMNODES 8

main(int argc, char **argv) {
    int n, i;
    pid_t pids[NUMNODES];

    /* create application */
    n = nx_initve("mypart", NUMNODES, NULL, &argc, argv);
    if(n == -1) {
        /* nx_initve() failed */
        perror("nx_initve");
        exit(1);
    }

    /* set application priority to 7 */
    n = nx_pri(0, 7); /* 0 specifies calling application */
    if(n == -1) {
        /* nx_pri() failed */
        perror("nx_pri");
        exit(1);
    }

    /* fork child processes onto all nodes of application */
    n = nx_nfork(NULL, -1, 0, pids);
    if (n == -1) {
        /* nx_nfork() failed */
        perror("nx_nfork");
        exit(1);
    } else if(n == 0) {
        /* child process: print "Hello" and exit */
        printf("Hello from node %d!\n", mynode());
        exit(0);
    } else {
        /* parent (controlling process): wait for all children */
        nx_waitall();
        /* now print "Hello" and exit */
        printf("Hello from controlling process!\n");
        exit(0);
    }
}
```



## Message Passing Between Controlling Process and Application Processes

Synopsis	Description
<code>myhost()</code>	Obtain the controlling process's node number.

A controlling process can exchange messages with its child processes using the Paragon OSF/1 message-passing calls described in Chapter 3.

- The controlling process's node number is equal to `numnodes()`. (The maximum node number within the application is `numnodes() - 1`.) The controlling process's node number is also returned by `myhost()` in any process in the application. In the controlling process, `myhost()`, `mynode()`, and `numnodes()` all return the same number.
- The controlling process's process type is initially `INVALID_PTYPE`, but you can change it to a valid value by calling `setptype()`. For best performance, you should not call `setptype()` until *after* you have created all application processes with `nx_nfork()`, `nx_load()`, or `nx_loadve()`, and you should not call `setptype()` at all unless you need to exchange messages with application processes.

Although the controlling process can exchange messages with the application processes, it does not participate in global operations:

- The controlling process may not make any of the calls described under "Global Operations" on page 3-27.
- The controlling process does not participate when the application processes make any of the calls described under "Global Operations" on page 3-27.
- The controlling process does not get messages sent to node number -1 (all nodes).

A send to node -1 (all nodes) sends the message to all the nodes in the application (except the calling process's node), but not the controlling process. This applies whether the message is sent by a node process or by the controlling process itself. On the other hand, an extended receive that specifies node -1 (any node) as the sending node *will* match a message from the controlling process.

Here is an application, written in Fortran, that demonstrates message-passing between the controlling process and the application processes. This application multiplies two numbers (in a very inefficient way). It consists of two programs, `control.f` and `app.f`. You must link `control.f` with `-lnx`, not `-nx`; `app.f` can be linked with either `-lnx` or `-nx`.

The controlling process (*control.f*) requests a number of nodes and an integer value from the user. It creates an application of the specified number of nodes on the partition *mypart* and loads the program *app* onto each node. It then sends the user's integer value to each node as a message (note that the node number -1 sends to all nodes, not including the controlling process) and waits for a return message with the result. When the result is received, the controlling process prints its value and then exits.

```

include 'fnx.h'

integer      num_nodes, n, i
integer      nodes(128), pids(128)
integer      app_ptype
parameter    (app_ptype = 0)
integer      data, result
integer      result_type, data_type
parameter    (result_type = 1)
parameter    (data_type = 2)

c get number of nodes (limited to size of "nodes" and "pids" arrays)
1  print *, "Enter number of nodes (must not be greater than 128)"
   read(*,*) num_nodes
   if(num_nodes .gt. 128) goto 1

c create application of specified size
   n = nx_initve("mypart", num_nodes, "", 0, 0)
   if(n .eq. -1) then
       call nx_perror("nx_initve")
       stop
   end if

c fill in node array for nx_load()
   do 2, i = 1, num_nodes
       nodes(i) = i - 1
2  continue

c load program "app" onto the nodes of the application
   n = nx_load(nodes, num_nodes, app_ptype, pids, "app")
   if(n .eq. -1) then
       call nx_perror("nx_load")
       stop
   end if

c get an integer from the user
   print *, "Enter value to be summed"
   read(*,*) data

c set my process type
   call setptype(app_ptype)

```

```

c send integer to all the nodes
  call csend(data_type, data, 4, -1, app_ptype)

c receive the result
  call crecv(result_type, result, 4)

c print the result
  print *, "The sum of ",data," over ",num_nodes," nodes is ",result

end

```

The application process (*app.f*) waits for a message and performs a **gisum()** on the value received. (Note that the controlling process does not participate in the **gisum()**.) The process on node 0 sends the result to the controlling process, then all the application processes exit.

```

include 'fnx.h'

integer    val, work
integer    result_type, data_type
parameter  (result_type = 1)
parameter  (data_type = 2)

c get an integer from the controlling process
  call crecv(data_type, val, 4)

c sum it over all nodes
  call gisum(val, 1, work)

c if I'm node 0, send the result back to the controlling process
  if(mynode() .eq. 0) call csend(result_type, val, 4, myhost(), 0)

end

```

## Managing Partitions

Paragon OSF/1 provides system calls that let you create and remove partitions, get information about partitions, and change their characteristics, like the **mkpart**, **rmpart**, **showpart**, and **chpart** commands described in Chapter 2. See "Managing Partitions" on page 2-25 for introductory information on partitions.

## Making Partitions

Synopsis	Description
<code>nx_mkpart(partition, size, type)</code>	Create a partition with a particular number of nodes.
<code>nx_mkpart_rect(partition, rows, cols, type)</code>	Create a partition with a particular height and width.
<code>nx_mkpart_map(partition, numnodes, node_list, type)</code>	Create a partition with a specific set of nodes.

To create a partition, use `nx_mkpart()`, `nx_mkpart_rect()`, or `nx_mkpart_map()`. These calls all create a partition, but they use different methods to specify the nodes allocated to the new partition:

- `nx_mkpart()` works like the `mkpart` command's `-sz size` switch.
- `nx_mkpart_rect()` works like the `mkpart` command's `-sz hXw` switch.
- `nx_mkpart_map()` works like the `mkpart` command's `-nd nodespec` switch (except that only node numbers can be specified).

See "Specifying the Nodes Allocated to the Partition" on page 2-40 for more information on these switches.

These calls have the following parameters:

<i>partition</i>	The new partition's relative or absolute pathname. The specified new partition must not exist; the parent partition of the specified new partition must exist and must give write permission to the calling process. See "Partition Pathnames" on page 2-28 for more information on partition pathnames; see "Owner, Group, and Protection Modes" on page 2-32 for more information on partition permissions.
<i>size</i>	The number of nodes of the new partition, or -1 to specify "all the nodes of the parent partition." If you specify a size smaller than that of the parent partition, the nodes are selected by the system (and are not necessarily contiguous).

`nx_mkpart()` attempts to allocate a square group of nodes if it can. If this is not possible, it attempts to allocate a rectangular group of nodes that is either twice as wide as it is high or twice as high as it is wide. If this is not possible, it allocates any available nodes. In this case, nodes allocated to the partition may not be contiguous.

*rows* and *cols*    The height and width of the new partition. The new partition is a rectangle with the specified number of rows and columns, but its location within the parent partition is selected by the system.

*numnodes* and *node\_list*    The exact node numbers within the parent partition for the new partition. The *node\_list* parameter is an array of node numbers; the *numnodes* parameter specifies the number of elements in *node\_list*.

*type*    The new partition's scheduling type: **NX\_STD** to specify standard scheduling, **NX\_GANG** to specify gang scheduling, or **NX\_SPS** to specify space sharing. The names **NX\_STD**, **NX\_GANG**, and **NX\_SPS** are defined in *nx.h* and *fnx.h*. See "Scheduling Characteristics" on page 2-33 for more information on the different scheduling types.

**nx\_mkpart()**, **nx\_mkpart\_rect()**, and **nx\_mkpart\_map()** return the number of nodes in the new partition, or -1 if any error occurs.

The new partition's owner and group are set to the owner and group of the calling process. All other partition characteristics not specified in the call (such as protection modes and rollin quantum) are set to the same values as the parent partition. Once the partition is created, you can use the **nx\_chpart...()** calls to set these characteristics to different values, as discussed under "Changing Partition Characteristics" on page 4-36.

For example, the following Fortran call creates a new gang-scheduled partition called *newpart* whose parent partition is the compute partition (using a relative partition pathname) and which consists of all the nodes in the compute partition:

```
include 'fnx.h'
integer n

n = nx_mkpart("newpart", -1, NX_GANG)
```

The following C call creates a new space-shared partition called *mypart* whose parent partition is the compute partition (using an absolute partition pathname) and which has 54 nodes:

```
#include <nx.h>
int n;

n = nx_mkpart(".compute.mypart", 54, NX_SPS);
```

The following C call creates a new gang-scheduled partition called *rect* whose parent partition is *mypart* and which is 3 nodes high and 4 nodes wide:

```
#include <nx.h>
int n;

n = nx_mkpart_rect(".compute.mypart.rect", 3, 4, NX_GANG);
```

The following C call creates a new space-shared partition called *corners* whose parent partition is *rect* and which consists of the four nodes at the corners of *rect*:

```
#include <nx.h>
long nodes[4];
int n;

nodes[0] = 0;
nodes[1] = 3;
nodes[2] = 8;
nodes[3] = 11;
n = nx_mkpart_map( ".compute.mypart.rect.corners", 4,
                  nodes, NX_SPS);
```

In each of these examples, the variable *n* is assigned the number of nodes in the new partition, or -1 if any error occurred.

## Removing Partitions

Synopsis	Description
<code>nx_rmpart(partition, force, recursive)</code>	Remove a partition.

To remove a partition, use `nx_rmpart()`. The parameters of `nx_rmpart()` have the following meanings:

<i>partition</i>	The relative or absolute pathname of the partition to be removed. The parent partition of the specified partition must give write permission to the calling process. See “Partition Pathnames” on page 2-28 for more information on partition pathnames; see “Owner, Group, and Protection Modes” on page 2-32 for more information on partition permissions.
<i>force</i>	Specifies whether to remove the partition if it contains running applications: if <i>force</i> is 0, the partition will not be removed if it contains any applications; if <i>force</i> is any value other than 0, the partition will be removed even if it contains applications.
<i>recursive</i>	Specifies whether to remove the partition if it contains subpartitions: if <i>recursive</i> is 0, the partition will not be removed if it contains any subpartitions; if <i>recursive</i> is any value other than 0, the partition will be removed along with all its subpartitions, sub-subpartitions, and so on. This is an “all or nothing” operation: if any subpartitions cannot be removed, the call fails and no subpartitions are removed.

If the partition contains both subpartitions and applications, or contains subpartitions that contain applications, you must set both *force* and *recursive* to a nonzero value to remove it.

**nx\_rmupart()** returns 0 for success, or -1 if any error occurs.

For example, the following Fortran call removes the partition called *newpart* whose parent partition is the compute partition (using a relative partition pathname), but only if it does not contain any running applications or subpartitions:

```
include 'fnx.h'
integer n

n = nx_rmupart("newpart", 0, 0)
```

After this call, the variable *n* contains 0 if the partition was removed, or -1 if it was not removed for any reason (for example, if the partition contained applications or subpartitions).

The following C call removes the partition called *mypart* whose parent partition is the compute partition (using an absolute partition pathname), even if it contains running applications; however, it does not remove *mypart* if the partition contains subpartitions:

```
#include <nx.h>
int n;

n = nx_rmupart(".compute.mypart", 1, 0);
```

After this call, the variable *n* contains 0 if the partition was removed, or -1 if it was not removed for any reason (for example, if the partition contained subpartitions, or if the partition does not exist).

## Getting Information About Partitions

Synopsis	Description
<b>nx_part_attr</b> ( <i>partition, attributes</i> )	Get a partition's attributes.
<b>nx_part_nodes</b> ( <i>partition, node_list, list_size</i> )	List the root node numbers for the nodes of a partition.

To get information about a partition, use **nx\_part\_attr()** or **nx\_part\_nodes()**. **nx\_part\_attr()** returns the attributes of a partition, and **nx\_part\_nodes()** returns a list of the nodes in the specified partition.

## NOTE

Do not call `nx_part_attr()` or `nx_part_nodes()` on more than a few nodes at once.

If many nodes use the partition information calls at the same time, the allocator daemon can become overwhelmed with requests, which could slow down your application or reduce system stability. If all the nodes in your application need this information, you should have one node make the call and then distribute the information to the other nodes.

## Determining a Partition's Attributes with `nx_part_attr()`

`nx_part_attr()` returns the attributes of a partition. It has the following parameters:

<i>partition</i>	The relative or absolute pathname of the partition. The specified partition must exist and must give read permission to the calling process.																
<i>attributes</i>	A structure of type <code>nx_part_info_t</code> (you must allocate the space for this structure). The <code>nx_part_info_t</code> structure is defined in <code>allocsys.h</code> , which is automatically included by <code>nx.h</code> and <code>fnx.h</code> . It includes the following elements: <table> <tr> <td><i>uid</i></td> <td>The numeric user ID of the partition's owner.</td> </tr> <tr> <td><i>gid</i></td> <td>The numeric group ID of the partition's group.</td> </tr> <tr> <td><i>access</i></td> <td>The access permissions of the partition, expressed as a three-digit octal number.</td> </tr> <tr> <td><i>sched</i></td> <td>The scheduling type of the partition: <code>NX_STD</code>, <code>NX_GANG</code>, or <code>NX_SPS</code>. (These are constants defined in <code>nx.h</code> or <code>fnx.h</code>).</td> </tr> <tr> <td><i>rq</i></td> <td>The rollin quantum of the partition, expressed as an integer number of milliseconds (0 for a standard-scheduled or space-shared partition).</td> </tr> <tr> <td><i>epl</i></td> <td>The effective priority limit of the partition (20 for a standard-scheduled partition).</td> </tr> <tr> <td><i>nodes</i></td> <td>The number of nodes in the partition.</td> </tr> <tr> <td><i>mesh_x</i></td> <td>The width of the partition (columns), or -1 if the partition is not rectangular.</td> </tr> </table>	<i>uid</i>	The numeric user ID of the partition's owner.	<i>gid</i>	The numeric group ID of the partition's group.	<i>access</i>	The access permissions of the partition, expressed as a three-digit octal number.	<i>sched</i>	The scheduling type of the partition: <code>NX_STD</code> , <code>NX_GANG</code> , or <code>NX_SPS</code> . (These are constants defined in <code>nx.h</code> or <code>fnx.h</code> ).	<i>rq</i>	The rollin quantum of the partition, expressed as an integer number of milliseconds (0 for a standard-scheduled or space-shared partition).	<i>epl</i>	The effective priority limit of the partition (20 for a standard-scheduled partition).	<i>nodes</i>	The number of nodes in the partition.	<i>mesh_x</i>	The width of the partition (columns), or -1 if the partition is not rectangular.
<i>uid</i>	The numeric user ID of the partition's owner.																
<i>gid</i>	The numeric group ID of the partition's group.																
<i>access</i>	The access permissions of the partition, expressed as a three-digit octal number.																
<i>sched</i>	The scheduling type of the partition: <code>NX_STD</code> , <code>NX_GANG</code> , or <code>NX_SPS</code> . (These are constants defined in <code>nx.h</code> or <code>fnx.h</code> ).																
<i>rq</i>	The rollin quantum of the partition, expressed as an integer number of milliseconds (0 for a standard-scheduled or space-shared partition).																
<i>epl</i>	The effective priority limit of the partition (20 for a standard-scheduled partition).																
<i>nodes</i>	The number of nodes in the partition.																
<i>mesh_x</i>	The width of the partition (columns), or -1 if the partition is not rectangular.																



- mesh\_y* The height of the partition (rows), or -1 if the partition is not rectangular.
- enclose\_mesh\_x* The width of the smallest rectangle that completely encloses the partition.
- enclose\_mesh\_y* The height of the smallest rectangle that completely encloses the partition.

`nx_part_attr()` returns 0 for success, or -1 if any error occurs.

For example, the following C program fragment prints the rollin quantum and effective priority limit for the partition *mypart*:

```
#include <nx.h>

nx_part_info_t info;
int            status;

status = nx_part_attr("mypart", &info);

if(status != 0) {
    nx_perror("nx_part_attr()");
    exit(1);
}

printf("rq = %d, epl = %d\n", info.rq, info.epl);
```

Note the use of the `&` operator on the structure *info* in the call to `nx_part_attr()`. The equivalent Fortran code is as follows:

```
include 'fnx.h'

record /nx_part_info_t/ info
integer status

status = nx_part_attr("mypart", info)

if(status .ne. 0) then
    call nx_perror("nx_part_attr()")
    stop
end if

print *, "rq =", info.rq, ", epl =", info.epl
```

If the partition is not a contiguous rectangle, the values of *mesh\_x* and *mesh\_y* are -1 and the rectangle described by *enclose\_mesh\_x* and *enclose\_mesh\_y* includes nodes that are not part of the partition. For example, Figure 4-1 shows a non-rectangular partition called *mypart*. For this partition:

- *nodes* is 4.
- *mesh\_x* and *mesh\_y* are both -1.
- *enclose\_mesh\_x* is 3.
- *enclose\_mesh\_y* is 2.

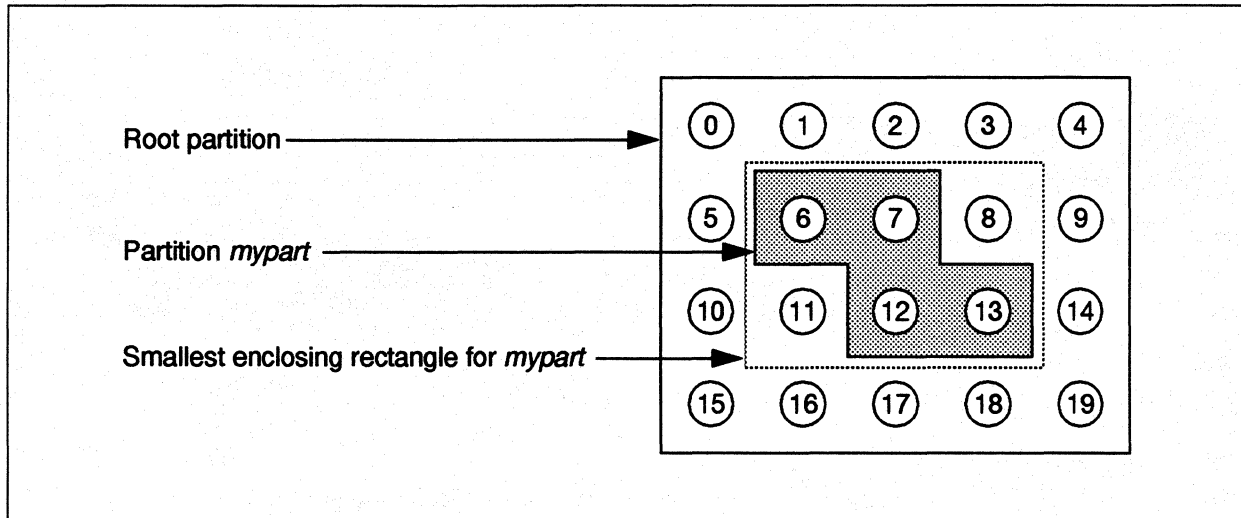


Figure 4-1. Sample Partition for `nx_part_attr()` and `nx_part_nodes()`

## Determining a Partition's Nodes with `nx_part_nodes()`

`nx_part_nodes()` returns a list of the nodes in the specified partition. You might want to do this to determine whether or not the partition includes a certain node which has special hardware characteristics such as extra memory or an I/O interface. `nx_part_nodes()` has the following parameters:

<i>partition</i>	The relative or absolute pathname of the partition. The specified partition must exist and must give read permission to the calling process.
<i>node_list</i>	Pointer variable into which <code>nx_part_nodes()</code> stores the address of the list of nodes. <code>nx_part_nodes()</code> allocates the memory for this list; when you are finished using the information, you should release this memory by calling <code>free()</code> .
<i>list_size</i>	Variable into which <code>nx_part_nodes()</code> stores the number of entries in <i>node_list</i> .

`nx_part_nodes()` returns 0 for success, or -1 if any error occurs.

The node numbers returned by `nx_part_nodes()` are node numbers from the root partition. For example, `nx_part_nodes()` for the partition *mypart* shown in Figure 4-1 would return node numbers 6, 7, 12, and 13. This is true even if the root partition is not the direct parent partition of *mypart*.

For example, the following Fortran program fragment prints the root node numbers for the partition *mypart*:

```

include 'fnx.h'

integer*4  mynodes(1)
pointer    (ptr, mynodes)
integer    nnodes
integer    i, status

status = nx_part_nodes("mypart", ptr, nnodes)

if(status .ne. 0) then
    call nx_perror("nx_part_nodes()")
    stop
end if

do 2, i = 1, nnodes
    print *, mynodes(i)
2   continue

call free(ptr)

```

The equivalent C code is as follows:

```
#include <nx.h>

nx_nodes_t    mynodes;
unsigned long  nnodes;
int           i, status;

status = nx_part_nodes("mypart", &mynodes, &nnodes);

if(status != 0) {
    nx_perror("nx_part_nodes()");
    exit(1);
}

for(i = 0; i < nnodes; i++) {
    printf("%d\n", mynodes[i]);
}

free(mynodes);
```

Note the use of the & operator on the variables *mynodes* and *nnodes* in the call to `nx_part_nodes()`.

## Changing Partition Characteristics

Synopsis	Description
<code>nx_chpart_name(partition, name)</code>	Change a partition's name.
<code>nx_chpart_mod(partition, mode)</code>	Change a partition's protection modes.
<code>nx_chpart_epl(partition, priority)</code>	Change a partition's effective priority limit.
<code>nx_chpart_rq(partition, rollin_quantum)</code>	Change a partition's rollin quantum.
<code>nx_chpart_owner(partition, owner, group)</code>	Change a partition's owner and group.
<code>nx_chpart_sched(partition, sched_type)</code>	Change a partition's scheduling type.

To change a partition's characteristics, use `nx_chpart_name()`, `nx_chpart_mod()`, `nx_chpart_epl()`, `nx_chpart_rq()`, `nx_chpart_owner()`, or `nx_chpart_sched()`. Each of these calls changes one characteristic, and leaves the other characteristics unchanged. These calls have the following parameters:

*partition*      The relative or absolute pathname of the partition to change. The specified partition must exist; the permissions required depend on the operation.

**name (nx\_chpart\_name() only)**

The new name for the partition, expressed as a string of any length containing only uppercase letters, lowercase letters, digits, and underscores. Note that **nx\_chpart\_name()** can only change the partition's name "in place;" there is no way to move a partition to a different parent partition.

The calling process must have write permission on the parent partition of the specified partition to use **nx\_chpart\_name()**.

**mode (nx\_chpart\_mod() only)**

The new protection modes of the partition, expressed as an octal number. See **chmod()** in the *OSF/1 Programmer's Reference* for more information on specifying protection modes; see "Owner, Group, and Protection Modes" on page 2-32 for more information on protection modes for partitions.

The calling process must be the owner of the partition or the system administrator to use **nx\_chpart\_mod()**.

**priority (nx\_chpart\_epl() only)**

The new effective priority limit for the partition, expressed as an integer from 0 to 10 inclusive. See "Scheduling Characteristics" on page 2-33 for more information on effective priority limits.

The calling process must have write permission for the partition to use **nx\_chpart\_epl()**.

**rollin\_quantum (nx\_chpart\_rq() only)**

The new rollin quantum for the partition, expressed as an integer number of milliseconds, or 0 to specify an "infinite" rollin quantum. The specified value must not be greater than 86,400,000 milliseconds (24 hours) and must not be less than the minimum rollin quantum for your system (determined by your system administrator). If it is not a multiple of 100, it is silently rounded up to the next multiple of 100. See "Scheduling Characteristics" on page 2-33 for more information on rollin quanta.

The calling process must have write permission for the partition to use **nx\_chpart\_rq()**.

**owner and group (nx\_chpart\_owner() only)**

The new user and group for the partition, expressed as a numeric user ID (UID) and group ID (GID). You can also specify -1, meaning "leave owner/group unchanged," for either or both. See "Owner, Group, and Protection Modes" on page 2-32 for more information on partition ownership.

The permissions required for `nx_chpart_owner()` depend on the operation. To change the partition's ownership, the calling process must be the system administrator. To change the partition's group, the calling process must either be the system administrator or must be the partition's owner and changing the group to a group that the calling process belongs to.

*sched\_type* (`nx_chpart_sched()` only)

The new scheduling type for the partition, which must be `NX_GANG` or `NX_SPS` (constants defined in `nx.h` or `fnx.h`). See "Scheduling Characteristics" on page 2-33 for more information on gang-scheduling and space sharing.

The specified partition must not be standard-scheduled. A space-shared partition can be changed to gang-scheduled at any time; a gang-scheduled partition can only be changed to space-shared if it contains no applications and no overlapping subpartitions.

The calling process must have write permission for the partition to use `nx_chpart_sched()`.

`nx_chpart_name()`, `nx_chpart_mod()`, `nx_chpart_epl()`, `nx_chpart_rq()`, `nx_chpart_owner()`, and `nx_chpart_sched()` return 0 for success, or -1 if any error occurs.

For example, the following Fortran call changes the name of *mypart* to *newpart*:

```
include 'fnx.h'
integer n

n = nx_chpart_name("mypart", "newpart")
```

The following C call has the same effect, but uses an absolute partition pathname:

```
#include <nx.h>
int n;

n = nx_chpart_name(".compute.mypart", "newpart");
```

Note that the second parameter of `nx_chpart_name()` is always a partition name, never a partition pathname. There is no way to move a partition from one parent partition to another.

The following C call sets the permissions of *mypart* to `rxrx-x---` (750 octal):

```
#include <nx.h>
int n;

n = nx_chpart_mod("mypart", 0750);
```

The following Fortran call has the same effect, but uses an absolute partition pathname:

```
include 'fnx.h'
integer n

n = nx_chpart_mod(".compute.mypart", '750'0)
```

The following C call sets *mypart*'s effective priority limit to 7:

```
#include <nx.h>
int n;

n = nx_chpart_epl("mypart", 7);
```

The following Fortran call sets *mypart*'s rollin quantum to 10 minutes (600,000 microseconds):

```
include 'fnx.h'
integer n

n = nx_chpart_rq("mypart", 600000)
```

The following C calls set *mypart*'s owner to *fred* and its group to *devel* (see the *OSF/1 Programmer's Reference* for information on `getpwnam()` and `getgrnam()`, which get the numeric user and group IDs based on their names):

```
#include <stdio.h>
#include <pwd.h>
#include <grp.h>
#include <nx.h>

struct passwd *user;
struct group *group;
int n;

user = getpwnam("fred");
group = getgrnam("devel");
n = nx_chpart_owner("mypart", user->pw_uid, group->gr_gid);
```

The following Fortran call changes *mypart* to a gang-scheduled partition (it must currently be either gang-scheduled or space-shared):

```
include 'fnx.h'
integer n

n = nx_chpart_sched("mypart", NX_GANG)
```

In each of these examples, the variable *n* is assigned 0 if the call succeeded, or -1 if any error occurred.

## Listing Unusable Nodes

Synopsis	Description
<b>nx_empty_nodes</b> ( <i>node_list</i> , <i>list_size</i> )	List the nodes that are empty slots.
<b>nx_failed_nodes</b> ( <i>node_list</i> , <i>list_size</i> )	List the nodes that failed to boot.

To find out which nodes in the system are unusable, use **nx\_empty\_nodes()** and **nx\_failed\_nodes()**. (See “Unusable Nodes” on page 2-31 for more information on unusable nodes.)

- **nx\_empty\_nodes()** returns a list of the nodes that are part of the root partition but do not have a node board installed in the corresponding slot (these are shown as “-” in the output of **showpart**).
- **nx\_failed\_nodes()** returns a list of the nodes that are part of the root partition but failed to boot for some reason (these are shown as “X” in the output of **showpart**).

### NOTE

Do not call **nx\_empty\_nodes()** or **nx\_failed\_nodes()** on more than a few nodes at once.

If many nodes use these calls at the same time, the allocator daemon can become overwhelmed with requests, which could slow down your application or reduce system stability. If all the nodes in your application need this information, you should have one node make the call and then distribute the information to the other nodes.

Both these calls have the following parameters:

<i>node_list</i>	Pointer variable into which the call stores the address of the list of nodes. The call allocates the memory for this list; when you are finished using the information, you should release this memory by calling <b>free()</b> .
<i>list_size</i>	Variable into which the call stores the number of entries in <i>node_list</i> .

The node numbers returned by these calls are node numbers from the root partition. Both calls return 0 for success, or -1 if any error occurs.



For example, the following Fortran program fragment prints the node numbers of all empty slots in the root partition:

```

      include 'fnx.h'

      integer*4  empty(1)
      pointer    (ptr, empty)
      integer    nempty
      integer    i, status

      status = nx_empty_nodes(ptr, nempty)

      if(status .ne. 0) then
         call nx_perror("nx_empty_nodes()")
         stop
      end if

      do 2, i = 1, nempty
         print *, empty(i)
2      continue

      call free(ptr)

```

The following C program fragment prints the node numbers of all nodes in the root partition that failed to boot:

```

#include <nx.h>

nx_nodes_t    failed;
unsigned long  nfailed;
int           i, status;

status = nx_failed_nodes(&failed, &nfailed);

if(status != 0) {
    nx_perror("nx_failed_nodes()");
    exit(1);
}

for(i = 0; i < nfailed; i++) {
    printf("%d\n", failed[i]);
}

free(failed);

```

Note the use of the & operator on the variables *failed* and *nfailed* in the call to `nx_failed_nodes()`.

## Handling Errors

Synopsis	Description
<code>_call()</code>	Special version of <i>call</i> that returns error value to caller (C only).
<code>nx_perror(string)</code>	Print an error message corresponding to the current value of <i>errno</i> .

When an error occurs in a standard OSF/1 system call, the call indicates the error in one of two ways, depending on the error. For most errors, the call returns -1 and sets the variable *errno* to a value that describes the error. For certain severe errors (such as a segmentation violation caused by an invalid pointer parameter), the call sends a signal to the calling process; this signal may result in a *core dump*, as discussed under “Core Dumps” on page 4-44.

When an error occurs in a Paragon OSF/1 system call whose name begins with `nx_`, it uses the same two techniques as a standard OSF/1 system call. However, when an error occurs in a Paragon OSF/1 system call that is not a standard OSF/1 system call and whose name does *not* begin with `nx_`, the error is handled differently: the system prints a message on the terminal and terminates the calling process. (There are exceptions; see the manual page for each call in the *Paragon™ C System Calls Reference Manual* or *Paragon™ Fortran System Calls Reference Manual* for details.) If you program in C, you can get the same behavior as the `nx_` calls by calling the *underscore* version of the call. (Fortran does not have underscore versions.)

## Underscore Calls

The underscore version of a Paragon OSF/1 system call is the same as the standard version except that it has an underscore added to the beginning of its name. For example, `_crecv()` is the underscore version of `crecv()`. The underscore version returns -1 if the call encounters an error and 0 or a positive value if the call is successful.

If an error occurs, the underscore version also sets the system variable *errno* to indicate the cause of the error. The include file *errno.h* declares *errno* for you and defines constants for the possible *errno* values. For example, if `crecv()` receives a message that is larger than the size specified by its *len* parameter, an error message appears and the application terminates. If you use `_crecv()` instead, this does not occur; instead, the call to `_crecv()` returns -1 and the variable *errno* is set to the value `EQMSGLONG`.

There is a standard error message for each value of *errno*, which you can print out by calling `nx_perror()`. `nx_perror()` prints its argument (any string), the current node number and process type, and the error message associated with the current value of *errno* to the standard error output in the following format:

```
(node n, ptype p) string: error_message
```

Suppose you have a program where the user can specify the size of a certain buffer with a command-line argument. If a message is received that is too long for this buffer, you would like to be able to tell the user what happened and suggest that they increase the buffer size. The following example uses the underscore version of `crecv()` to do this:

```
#include <nx.h>
#include <errno.h>

char *transbuf;
int transbuf_size;
.
.
.
if(_crecv(1, transbuf, transbuf_size) == -1) {
    if(errno == EQMSGLONG) {
        /* received message too long for buffer */
        printf("Message exceeded transit buffer size!\n");
        printf("Use -t to specify a larger transit buffer.\n");
        exit(1);
    } else {
        /* some other error, print a standard error
        message and exit*/
        nx_perror("crecv");
        exit(1);
    }
}
```

## Core Dumps

When certain severe errors occur in a standard OSF/1 system call or a Paragon OSF/1 system call whose name begins with **nx\_**, the operating system sends a signal to the calling process. The default action for the following signals is to cause a core dump:

<b>SIGABRT</b> (also called <b>SIGIOT</b> and <b>SIGLOST</b> )	Abort process (can be generated by the <b>abort()</b> system call).
<b>SIGBUS</b>	Bus error (specification exception).
<b>SIGEMT</b>	EMT instruction.
<b>SIGFPE</b>	Floating point exception.
<b>SIGILL</b>	Illegal instruction.
<b>SIGQUIT</b>	Quit (can be generated by the user typing <b>&lt;Ctrl-\&gt;</b> on the terminal).
<b>SIGSEGV</b>	Segmentation violation.
<b>SIGSYS</b>	Bad argument to system call.
<b>SIGTRAP</b>	Trace trap.

A *core dump* means that the process terminates immediately and writes a copy of its current memory contents to a file called *core* in the current working directory. You can prevent this default action by establishing a *signal handler* for the desired signal; see **signal()** in the *OSF/1 Programmer's Reference* for information on signal handlers.

### NOTE

No tools are currently provided for analyzing core files or debugging with core files.

If one process in an application dumps core, it may or may not terminate the rest of the application, as follows:

- If the application was linked with the **-nx** switch, when one process of the application is terminated by the signal **SIGBUS**, **SIGFPE**, **SIGILL**, **SIGSEGV**, or **SIGSYS** the whole application is terminated.
- If the application was not linked with **-nx** but does call **nx\_waitall()**, if **nx\_waitall()** detects that one of the processes being waited for has been terminated by the signal **SIGBUS**, **SIGFPE**, **SIGILL**, **SIGSEGV**, or **SIGSYS**, then **nx\_waitall()** terminates the whole application by sending a **SIGKILL** to the process group.
- In any other case (that is, if a process dumps core because of **SIGABRT**, **SIGEMT**, **SIGQUIT**, or **SIGTRAP**, the application was not linked with **-nx**, or the process that dumped core is not being waited for by **nx\_waitall()**), the other processes in the application are not directly affected by one process dumping core. However, if all the processes in the application are running the same code, all processes may dump core independently. If several processes in an application dump core, they will all write their core dumps to the same *core* file unless the processes have changed to different working directories.

## Controlling Floating-Point Behavior

---

Synopsis	Description
<b>isnan</b> ( <i>dsrc</i> )	Determine if a <b>double</b> value is Not-a-Number (C only).
<b>isnand</b> ( <i>dsrc</i> )	Determine if a <b>double</b> value is Not-a-Number (C only).
<b>isnanf</b> ( <i>fsrc</i> )	Determine if a <b>float</b> value is Not-a-Number (C only).
<b>fpgetround</b> ()	Get the floating-point rounding mode for the calling process (C only).
<b>fpsetround</b> ( <i>rnd_dir</i> )	Set the floating-point rounding mode for the calling process (C only).
<b>fpgetmask</b> ()	Get the floating-point exception mask for the calling process (C only).
<b>fpsetmask</b> ( <i>mask</i> )	Set the floating-point exception mask for the calling process.
<b>fpgetsticky</b> ()	Get the floating-point exception sticky flags for the calling process (C only).
<b>fpsetsticky</b> ( <i>sticky</i> )	Set the floating-point exception sticky flags for the calling process (C only).

---

Paragon OSF/1 supports a series of floating-point control calls compatible with those of UNIX System V.

### NOTE

Only **fpsetmask()** is available to Fortran programs. The other floating-point control calls are available only to C programs.

## Detecting Not-a-Number

The calls `isnan()`, `isnand()`, and `isnanf()` are used to determine whether a floating-point value is an IEEE NaN, or “Not-a-Number.” This value can be generated as a result of certain floating-point mathematical operations and system calls, when the operands are invalid or out of range. `isnan()` and `isnand()` take an argument of type `double`, and `isnanf()` takes an argument of type `float`. (`isnan()` and `isnand()` are identical except for the name.) All three calls return 1 if the argument is a NaN, and 0 otherwise.

### NOTE

These calls never generate an exception, even if the argument is a NaN.

## Controlling Floating-Point Behavior

The calls `fpgetround()`, `fpsetround()`, `fpgetmask()`, `fpsetmask()`, `fpgetsticky()`, and `fpsetsticky()` get and set the i860 microprocessor's floating-point control registers. The values of these registers are part of the process, and are saved and restored when the process is swapped in and out.

The `get` calls simply return the current value of the specified register for the calling process; the `set` calls set the register to the specified value for the calling process and return its previous value.

## Rounding Mode

`fpgetround()` and `fpsetround()` get and set the i860 microprocessor's *floating-point rounding mode*, which determines what happens when a floating-point value generated in a calculation cannot be represented exactly.

The i860 microprocessor has four rounding modes:

<code>FP_RN</code>	Round to nearest representable number (if two representable numbers are equidistant, round to the even one).
<code>FP_RM</code>	Round toward minus infinity.
<code>FP_RP</code>	Round toward plus infinity.
<code>FP_RZ</code>	Round toward zero (truncate).

These symbolic names are the values of the `enum` type `fp_rnd`, which is declared in `<ieeefp.h>`. The argument of `fpsetround()` and the return values of `fpsetround()` and `fpgetround()` are of this type.

## NOTE

When you convert a floating-point value to an integer type in C, it always truncates. The processor's rounding mode is ignored.

## Exception Mask and Sticky Flags

**fpgetsticky()** and **fpsetsticky()** get and set the i860 microprocessor's *floating-point exception sticky flags*, and **fpgetmask()** and **fpsetmask()** get and set the *floating-point exception mask*.

The i860 microprocessor defines five floating-point exceptions:

<b>FP_X_INV</b>	Invalid operation exception.
<b>FP_X_DZ</b>	Divide-by-zero exception.
<b>FP_X_OFL</b>	Overflow exception.
<b>FP_X_UFL</b>	Underflow exception.
<b>FP_X_IMP</b>	Imprecise (loss of precision) exception.

These symbolic names are the values of the **enum** type **fp\_except**, which is declared in *<ieeefp.h>*. The arguments of **fpsetsticky()** and **fpsetmask()** and the return values of **fpgetsticky()**, **fpsetsticky()**, **fpgetmask()**, and **fpsetmask()** are of this type.

The i860 microprocessor has five *exception sticky flags* and five *exception mask bits* corresponding to the five exception types. When a floating-point exception occurs, the corresponding exception sticky flag is set to 1. The corresponding exception mask bit is then examined; if it is set to 1, the exception is *trapped* and the appropriate trap handler is called.

## NOTE

After an exception, you must clear the corresponding sticky flag to recover from the trap and proceed.

If the sticky flag is not cleared before the next floating-point exception occurs, an incorrect exception type may be signaled. For the same reason, when you call **fpsetmask()**, you must be sure that the sticky flag corresponding to each exception being enabled is cleared.



## NOTE

**fpsetsticky()** and **fpsetmask()** set the sticky flags and exception mask to the specified values. Any bits not set in the call's argument are cleared.

To set or clear a particular mask or sticky flag, get the current mask or sticky flags, modify it, and then call **fpsetsticky()** or **fpsetmask()** with the modified mask or sticky flags.

## Fortran Exception Mask Values

Only the **fpsetmask()** call is supported in Fortran. You use the following numeric values with **fpsetmask()**:

0	No exceptions.
1	Invalid operation exception.
2	Divide-by-zero exception.
4	Overflow exception.
8	Underflow exception.
16	Imprecise (loss of precision) exception.

The argument and return value of **fpsetmask()** are integers whose values are the sum of some, none, or all of these values.

## Miscellaneous Calls

---

Synopsis	Description
<b>flick()</b>	Temporarily relinquish the CPU to another process.
<b>dclock()</b>	Return time in seconds since booting the system.

---

### Temporarily Releasing Control of the Processor

The **flick()** call temporarily releases control of the node processor to another process in the same application. If there are no other processes in the same application when a process calls **flick()**, control returns to the Paragon OSF/1 operating system. For example, if your node program has set up a number of **hrecv()**'s and has nothing else to do, it should issue **flick()**. The operating system can then more efficiently respond to an incoming message and wake up your process.

**flick()** does not have any effect on rollin and rollout of the application (see "Gang Scheduling" on page 2-35 for information on rollin and rollout).

### Timing Execution

**dclock()** returns the time in seconds since the system was last booted, as a double precision number. This time is obtained from the RPM global clock and is the same on every node.

Use **dclock()** to return a relative value that you can use to measure execution time. To time an interval in your program, first obtain an initial value. Then obtain a final value and take the difference. The actual values returned by the two **dclock()** calls are not important.

Here is an example that shows how to use **dclock()** to time the execution of an iteration loop:

```

/* C version */
double start_time, end_time, diff_time;
start_time = dclock();
for(i=0;i<imax;i++) {
    .
    .
    .
}
end_time = dclock();
diff_time = end_time - start_time;
printf("Timing = %e\n", diff_time);

```

```
c Fortran version
      double precision start_time, end_time, diff_time
      start_time = dclock()
      do 100 i=1, imax
        .
        .
        .
100   continue
      end_time = dclock()
      diff_time = end_time - start_time
      write(*, 10) diff_time
10    format('diff_time = ', D15.9)
```

## iPSC® and Touchstone DELTA Compatibility Calls

---

Synopsis	Description
<b>flushmsg</b> ( <i>typesel, nodesel, ptypesel</i> )	Flush specified messages from the system.
<b>ginv</b> ( <i>j</i> )	Return the position of an element in the binary-reflected gray code sequence. Inverse of <b>gray</b> ( <i>j</i> ).
<b>gray</b> ( <i>j</i> )	Return the binary-reflected gray code for an integer.
<b>hwclock</b> ( <i>hwtime</i> )	Place the current value of the hardware counter into a 64-bit unsigned integer variable.
<b>infopid</b> ()	Return the process type of the process that sent a pending or received message.
<b>killcube</b> ( <i>node, ptype</i> )	Terminate and clear node process(es).
<b>killproc</b> ( <i>node, ptype</i> )	Terminate a node process.
<b>led</b> ( <i>state</i> )	Does nothing; provided for compatibility only.
<b>load</b> ( <i>filename, node, ptype</i> )	Load a node process.
<b>mclock</b> ()	Return the time in milliseconds.
<b>msgcancel</b> ( <i>mid</i> )	Cancel an asynchronous send or receive operation.
<b>mypart</b> ( <i>rows, cols</i> )	Obtain the height and width of the rectangle of nodes allocated to the current application.
<b>mypid</b> ()	Return the process type of the calling process.
<b>nodedim</b> ()	Return the dimension of the current application (the number of nodes allocated to the application is $2^{\text{dimension}}$ ).
<b>restrictvol</b> ( <i>fileID, nvol, volist</i> )	Does nothing; provided for compatibility only.

---

The calls **flushmsg**(*typesel, nodesel, ptypesel*), **ginv**(*j*), **gray**(*j*), **hwclock**(*hwtime*), **infopid**(*node, ptype*), **killcube**(*node, ptype*), **killproc**(*node, ptype*), **led**(*state*), **load**(*filename, node, ptype*), **mclock**(*node, ptype*), **msgcancel**(*mid*), **mypart**(*rows, cols*), **mypid**(*node, ptype*), **nodedim**(*node, ptype*), and **restrictvol**(*fileID, nvol, volist*) are provided for compatibility with the iPSC series of supercomputers and Touchstone DELTA system from Intel Corporation.

These calls should not be used in new Paragon OSF/1 programs. They either provide the same functionality as Paragon OSF/1 calls (for example, **mypid()** is identical to **mytype()** but uses the iPSC system terminology), or provide functionality that is not needed in Paragon OSF/1 (for example, **gray()** is not useful in a machine without a hypercube architecture).

These calls work the same as the corresponding calls on the iPSC or Touchstone DELTA system, with the following exceptions:

- **flushmsg()** does nothing.
- The only valid use of **killcube()** is **killcube(-1,-1)**.
- The only valid use of **killproc()** is **killproc(-1,-1)**.
- **led()** does nothing.
- **load()** must be preceded by **nx\_initve()** (it is equivalent to **nx\_load()** but does not let you specify a list of nodes or find out the PIDs of the loaded processes).
- **msgcancel()** does nothing.
- If **numnodes()** is not a power of 2, **nodedim()** rounds it up to the next power of 2 and returns the dimension of a cube of that size. For example, if **numnodes()** is 7, **nodedim()** returns 3; if **numnodes()** is 9, **nodedim()** returns 4.
- **restrictvol()** does nothing. It always returns 0 (indicating success).

See your iPSC or Touchstone DELTA system documentation for more information on these calls.



# Using Parallel File I/O

5

## Introduction

The Paragon™ OSF/1 operating system provides two forms of parallel I/O to files:

- A special file system type called *PFS*, for *Parallel File System*, gives applications high-speed access to a large amount of disk storage. PFS file systems are optimized for simultaneous access by multiple nodes. Files in PFS file systems can be very large (up to several terabytes); the exact maximum depends on your system configuration. Access to PFS file systems also uses an I/O technique called *fast path I/O*, which gives superior performance for large I/O operations (64K bytes or more per read or write).
- Special I/O system calls, called *parallel I/O calls*, facilitate I/O from multiple nodes and permit I/O to very large files in PFS file systems. These calls can give applications better performance and more control over parallel file I/O than is offered by the standard C and Fortran file I/O features. These calls are compatible with the Concurrent File System™ (CFS™) calls provided by the iPSC® system.

A system running Paragon OSF/1 can have both PFS and non-PFS file systems. You can access files in PFS file systems with both parallel I/O calls and non-parallel I/O calls; you can use parallel I/O calls to access files in both PFS file systems and non-PFS file systems. However, in most cases you get the best performance when you use parallel I/O calls to access files in PFS file systems.

This chapter discusses both PFS file systems and parallel I/O calls. It also gives information on performing operations on tape devices in Paragon OSF/1. For information on getting the best performance from PFS file systems and parallel I/O calls, see “I/O Performance” on page 8-23.

## Disks and File Systems

Every Intel supercomputer has one or more *disk devices* attached to it. Each disk device is either a single hard disk or a *RAID subsystem*. RAID stands for Redundant Array of Inexpensive Disks; in a RAID subsystem, several hard disks are connected together into a unit that appears to the system as a single large disk drive. Files stored to a RAID subsystem are distributed, or *striped*, among the disks within it by the RAID controller hardware.

Each disk device is controlled by an *I/O node*: a compute node with an I/O connection. I/O nodes communicate with the other nodes in the system using the node-to-node message-passing network and with the disk drives using a SCSI interface (or other interface). The I/O nodes may or may not also run application processes; this is determined by your system administrator. Each I/O node can control up to seven disk devices, and the number of I/O nodes is limited only by the number of slots in the system, so the total amount of disk space that could be installed in an Intel supercomputer is a terabyte or more.

The set of disk devices connected to the Intel supercomputer's I/O nodes is divided into *file systems*. A file system can encompass anything from a portion of the space on one disk device to all of the space on several disk devices. A file system is made accessible by *mounting* it to a directory (this requires system administrator privileges). This directory is called the file system's *mount point*. For example, if the file system `/dev/io0/rz0f` is mounted on the directory `/home` (the directory `/home` is the file system's mount point), whenever you write a file in `/home` it is stored in the file system `/dev/io0/rz0f`.

Each file system has a *type* that describes its internal structure and determines some of the operations that can be performed on it. The supported file system types are:

UFS	UNIX File System, the standard file system type for OSF/1.
NFS	Network File System, a file system type that represents a file system on another computer on the network.
PFS	Parallel File System, a file system type that is optimized for access by parallel processes. This file system type is unique to Paragon OSF/1.

This chapter discusses how PFS file systems work and how you can use the parallel I/O system calls provided by Paragon OSF/1 to access files in file systems of all types.



## PFS File Systems and PFS Files

Internally, a file system of type PFS consists of one or more *stripe directories*. The stripe directories that make up a PFS file system are determined by the system administrator when the PFS file system is mounted.

Each stripe directory is usually the mount point of a separate UFS file system. Just as a RAID subsystem collects together several hard disks into a unit that behaves like a single large disk, a PFS file system collects together several file systems into a unit that behaves like a single large file system. A system running Paragon OSF/1 can have any number of PFS file systems.

The maximum storage capacity of a PFS file system is the sum of the capacities of the different file systems containing its stripe directories. For example, if a PFS file system consists of four stripe directories, each of which is the mount point of a UFS file system with a capacity of 100M bytes, the capacity of the PFS file system is 400M bytes. However, if another PFS file system also consists of four stripe directories, but two of them are directories in one UFS file system with a capacity of 100M bytes and the other two are directories in another UFS file system with a capacity of 100M bytes, the capacity of the PFS file system is only 200M bytes.

A *PFS file* is any ordinary file that is stored in a file system of type PFS. PFS files are distributed, or *striped*, across the stripe directories that make up the PFS file system. The amount of data from a PFS file that is stored in each stripe directory is determined by the PFS file system's *stripe unit*, a quantity that is set by the system administrator when the PFS file system is mounted. The maximum size of a file in a PFS file system is roughly 2G bytes times the number of file systems in the PFS file system.<sup>1</sup>

For example, suppose a PFS file system consists of four stripe directories and has a stripe unit of 4K bytes. When you write a 20K-byte file to this PFS file system, the first 4K bytes of the file are stored in the first stripe directory, the second 4K bytes in the second stripe directory, the third 4K bytes in the third stripe directory, the fourth 4K bytes in the fourth stripe directory, and the last 4K bytes back in the first stripe directory.

Objects in PFS file systems that are not ordinary files (objects such as directories, symbolic links, and device special files) are not striped; each such object exists on just one disk.

## PFS Filenames and Pathnames

Filenames and pathnames in PFS file systems work the same as pathnames in UFS file systems. The maximum length of a pathname is 1024 characters; the maximum length of a single filename is 255 characters.

---

1. The exact maximum size is given by the formula  $((((2G - 1) - r) \times n) + r)$ , where  $r$  is  $(2G - 1) \bmod \text{stripe\_unit}$  (that is, the remainder when the largest integer multiple of the stripe unit that is less than  $2G - 1$  is subtracted from  $2G - 1$ ) and  $n$  is the number of *different* file systems containing the PFS file system's stripe directories.

## PFS Limitations

In the current release, PFS file systems and parallel I/O calls have the following limitations:

- PFS files cannot be accessed from a remote system via NFS.
- PFS does not support executable files. If you copy a binary file to a PFS file system and try to execute it, an “Operation not supported by this file system” error occurs.
- PFS does not support core files. If a core dump occurs while your current directory is in a PFS file system, a core file of length 0 is created.
- PFS does not support the **quotaon** or **sysacct** commands or the **mmap()** system call.
- PFS file regions cannot be locked by the **fcntl()** system call. However, you can use the **flock()** system call to lock the entire file.
- The maximum number of open files per process at any given time is 64. This includes the standard input, standard output, and standard error. This means that there is a practical maximum of 61 open files per process.

## Using PFS Commands

In general, you use standard OSF/1 commands such as **ls**, **cat**, **cp**, and **mv** to manipulate files in PFS file systems. See the *OSF/1 Command Reference* for information on these commands. (Many commands do not work with files larger than 2G – 1 bytes, as described under “Using Extended Files” on page 5-33.) This section describes the additional file and file system commands provided by Paragon OSF/1.

### Displaying File System Attributes

---

#### Command Synopsis

#### Description

**showfs** [ **-k** ] [ **-t type** ] [ *filesystem* | *directory* ] Display file system attributes.

---

The command **showfs** with no arguments lists the file systems on your system, together with information on each. For example:

```
% showfs
Mounted on          512-blks      avail  capacity  sunit  sfactor
/                   1458308      719276   45%
/home               4060838      3373782   8%
/usr                2379194      1948124   9%
/home/.sdirs/vol0   598622       574464   4%
/home/.sdirs/vol1   598622       574464   4%
/home/.sdirs/vol2   598622       574464   4%
/home/.sdirs/vol3   598622       574464   4%
/pfs                2394488      2297856   4%          8192      4
  sdirectories: /home/.sdirs/vol0
                /home/.sdirs/vol1
                /home/.sdirs/vol2
                /home/.sdirs/vol3
```

In this case, the system has eight file systems. The seven file systems mounted on the directories / (root), /home, /usr, /home/.sdirs/vol0, /home/.sdirs/vol1, /home/.sdirs/vol2, and /home/.sdirs/vol3 are non-parallel file systems (type UFS or NFS); the file system mounted on the directory /pfs is a PFS file system.

### NOTE

There's nothing special about the name /pfs; your PFS file systems can have any name. However, the rest of this chapter uses the convention that pathnames beginning with /pfs are in a PFS file system.

The **showfs** command shows the following information for every file system:

<b>Mounted on</b>	The directory where the file system is mounted (its <i>mount point</i> ). If you need to know the file system's device name, use the standard OSF/1 command <b>mount</b> or <b>df</b> .
<b>512-blks</b>	The total capacity of the file system in 512-byte disk blocks.
<b>avail</b>	The number of disk blocks currently available.
<b>capacity</b>	The approximate percentage of the file system's capacity currently in use.

In this example, the file system mounted on */usr* has a size of 2,379,194 512-byte disk blocks, of which 1,948,124 blocks are currently unused, so that the file system is approximately 9% full.

The **showfs** command shows the following additional information for each PFS file system:

<b>sunit</b>	The file system's stripe unit, in bytes.
<b>sfactor</b>	The number of stripe directories within the PFS file system.
<b>sdirectories</b>	The stripe directories (usually mount points of UFS file systems) within the PFS file system.

In this example, the PFS file system mounted on */pfs* has a stripe unit of 8K bytes and consists of the four UFS file systems mounted on */home/.sdirs/vol0*, */home/.sdirs/vol1*, */home/.sdirs/vol2*, and */home/.sdirs/vol3*.

The **showfs** command accepts the following optional arguments:

<b>-k</b>	Display capacity and available capacity in 1024-byte disk blocks instead of 512-byte disk blocks. The header "512-blks" changes to "kbytes".
<b>-t type</b>	Display information about all file systems of type <i>type</i> , where <i>type</i> is any recognized file system type in lowercase ( <b>pfs</b> , <b>ufs</b> , or <b>nfs</b> ).
<b>filesystem</b>	Display information about the file system whose device name is <i>filesystem</i> .
<b>directory</b>	Display information about the file system mounted on <i>directory</i> .

The *filesystem* or *directory* argument overrides **-t type** if used together.

## NOTE

You should use **showfs**, not **df**, to get information about the cumulative amount of free space in a PFS file system. Using the standard **df** command on a PFS file system only gives information about the single disk partition on which the PFS file system is mounted, so does not indicate how much space is actually available for file striping.

## Increasing the Size of a File

Command Synopsis	Description
<b>lsize</b> [ <b>-a</b> ] <i>size file</i> [ <i>file ...</i> ]	Change the size of a file or files.

The **lsize** command changes the amount of disk space allocated to each specified file. You can use this command to allocate all the space you will need for a large file before you run the application that writes to the file. This makes sure that there is enough room in the file system for the file, and can also increase file I/O performance.

The **lsize** command has two forms:

<b>lsize</b> <i>size file</i> [ <i>file ...</i> ]	Sets the size of the <i>file(s)</i> to <i>size</i> bytes.
<b>lsize -a</b> <i>size file</i> [ <i>file ...</i> ]	Increases the size of the <i>file(s)</i> by <i>size</i> bytes.

If the specified *file* does not exist, it is created with the specified *size*. The *size* can be a simple integer to represent a number of bytes, or an integer followed by the letter **k**, **m**, or **g** to represent a number of kilobytes (1024 bytes), megabytes (1024K bytes), or gigabytes (1024M bytes).

For example, the following command sets the size of the file *mydat* to 5M bytes:

```
% lsize 5m mydat
```

The following command increases the size of the file *mydat* by 200K bytes:

```
% lsize -a 200k mydat
```

The additional space is allocated to the file from the file system, but it is not initialized (its contents are undefined).

**lsize** will not decrease the size of a file. If the specified size is smaller than the file's current size, the command has no effect.

## Using Parallel I/O Calls

The rest of this chapter discusses the *parallel I/O calls* you can use in parallel applications to access both PFS and non-PFS files.

The term *parallel I/O calls* refers to all I/O calls that are provided by Paragon OSF/1 but not by standard OSF/1. These calls facilitate I/O on multiple nodes and permit I/O to very large files in PFS file systems. They are part of the library *libnx.a*, which is automatically searched when you link an application with the **-nx** switch. You can also use the switch **-lnx** to search *libnx.a* without using **-nx**. See “Compiling and Linking Applications” on page 2-5 for more information on these switches.

Most of the parallel I/O calls can only be used in programs running in the compute partition. They will not work, or will give unexpected results, if used in a program running in the service partition. (See “Managing Partitions” on page 2-25 for more information on the service and compute partitions.)

### NOTE

The parameter *fileID* in the system call synopses in this chapter is an integer that represents an open file: a *unit* in Fortran, or a *file descriptor* in C.

A call description at the beginning of each section or subsection gives a language-independent synopsis (call name, parameter names, and brief description) of each call discussed in that section. Differences between C and Fortran are noted where applicable. See Appendix A for information on call and parameter types; see the *Paragon™ C System Calls Reference Manual* or the *Paragon™ Fortran System Calls Reference Manual* for complete information on each call.

## Opening Files in Parallel

Synopsis	Description
<b>gopen</b> ( <i>path</i> , <i>oflag</i> , <i>iomode</i> [ , <i>perms</i> ] ) <b>gopen</b> ( <i>unit</i> , <i>path</i> , <i>iomode</i> )	(C) (Fortran)      Open a file on all nodes and set its I/O mode.

To open a file for use by all the nodes in your application, call **gopen()**. You can use **gopen()** to open files in both PFS and non-PFS file systems. **gopen()** works like the standard **open()** operation, with the following exceptions:

- It is a *global call*. All the nodes in the application must call **gopen()**, and all must call it with the same arguments.
- It is a *synchronizing call*. Each node blocks at the **gopen()** until all the nodes have called it.
- It sets the *I/O mode* of the file, as described under “Using I/O Modes” on page 5-13.
- When called on a large number of nodes, it offers better performance and causes less system overhead.

Note that **gopen()** must be called by *all* the nodes in the application, even those that do not actually perform any I/O. For example, suppose that your application has a “manager” node that assigns I/O work to the “worker” nodes, but does no I/O itself. If you want to use **gopen()**, all the nodes, even the manager, must open the file.

## Using `gopen()` in C

The C version of `gopen()` opens the specified file and returns a file descriptor, like the standard OSF/1 system call `open()`. In addition to being a global synchronizing call and setting the I/O mode of the file, as discussed earlier, the C version of `gopen()` has the following differences from the standard `open()`:

- It can only be used to open an ordinary file (not a directory or a device special file).
- If an error occurs, it prints an error message and terminates the calling process.

`gopen()` is otherwise equivalent to `open()`. For example, the following C call opens the file `/pfs/mydat` for reading and writing, creating it if it does not exist, and returns a file descriptor that you can use to access it. The file's I/O mode is set to `M_GLOBAL`, and if the file is created it is given permissions 644 octal (`rw-r--r--`).

```
#include <fcntl.h>
#include <nx.h>
int fd;
fd = gopen("/pfs/mydat", O_RDWR | O_CREAT, M_GLOBAL, 0644);
```

The symbolic names for *oflag* (such as `O_CREAT`) are defined in the header file `fcntl.h`, and the symbolic names for *iomode* (such as `M_GLOBAL`) are defined in the header file `nx.h`.

See `open()` in the *OSF/1 Programmer's Reference* for information on the *oflag* parameter; see "Using I/O Modes" on page 5-13 for information on the *iomode* parameter; see `chmod()` in the *OSF/1 Programmer's Reference* for information on the *perms* parameter.

## Using `gopen()` in Fortran

The Fortran version of `gopen()` opens the specified file for unformatted I/O on a specified unit. It is equivalent to the following Fortran `open()` statement:

```
OPEN(unit, path, status='unknown', form='unformatted',
x      access='sequential')
```

However, it differs from the standard Fortran `open()` in that it is a subroutine. Also, as discussed earlier, it is a global synchronizing call and sets the I/O mode of the file.

For example, the following Fortran call opens the file `/pfs/mydat` on unit 10 in I/O mode `M_GLOBAL`:

```
include 'fnx.h'
call gopen(10, "/pfs/mydat", M_GLOBAL)
```

The symbolic names for *iomode* (such as `M_GLOBAL`) are defined in the header file `fnx.h`.



## Opening Files with Standard Operations

PFS and non-PFS files can also be opened and closed with the standard OSF/1 system calls and Fortran routines. For example, to open the file */pfs/mydat* for read and write access:

```
/* C version */
fd = open("/pfs/mydat", O_CREAT | O_RDWR, 0644);

c      Fortran version
      open(unit=10, file = '/pfs/mydat',
x      status = 'new', form='unformatted')
```

Use this method when not all nodes open the same file at the same time, or when source compatibility with other systems is necessary. (Note that, if you want to use any synchronizing calls, all nodes must open the file.)

### NOTE

In Fortran, you must open the file with **form='unformatted'** in order to use any parallel I/O calls on the file.

The following section discusses additional special considerations for Fortran.

## Special Considerations for Fortran

This section describes the special considerations that apply when you open files with the standard Fortran **open()** instead of **gopen()**.

### Formatted Versus Unformatted I/O

If you call **open()** with **form='formatted'** (the default):

- You must use only Fortran I/O statements to access the file. You cannot use any of the parallel I/O calls described in this chapter on the file.
- Only one node may perform I/O to the file. If you perform formatted I/O to the same file from multiple nodes, the results are undefined.

If you open a file with **form='unformatted'**, you can use either Fortran I/O statements or parallel I/O calls to access the file. However, you must pick either one or the other: mixing Fortran I/O and parallel I/O to the same file can give unexpected results.

For the best I/O performance, you should use **gopen()**, or **open()** with **form='unformatted'**, and use parallel I/O calls for all file I/O.

If compatibility with other programs that use formatted I/O is required, you can perform formatted I/O to an internal file or a string and then use **cwrite()** to write the data to a file. However, if you use a string you must add a newline (ASCII character 10) to the end of the string using the function **char()**, since neither formatted I/O to a string nor **cwrite()** will add these for you. For example:

```
include 'fnx.h'
character*20 msgbuffer

write(msgbuffer, 26) answer, char(10)
26  format('The answer is: ', i4, a1)
call cwrite(10, msgbuffer, 20)
```

Alternatively, you can write a small program that translates your data files from unformatted to formatted and vice versa, and run it only when you need to share data with other programs.

## New Files

If you call **open()** with **status='new'**, the result depends on whether or not the program is running on multiple nodes:

- If the program is running on one node (**numnodes()** is 1 or undefined), the **open()** fails if the file exists, as specified by the ANSI standard.
- If the program is running on multiple nodes (**numnodes()** is greater than 1), the file is truncated if it exists, as though you had specified **status='unknown'**.

This change makes it possible to specify **status='new'** when multiple nodes are opening a file that does not yet exist; with the standard Fortran semantics for **status='new'**, the first node to execute the **open()** statement would create the file, and the other nodes would fail because the file already exists. You can use the system call **stat()** to determine if a file exists before you open it.

## Unnamed Files

If you call `open()` with no filename, the result depends on whether or not you specified `status='scratch'`:

- If you did not specify `status='scratch'`, the file is created in the current working directory with the filename `fort.nnn`, where `nnn` is the unit number. The file remains after the program terminates.
- If you specified `status='scratch'`, the file is created in the directory `/usr/tmp` with the filename `FTNxxxxxx.nn`, where `xxxxxx` is the OSF/1 process ID of the creating process and `nn` is the unit number. The file does not remain after the program terminates, whether it terminated normally or abnormally.

For compatibility with the iPSC system, if you specified `status='scratch'` and the directory specified by the variable `CFS_MOUNT` exists (or, if `CFS_MOUNT` is not defined, if the directory `/cfs` exists), the file `FTNxxxxxx.nn` is created in `$CFS_MOUNT` (or `/cfs`) instead of `/usr/tmp`.

## Using I/O Modes

Synopsis	Description
<code>setiomode(fileID, iomode)</code>	Set the I/O mode for a file.
<code>iomode(fileID)</code>	Return the current I/O mode for a file.

A parallel application can access a file in one of five I/O modes. You can specify a file's I/O mode when you open it with `gopen()`, and you can use `setiomode()` to change the I/O mode of a file that is already open. You can use `iomode()` to determine an open file's current mode.

Like `gopen()`, `setiomode()` is a global synchronizing call. When a node calls `setiomode()`, it blocks until all the other nodes in the application call `setiomode()` with the same arguments. `setiomode()` must be called by *all* the nodes in the application, even those that do not actually perform any I/O (this means that all nodes must open the file). Also, `setiomode()` can only be used on an ordinary file, not a directory or a device special file.

A file's I/O mode actually belongs to the file descriptor or unit through which the file is accessed, not to the file itself. The I/O mode is not stored with the file, and different programs can access the same file with different I/O modes (even at the same time).

A file's I/O mode is *not* inherited across a `fork()` (after a `fork()` all files in the child process have I/O mode `M_UNIX`).

There are five I/O modes, each of which has a name and a number:

- M\_UNIX (0)** In this mode, each node has its own file pointer and file operations are performed on a first-come, first-served basis. If you open a file with the C **open()** call or Fortran **open** statement, it is opened with this mode (but you can change it with **setiomode()**).
- M\_LOG (1)** In this mode, all nodes share the same file pointer and file operations are performed on a first-come, first-served basis.
- M\_SYNC (2)** In this mode, all nodes share the same file pointer and file operations are performed in order by node number. Records may be of variable length.
- M\_RECORD (3)** In this mode, each node has its own file pointer and file operations are performed on a first-come, first-served basis. However, records are stored in the file in order by node number. Records must be of a fixed length.
- M\_GLOBAL (4)** In this mode, all nodes share the same file pointer and must perform the same file operations at the same time. All file operations are performed by a single node, which then distributes the results to the other nodes over the internode network.

The names **M\_UNIX**, **M\_LOG**, **M\_SYNC**, **M\_RECORD**, and **M\_GLOBAL** are constants defined in the header files *nx.h* (for C) and *fix.h* (for Fortran). You can use either these names or the corresponding numbers in your programs (using the names is recommended).

The I/O mode you choose for a file determines which, if any, parallel I/O calls become *synchronizing operations* (that is, each node blocks until all nodes have made the call). The synchronizing operations for each mode are described in the following sections and are summarized under “Synchronization Summary” on page 5-48.

## **M\_UNIX (Mode 0)**

In mode **M\_UNIX** (number 0), each node maintains its own file pointer. File access requests are honored on a first-come, first-served basis. If two nodes write to the same place in the file, the second node overwrites the data written by the first node. This mode is the default for files opened with the standard **open** call or statement.

Use this mode in applications where each node performs I/O on disjoint segments of the file, or where I/O accesses are synchronized by other means (such as message-passing inherent to the application).

## M\_LOG (Mode 1)

In mode **M\_LOG** (number 1), all nodes share a single file pointer. File accesses are performed on a first-come, first-served basis. Whenever any node reads, writes, or moves the pointer, it affects the pointer position for all nodes. This may change the results of subsequent reads, writes, or moves by other nodes. This mode is useful for parallel log files.

Closing a file in this mode is a synchronizing operation. When a node closes a file, the operation blocks until all the other nodes also close the file.

## M\_SYNC (Mode 2)

In mode **M\_SYNC** (number 2), all nodes share a single file pointer and the nodes access the file in a synchronized round-robin fashion.

- All nodes share a single file pointer, as for **M\_LOG**.
- All the nodes in the application must open the file, and all must perform the same operations on the file in the same order. Reads and writes can be of variable sizes.
- All file operations are synchronizing.

Closing a file is a synchronizing operation, as for **M\_LOG**. In addition, reading, writing, seeking (using **lseek()**) and detecting end-of-file (using **iseof()**) become synchronizing operations—they block until all nodes have called them. For example, when a node reads from a file with the parallel I/O call **cread()**, the node blocks and the read request is not honored until all other nodes have called **cread()**.

- All reads and writes to the file are performed in order by node number.

For example, suppose node 3 in an application running on four nodes writes to a file with the parallel I/O call **cwrite()** before any of the other nodes. The node blocks until all the other nodes have called **cwrite()**. When all nodes have called **cwrite()**, the data from node 0 is written to the file first, followed by the data from node 1, then the data from node 2, and finally the data from node 3.

- The only valid use for **lseek()** is for all nodes to seek to the same position in the file. If nodes attempt to seek to different positions, an error occurs.

## M\_RECORD (Mode 3)

Mode **M\_RECORD** (number 3) gives results that are similar to **M\_SYNC**, but it operates more efficiently. However, **M\_RECORD** requires a fixed record size.

- Each node has its own file pointer, as for **M\_UNIX**.
- All the nodes in the application must open the file, and all must perform the same operations on the file in the same order, as for **M\_SYNC**.
- Corresponding reads and writes must be of the same size on all nodes.

When a node reads or writes to the file for the *n*th time, it must read or write the same number of bytes as the *n*th read or write by every other node. For example, if node 0 writes 100 bytes to the file with its first call to **cwrite()** and 50 bytes with its second call to **cwrite()**, then all nodes must write 100 bytes with their first call to **cwrite()** and 50 bytes with their second call to **cwrite()**.

### NOTE

No verification is performed. You must make sure that all the nodes in the application make the same calls and read and write the same number of bytes.

If different nodes read different amounts of data, incorrect data will be read. If different nodes write different amounts of data, the output of different nodes will overwrite each other and/or leave areas of the file with uninitialized data.

- All reads and writes *appear* to be performed in order by node number.

Because reads and writes are of a known length, the operating system on each node can determine where in the file it should be reading from or writing to independently of the other nodes. The results of reading or writing a file with **M\_RECORD** are the same as **M\_SYNC**, but **M\_RECORD** is more efficient because no synchronization is required. No seeking is required by the application; the file system automatically reads or writes file data to or from the proper offset in the file.

For example, suppose node 2 in an application running on four nodes writes a 10-byte record. Node 2's file pointer is first moved forward by 20 bytes to leave room for the records from nodes 0 and 1. Next, node 2's record is written to the file (which advances the file pointer by 10 bytes). Finally, node 2's file pointer is moved forward by 10 bytes to leave room for node 3's record. The other nodes can fill in their "slots" at any time (earlier or later); no synchronization or communication between nodes is required.

- Closing a file is a synchronizing operation, as for **M\_LOG** and **M\_SYNC**.
- As for **M\_SYNC**, **lseek()** becomes a synchronizing call, and the only valid use for **lseek()** is for all nodes to seek to the same position in the file. If nodes attempt to seek to different positions, an error occurs.

## M\_GLOBAL (Mode 4)

In mode **M\_GLOBAL** (number 4), all nodes must read and write the same data to the same parts of the file at the same time. This mode gives excellent performance for programs that work this way, such as a program where every node reads in the entire contents of a large input file.

- All nodes share a single file pointer.
- All the nodes in the application must open the file, and all must perform the same operations on the file at the same time.
- All file operations are synchronizing.
- Corresponding reads and writes must be of the same size on all nodes.
- The only valid use for **lseek()** is for all nodes to seek to the same position in the file.
- When the nodes write to a file, only the data written by a single node is actually written. Data written by other nodes is ignored.

The way that this mode is implemented is that only a single node actually reads from and writes to the disk. After a read, that node distributes the data to the other nodes over the internode network. This eliminates the contention for the disk device that would otherwise occur when many nodes attempt to read from the same place in a file at the same time.

## An I/O Mode Example

This section provides a small example program (in Fortran and C) that you can compile and execute to illustrate the differences between the various I/O modes. The source for this program can be found on the Intel supercomputer in */usr/share/examples/fortran/iomodes/iomodes.f* (Fortran version) or */usr/share/examples/ciomodes/iomodes.c* (C version).

The example program works as follows: node 0 gets an I/O mode from the user (specified as a number), and sends it to the other nodes. Then all nodes call **gopen()** to open the file *mydat* in the current directory (which could be in either a PFS file system or a non-PFS file system) with the specified I/O mode.

Each node then writes 10 records to the file. Each record contains the time in seconds since the file was opened, to four decimal places, and the message "Hello from node *x*." Node 0 waits one second before each write to the file; the other nodes write as fast as they can (this demonstrates how writes to the file are differently synchronized in the different modes). When each node finishes writing, it writes a "done writing" message to the screen. Then it closes the file and writes a "finished" message to the screen (the two messages show that, in some modes, `close()` is a synchronizing operation).

## Fortran Example

```

program iomodes

include 'fnx.h'

integer nunit, mode, iam
double precision start, now, loop_time, loop_start
character*16 msg
character*29 msgbuffer

msg = 'Hello from node '
nunit = 12
iam = mynode()

if(iam .eq. 0) then
  print *, 'Enter I/O mode (0, 1, 2, 3, or 4):'
  read(*, 11) mode
11  format(i1)
  call csend(1, mode, 4, -1, myptype())
else
  call crecv(1, mode, 4)
endif

call gopen(nunit, "mydat", mode)
print 13, iam, iomode(nunit)
13  format('Node ', i4, ' using mode ', i1)

start = dclock()
do 100 i = 1, 10
c   *** if node 0, do nothing for 1.0 seconds ***
  if(iam .eq. 0) then
    loop_start = dclock()
101  loop_time = dclock() - loop_start
    if (loop_time .lt. 1.0) goto 101
  endif

```



```

c      *** all nodes now write a record to the file ***
102     now = dclock() - start
        write(msgbuffer, 14) now, msg, iam, char(10)
14      format(f7.4, a17, i4, a1)
        call cwrite(nunit, msgbuffer, 29)
100    continue

        print 15, iam
15      format('Node ', i3, ' done writing')
        close(nunit)
        print 16, iam
16      format('Node ', i3, ' finished')
        end

```

## C Example

```

#include <fcntl.h>
#include <stdio.h>
#include <nx.h>

main()
{
    int    i, fd;
    double start, now;
    double loop_start, loop_cur;
    long   mode, iam;
    char   instr[40], msg[40];

    iam = mynode();

    if(iam == 0) {
        printf("Enter I/O mode (0, 1, 2, 3, or 4):\n");
        gets(instr);
        sscanf(instr, "%ld", &mode);
        csend(1, &mode, sizeof(mode), -1, mytype());
    } else {
        crecv(1, &mode, sizeof(mode));
    }

    fd = fopen("mydat", O_WRONLY | O_CREAT | O_TRUNC, mode, 0666);
    printf("Node %d using mode %d\n", iam, iomode(fd));

```

```

start = dclock();
for(i=0;i<10;i++) {
    if(iam==0) {
        loop_start = dclock();
        loop_cur = loop_start;
        while(loop_cur - loop_start < 1.0) {
            loop_cur = dclock();
        }
    }
    now = dclock() - start;
    sprintf(msg, "%7.4f Hello from node %4ld\n", now, iam);
    cwrite(fd, msg, strlen(msg));
}

printf("Node %d done writing\n", iam);
close(fd);
printf("Node %d finished\n", iam);
}

```

## Compiling and Running the Example

To compile this program to a parallel application, use the following **if77** or **icc** command:

```
% if77 -nx iomodes.f -o iomodes
```

or

```
% icc -nx iomodes.c -o iomodes
```

When you run the resulting application, you may find the output easier to understand if you run the example on four or fewer nodes. Use the **-sz** switch to determine the number of nodes on which the application runs (see “Controlling the Application’s Execution Characteristics” on page 2-13 for information on **-sz** and other application switches).

For example, to run the application on two nodes of your default partition with I/O mode 1 (**M\_LOG**):

```
% iomodes -sz 2
Enter I/O mode (0, 1, 2, 3, or 4):
1
Node 0 using mode 1
Node 1 using mode 1
Node 1 done writing
Node 0 done writing
Node 1 finished
Node 0 finished
%
```

The following example outputs came from the C version of the example, run on two nodes.

## **M\_UNIX Output**

In mode **M\_UNIX** (0), each node has its own file pointer. Node 1 finishes right away. Node 0 waits before each write and overwrites the message from node 1. As a result, the file contains only the writes from node 0.

```
1.0000 Hello from node 0
2.0087 Hello from node 0
.
.
.
9.0711 Hello from node 0
10.0797 Hello from node 0
```

## M\_LOG Output

In mode **M\_LOG** (1), the nodes share a common file pointer, but there is no synchronization. As in mode **M\_UNIX**, node 1 finishes right away; but this time, node 0 appends its data to the file rather than overwriting the data from node 1.

```

0.0000 Hello from node 1
0.0382 Hello from node 1
.
.
.
0.0990 Hello from node 1
0.1076 Hello from node 1
1.0000 Hello from node 0
2.0086 Hello from node 0
.
.
.
9.0712 Hello from node 0
10.0804 Hello from node 0

```

If the output file were large enough so that node 0 started before node 1 finished, the output of the two nodes would be interleaved in the middle of the file.

## M\_SYNC Output

In mode **M\_SYNC** (2), the nodes share a common file pointer, and there is synchronization. Nodes 1 and 0 finish at around the same time. Because node 1 waits for node 0 on each write, the writes are interleaved within the file.

```

1.0000 Hello from node 0
0.0000 Hello from node 1
2.0278 Hello from node 0
1.1105 Hello from node 1
.
.
.
9.2262 Hello from node 0
8.1641 Hello from node 1
10.2535 Hello from node 0
9.1914 Hello from node 1

```

Note that node 0's records appear *earlier* in the file than node 1's, but the time value shown for each record from node 0 is *later* than for the corresponding record from node 1. This is because the value shown is the time at which `cwrite()` was called, but node 1's record was not actually written to the file until node 0 had written its record.

In this case, node 1 called `cwrite()` for the first time immediately after opening the file, at time 0, but the `cwrite()` blocked and the record was not written to the file until after node 0 called `cwrite()` for the first time, at time 1.0000 (1.0000 seconds after the file was opened). Node 1 then called `cwrite()` for the second time, at time 1.1105, but that `cwrite()` again blocked until after node 0 called `cwrite()` again at time 2.0278, and so on.

## M\_RECORD Output

In mode `M_RECORD` (3), the nodes access the file in round-robin fashion, but there is no lock-step synchronization. Node 1 finishes first. Then, node 0 goes into the file and fills in its data in the correct places. Because the records are of a fixed length, node 0 has no trouble doing this. The result is that the records are in the same order as in mode `M_SYNC`, but node 1 did not spend any time waiting for node 0.

```

1.0000 Hello from node 0
0.0000 Hello from node 1
2.0208 Hello from node 0
0.0505 Hello from node 1
.
.
.
9.1637 Hello from node 0
0.1955 Hello from node 1
10.1841 Hello from node 0
0.2158 Hello from node 1

```

Note that node 1 finished in only 0.2158 seconds, without having to wait for node 0.

## M\_GLOBAL Output

In mode `M_GLOBAL` (4), writes by all nodes but one (node 0 in this case) are ignored. As a result, the file contains only the writes from that node.

```

1.0000 Hello from node 0
2.0087 Hello from node 0
.
.
.
9.0711 Hello from node 0
10.0797 Hello from node 0

```

This output is the same as the output of `M_UNIX`, but the other nodes do not compete with node 0 for access to the disk, so this mode is more efficient. However, because this program uses such a small data file, the difference in execution time is probably not noticeable.

Note that `M_GLOBAL` is usually used for reading, not writing.

## Reading and Writing Files in Parallel

You can read and write files with the familiar OSF/1 system calls and Fortran routines. For example, here is a Fortran code fragment that opens a file whose pathname is */pfs/mydat* and reads some data into an array called *array* using the Fortran **read** statement:

```
open(unit=10, file='/pfs/mydat', form='unformatted')
read 10, (array(j), j=1, n)
```

In addition to the usual I/O facilities, the Paragon OSF/1 operating system offers a series of parallel I/O calls, which are discussed in the following pages. These calls can be used on files in both PFS and non-PFS file systems.

Like the message-passing calls, the parallel I/O calls offer you the choice of *synchronous* or *asynchronous I/O*. The synchronous calls begin with **c** (for “complete”) and do not return until the operation is complete. The asynchronous calls begin with **i** (for “incomplete”) and return immediately; you use the call **iodone()** or **iowait()** to determine when the operation is complete.

If you program in Fortran, you should use the parallel I/O calls rather than Fortran I/O whenever you can. These calls offer better performance than the Fortran I/O routines, and you can test for the end of a file with **iseof()**. (This does not apply to C programmers; the usual C I/O calls are as efficient as their parallel I/O counterparts.) However, if you use parallel I/O calls on a file, you must not use Fortran file I/O statements on the same file (for example, you must not mix **write** and **cwrite()** on the same file).

### NOTE

Parallel I/O to NFS files may give poor performance or unexpected results.

The Intel supercomputer's disk I/O hardware and software are designed to support simultaneous access by large numbers of nodes. However, a remote NFS server may not be configured to support this level of access. If you perform large parallel I/O operations from large numbers of nodes to a file that is NFS-mounted from another computer, you may overload the network or the NFS server, resulting in poor performance or unexpected results.

## Synchronous File I/O

Synopsis	Description
<code>cread(fileID, buffer, nbytes)</code>	Read from a file, waiting for completion.
<code>cwrite(fileID, buffer, nbytes)</code>	Write to a file, waiting for completion.
<code>creadv(fileID, iov, iovcnt)</code>	Read from a file to irregularly-scattered buffers, waiting for completion.
<code>cwritev(fileID, iov, iovcnt)</code>	Write to a file from irregularly-scattered buffers, waiting for completion.

The calls `cread()`, `cwrite()`, `creadv()`, and `cwritev()` perform synchronous file I/O. They are equivalent to the standard OSF/1 calls `read()`, `write()`, `readv()`, and `writev()`, except that they follow the same naming and error-handling conventions as the Paragon OSF/1 message-passing calls (see “Names of Send and Receive Calls” on page 3-7 for information on the Paragon OSF/1 system call naming conventions; see “Handling Errors” on page 4-42 for information on the Paragon OSF/1 error-handling conventions). Unlike their standard OSF/1 equivalents, these calls are available to Fortran programs (as well as C).

For example, here is a C code fragment that writes the message “Hello from node *x*” to the file `/pfs/hello`:

```
fd = open("/pfs/hello", O_RDWR, 0644);
    .
    .
    .
printf(buffer, "Hello from node %d\n", iam);
cwrite(fd, buffer, strlen(buffer));
```

Here is a slightly more complicated example: a Fortran code fragment that opens a file whose pathname is `/pfs/mydat`, seeks to a location, and reads some data using the synchronous call `cread()`. The data represents a matrix stored in rows of *n* four-byte elements. Each node reads *m* rows and performs a calculation with each row (calling the Basic Linear Algebra Subroutines routine `sdot()` to get the dot product of two vectors). Because each node seeks to a different place in the file, you must use I/O mode `M_UNIX` (the default).

```
open(unit=10, file='/pfs/mydat', form='unformatted')
lseek(10, 4*mynode()*n*m, 0)

do 10 i = 1, m
    call cread(10, arow, n*4)
    y(i) = sdot(n, arow, 1, xtotal, 1)
10 continue
```

Note that when you open a file in Fortran, you must open it as sequential and unformatted to be able to use **cread()** and **cwrite()**. (Sequential is the default access, but you must specify **form='unformatted'**.)

## NOTE

Unlike their OSF/1 equivalents, these calls do not return the number of bytes read or written. If any error occurs, these calls print an error message and terminate the calling process.

Reading past the end of a file is considered an error, so you must be certain you know how many bytes remain in the file before you read from it. You can use **iseof()**, to detect end-of-file, after each **cread()** or **creadv()**. You can also use the following call to determine the length of a file:

```
length = lseek(unit, 0, SEEK_END)
```

This call sets the file pointer to the end of the file and returns the current position of the file pointer (that is, the file's length). You can then use **lseek(unit, 0, SEEK\_SET)** to return the file pointer to the beginning of the file. (If the file might be larger than 2G - 1 bytes, use **eseek()** instead of **lseek()**; see "Manipulating Extended Files" on page 5-36 for more information.)

If you need to detect errors in reading and writing, you must program in C and use either the standard OSF/1 calls (**read()**, **write()**, **readv()**, and **writev()**, described in the *OSF/1 Programmer's Reference*) or the underscore versions of the parallel I/O calls (**\_cread()**, **\_cwrite()**, **\_creadv()**, and **\_cwritev()**, described under "Handling Errors" on page 4-42). The underscore versions do return the number of bytes read or written.



## Asynchronous File I/O

---

Synopsis	Description
<b>iread</b> ( <i>fileID</i> , <i>buffer</i> , <i>nbytes</i> )	Read from a file without waiting for completion.
<b>iwrite</b> ( <i>fileID</i> , <i>buffer</i> , <i>nbytes</i> )	Write to a file without waiting for completion.
<b>ireadv</b> ( <i>fileID</i> , <i>iov</i> , <i>iovcnt</i> )	Read from a file to irregularly-scattered buffers, without waiting for completion.
<b>iwritev</b> ( <i>fileID</i> , <i>iov</i> , <i>iovcnt</i> )	Write to a file from irregularly-scattered buffers, without waiting for completion.
<b>iodone</b> ( <i>id</i> )	Determine whether an asynchronous I/O operation is complete. If complete, release the I/O ID.
<b>iowait</b> ( <i>id</i> )	Wait for completion of an asynchronous I/O operation and release the I/O ID.

---

The calls **iread()**, **iwrite()**, **ireadv()**, and **iwritev()** perform asynchronous file I/O. They work like **cread()**, **cwrite()**, **creadv()**, and **cwritev()**, but they return immediately, without waiting for the read or write to complete. The asynchronous I/O calls return an I/O ID much like the message ID returned by the asynchronous message passing calls. You can pass this I/O ID to **iodone()** or **iowait()** to determine when the asynchronous file I/O operation has completed.

### NOTE

The number of I/O IDs is limited, so you *must* use **iodone()** or **iowait()** to release each ID after you use it.

To check if an asynchronous I/O operation has completed, use the `iodone()` call. It returns 1 if the asynchronous operation has completed and 0 otherwise. You can also decide to block on the completion of an asynchronous call. Use the `iowait()` call for this. Both `iodone()` and `iowait()` take the I/O ID as an input parameter. For example (in Fortran):

```

c   Write to a file
      ioid = iwrite(12, sbuf, size)
      .
      .
      .
c   Do some calculation...
      .
      .
      .
c   Wait until the write completes
      call iowait(ioid)

```

The number of available I/O IDs is limited; be sure to release IDs that are no longer needed. There are two ways to release an I/O ID: you can issue an `iowait()`, as shown in the previous example, or you can keep issuing `iodone()`s until an `iodone()` returns 1.

## NOTE

To preserve data integrity, all I/O requests that use or affect the file pointer are processed on a "first-in, first-out" basis.

This means that if an asynchronous I/O call is followed by a synchronous read, write, or seek on the same file, the synchronous call will block until the asynchronous operation has completed.

## Closing Files in Parallel

It's always a good idea to close a file when you are finished using it. Whether you used `open()` or `gopen()` to open a file, and whether the file is a PFS file or a non-PFS file, you use the standard OSF/1 system calls or Fortran routines to close it.

For example, to close the file open on file descriptor `fd` (C) or unit 10 (Fortran):

```

/* C version */
close(fd);

c   Fortran version
      close(unit=10)

```

**NOTE**

If the I/O mode of the file being closed is anything other than **M\_UNIX**, closing the file is a synchronizing operation.

See "Using I/O Modes" on page 5-13 for more information.

## Detecting End-of-File and Moving the File Pointer

Synopsis	Description
<code>iseof(fileID)</code>	Test for end-of-file.
<code>lseek(fileID, offset, whence)</code>	Move the read/write file pointer.

The calls `iseof()` and `lseek()` are provided for both C and Fortran programmers. If you use parallel I/O calls to perform file I/O in a Fortran program, you must use `iseof()` and `lseek()` instead of the equivalent Fortran features.

The `iseof()` call returns 1 if the given file is at the end of the file and 0 otherwise. For example, the following Fortran code reads characters from the file open on unit 12, writing each one to the screen, until it reaches the end of the file:

```

do while(iseof(12) .eq. 0)
  call cread(12, char, 1)
  print 300, iam, char
300  format('Node ', i3, ' read: ', a1)
end do

```

The `lseek()` call moves the file pointer to *offset* bytes from the point specified by *whence*, which can be either a name or a number:

- If *whence* is **SEEK\_SET**, `lseek()` moves the pointer to *offset* bytes from the beginning of the file.
- If *whence* is **SEEK\_CUR**, `lseek()` moves the pointer forward *offset* bytes from its current position.
- If *whence* is **SEEK\_END**, `lseek()` moves the pointer to *offset* bytes after the end of the file.

The names **SEEK\_SET**, **SEEK\_CUR**, and **SEEK\_END** are constants defined in the header files `unistd.h` (for C) and `fnx.h` (for Fortran). For compatibility with the iPSC system, the numeric values 0, 1, and 2 are also accepted (but using the symbolic names is recommended).

**lseek()** returns the new position of the file pointer (measured in bytes from the beginning of the file).

For example, the following C call moves the file pointer of the file open on file descriptor *fd* to the beginning of the file:

```
#include <unistd.h>

newpos = lseek(fd, 0, SEEK_SET);
```

The following Fortran call moves the file pointer of the file open on unit 12 forward 500 bytes:

```
include 'fnx.h'

newpos = lseek(12, 500, SEEK_CUR)
```

## NOTE

If the I/O mode of the file is **M\_SYNC**, **M\_RECORD**, or **M\_GLOBAL**, seeking is a synchronizing operation.

See “Using I/O Modes” on page 5-13 for more information.

## Flushing Fortran Buffered I/O

Synopsis	Description
<b>forceflush()</b>	Cause all buffered I/O to be flushed if an exception occurs.
<b>forflush(<i>unit</i>)</b>	Flush all buffered I/O on a particular unit.

The subroutines **forceflush()** and **forflush()** let Fortran programmers make sure that buffered I/O actually goes to the associated file or device. These subroutines are not available to C programs.

Fortran I/O to files and devices other than the user's terminal is *buffered*—that is, when you write to a file, the data is stored in a memory buffer, and only written to the corresponding file or device when the buffer is full. However, if another node is waiting for some data to appear in a file, you might

want to force the contents of a unit's buffer to be written immediately. You can do this by calling **forflush()** on the unit. For example, to flush all buffered I/O on unit 9 to the corresponding file or device:

```
call forflush(9)
```

Another possible problem with buffered I/O is that if the program is interrupted by an exception, buffered data that has not yet been written to the file is lost. The subroutine **forceflush()** establishes a signal handler that flushes all buffered I/O in case of an exception. You call it as follows:

```
call forceflush
```

Note that you must call **forceflush()** *before* the exception occurs. You can use **fpsetmask()** (described under "Controlling Floating-Point Behavior" on page 4-46) to control whether or not an exception occurs in case of certain floating-point errors.

Fortran I/O to the user's terminal is not buffered. You can avoid buffering to files and devices by using parallel file I/O calls such as **cwrite()** and **iwrite()** instead of Fortran I/O. These calls do not buffer I/O into the Fortran I/O memory buffer; when the call returns, you can be sure the data has been sent to the specified file or device. (However, there may be some buffering within the operating system, which cannot be avoided.)

## Using "###" Filenames

If you perform certain standard file operations (shown in Table 5-1) on a file that contains three or more # symbols in its filename, the series of # symbols is automatically replaced by the node number (within the application) of the node that opens the file.

**Table 5-1. File Operations that Accept "###" Filenames**

<b>access()</b>	<b>mknod()</b>	<b>truncate()</b>
<b>chdir()</b>	<b>open()</b>	<b>unlink()</b>
<b>chmod()</b>	<b>readlink()</b>	<b>utimes()</b>
<b>chown()</b>	<b>rmdir()</b>	<b>link()</b>
<b>creat()</b>	<b>stat()</b>	<b>rename()</b>
<b>mkdir()</b>	<b>statfs()</b>	<b>symlink()</b>

The Fortran equivalents of these operations also support "###" filenames. Note, though, that **gopen()** does *not* appear in this list.

For example, assume that you have the same program running on all the nodes of your application, and each node calls **open()** to open a file called *file###*. The result is that each node opens a separate file. Node 0 opens *file000*, node 1 opens *file001*, node 2 opens *file002*, and so on. If an application opens *file###* for reading, the specified files (*file000*, *file001*, *file002*, and so on) must exist.

If you use a “###” filename in a non-parallel program running in the service partition, it uses node number 0. For example, opening a file called *file###* from a service node opens *file000*. Note that this also affects standard commands that make these calls; for example, since the **rm** command calls **unlink()**, the command **rm file###** will attempt to remove the file *file000*.

Filenames containing a sequence of one or two # symbols are not affected. For example, the file *file##* is a single file that is accessible by each node.

If the number of digits in a node number is less than the number of # symbols in the filename, the node number is padded with zeros to the length of the sequence of # symbols. If the number of digits in a node number exceeds the number of # symbols in the filename, the filename is extended, but only when necessary. For example, calling **unlink()** on the file *data.###* in every node of an application running on 2000 nodes unlinks files *data.000*, *data.001*, *data.002* ... *data.998*, *data.999*, *data.1000*, *data.1001* ... *data.1998*, and *data.1999*.

There is nothing special about files created in this way; each file created is a single ordinary file. For example, suppose an application uses **open()** or **creat()** to create *###myfile*, writes into it, and then closes the file. This creates a series of files called *000myfile*, *001myfile*, *002myfile*, and so on. Each of these files is an ordinary file; for example, you can delete one without affecting the others, and there's nothing to prevent node 1 from opening *005myfile*.

## Increasing the Size of a File

Synopsis	Description
<b>lsize</b> ( <i>fileID</i> , <i>offset</i> , <i>whence</i> )	Increase size of a file.

You can allocate more space to a file with **lsize()**. The **lsize()** call sets the file's size as specified by *offset* and *whence*:

- If *whence* is **SIZE\_SET**, **lsize()** sets the file's size to *offset* bytes.
- If *whence* is **SIZE\_CUR**, **lsize()** sets the file's size to the current file pointer position plus *offset* bytes.
- If *whence* is **SIZE\_END**, **lsize()** increases the file's size by *offset* bytes.

The names **SIZE\_SET**, **SIZE\_CUR**, and **SIZE\_END** are constants defined in the header files *nx.h* (for C) and *fnx.h* (for Fortran). For compatibility with the iPSC system, the numeric values 0, 1, and 2 are also accepted (but using the symbolic names is recommended).

For example, the following Fortran call increases the size of the file open on *unit1* to one million bytes:

```
include 'fnx.h'  
  
size = lsize(unit1, 1000000, SIZE_SET)
```

The following C call increases the size of the file open on file descriptor *fd* by 500,000 bytes:

```
#include <unistd.h>  
#include <nx.h>  
int size, fd;  
  
size = lsize(fd, 500000, SIZE_CUR)
```

The additional space is allocated to the file from the file system, but it is not initialized (its contents are undefined).

**lsize()** will not decrease the size of a file. If the size specified by *offset* and *whence* is smaller than the file's current size, the call has no effect.

The major use of this call is to ensure that enough disk space is available before you begin a lengthy calculation. Pre-allocating disk space can also improve disk performance.

## Using Extended Files

A PFS file greater than or equal to 2G bytes in size is called an *extended file*. These files are stored in the same way as non-extended PFS files. However, some of the file parameters (like the file pointer and file size) do not fit into a 32-bit integer. This means that standard OSF/1 calls and commands that use these parameters cannot be used on extended files. The following two sections list the calls and commands that do not support extended files.

## OSF/1 Calls that Do Not Support Extended Files

Most OSF/1 calls, such as `read()` and `write()`, don't care how big the file is and work perfectly well on extended files. The OSF/1 calls that have problems with extended files are shown in Table 5-2.

**Table 5-2. OSF/1 Calls Not Supporting Extended Files**

Call	Problem
<code>fcntl()</code>	Can't lock a file region larger than 2G – 1 bytes.
<code>fgetpos()</code>	Can't return an offset greater than 2G – 1 bytes.
<code>fseek()</code>	Can't specify an offset greater than 2G – 1 bytes.
<code>unlocked_fseek()</code>	Can't specify an offset greater than 2G – 1 bytes.
<code>fsetpos()</code>	Can't return an offset greater than 2G – 1 bytes.
<code>fstat()</code>	Can't be used on a file larger than 2G – 1 bytes. <sup>1</sup>
<code>ftell()</code>	Can't return an offset greater than 2G – 1 bytes.
<code>ftruncate()</code>	Can't specify a file size greater than 2G – 1 bytes.
<code>lseek()</code>	Can't specify an offset greater than 2G – 1 bytes.
<code>lstat()</code>	Can't be used on a file larger than 2G – 1 bytes. <sup>1</sup>
<code>madvise()</code>	Can't map a file larger than 2G – 1 bytes.
<code>mmap()</code>	Can't map a file larger than 2G – 1 bytes.
<code>mprotect()</code>	Can't map a file larger than 2G – 1 bytes.
<code>msync()</code>	Can't map a file larger than 2G – 1 bytes.
<code>munmap()</code>	Can't map a file larger than 2G – 1 bytes.
<code>stat()</code>	Can't be used on a file larger than 2G – 1 bytes. <sup>1</sup>
<code>truncate()</code>	Can't specify a file size greater than 2G – 1 bytes.

1. If you call `fstat()`, `lstat()`, or `stat()` on a file larger than 2G – 1 bytes, the call fails with the error `EFBIG`.

To manipulate extended files, Paragon OSF/1 provides special calls that perform the equivalent of `lseek()`, `stat()`, `fstat()`, and `lsize()` for extended files. These are discussed under “Manipulating Extended Files” on page 5-36.



## OSF/1 Commands that Do Not Support Extended Files

Many OSF/1 commands make one or more of the system calls in Table 5-2, so do not work on extended files. The commands **chgrp**, **chmod**, **chown**, **cp**, **ls**, **mv**, **tar**, and **rm** have been specially modified to support extended files; most other commands will fail if used on extended files. (Note that you must use the **-E** switch to archive an extended file with **tar**; see **tar** in the *Paragon™ Commands Reference Manual* for more information.)

Table 5-3 shows the OSF/1 commands that are known to have problems with extended files. (This list is not guaranteed to be complete; other commands, not listed here, may also have problems.)

**Table 5-3. OSF/1 Commands Not Supporting Extended Files**

Command	Problem
<b>cat</b>	Can't read a file larger than 2G – 1 bytes.
<b>compress</b>	Can't compress a file larger than 2G – 1 bytes.
<b>cpio</b>	Can't handle files or archives larger than 2G – 1 bytes.
<b>diff</b>	Can't compare a file larger than 2G – 1 bytes.
<b>du</b>	Can't show the size of a directory containing files larger than 2G – 1 bytes.
<b>ed</b>	Can't edit a file larger than 2G – 1 bytes.
<b>ex</b>	Can't edit a file larger than 2G – 1 bytes.
<b>find</b>	Can't find a file larger than 2G – 1 bytes with <b>-size</b> ; can't show the size of a file larger than 2G – 1 bytes with <b>-ls</b> .
<b>ftp</b>	Can't copy a file larger than 2G – 1 bytes.
<b>more</b>	Can't display a file larger than 2G – 1 bytes.
<b>newfs</b>	Can't create a UFS file system larger than 2G – 1 bytes.
<b>rcp</b> <sup>1</sup>	Can't copy a file larger than 2G – 1 bytes.
<b>tail</b>	Can't display a file larger than 2G – 1 bytes.
<b>vi</b>	Can't edit a file larger than 2G – 1 bytes.

1. Note that although **rcp** cannot copy an extended file, **cp** can.

## Manipulating Extended Files

Synopsis		Description
<b>eseek</b> ( <i>fildes, offset, whence</i> )	(C)	Move file pointer in extended file.
<b>eseek</b> ( <i>unit, offset, whence, newpos</i> )	(Fortran)	
<b>esize</b> ( <i>fildes, offset, whence</i> )	(C)	Increase size of extended file.
<b>esize</b> ( <i>unit, offset, whence, newsize</i> )	(Fortran)	
<b>estat</b> ( <i>path, buffer</i> )	(C only)	Get status of extended file from pathname.
<b>lstat</b> ( <i>path, buffer</i> )	(C only)	Get status of extended file or symbolic link from pathname.
<b>festat</b> ( <i>fildes, buffer</i> )	(C only)	Get status of open extended file from file descriptor.

The **e...**() calls perform file operations on extended files. They do this by having parameters that are *extended integers* (a data type capable of representing integers greater than  $2G - 1$ ). You must use the calls described under “Performing Extended Arithmetic” on page 5-37 to operate on extended integers.

- The call **eseek**() is like **lseek**() (discussed under “Detecting End-of-File and Moving the File Pointer” on page 5-29), except that the *offset* parameter is an extended integer. The C version of this call is a function that returns the new position as an extended integer; the Fortran version is a subroutine that stores the new position in its fourth parameter.
- The call **esize**() is like **lsize**() (discussed under “Increasing the Size of a File” on page 5-32), except that the *offset* parameter is an extended integer. The C version of this call is a function that returns the new size as an extended integer; the Fortran version is a subroutine that stores the new size in its fourth parameter.
- The calls **estat**(), **lstat**(), and **festat**() are like the standard OSF/1 calls **stat**(), **lstat**(), and **fstat**() (described in the *OSF/1 Programmer's Reference*), except that they use a structure called *estat*, defined in `<sys/estat.h>`, which is the same as the OSF/1 *stat* structure except that the file size is an extended integer. These calls are available only in C, not in Fortran.

You *must* use these calls to manipulate extended files (PFS files greater than or equal to 2G bytes in size). However, you can also use these calls on non-PFS files and on PFS files less than 2G bytes in size. You can use these calls or the standard OSF/1 calls on PFS files less than 2G bytes in size.

## Performing Extended Arithmetic

Synopsis		Description
<code>eadd(e1, e2)</code>	(C)	Add two extended integers.
<code>eadd(e1, e2, eresult)</code>	(Fortran)	
<code>ecmp(e1, e2)</code>		Compare two extended integers.
<code>ediv(e, n)</code>	(C)	Divide extended integer by integer.
<code>ediv(e, n, result)</code>	(Fortran)	
<code>emod(e, n)</code>	(C)	Give extended integer modulo an integer (remainder when <i>e</i> is divided by <i>n</i> ).
<code>emod(e, n, result)</code>	(Fortran)	
<code>emul(e, n)</code>	(C)	Multiply extended integer by integer.
<code>emul(e, n, eresult)</code>	(Fortran)	
<code>esub(e1, e2)</code>	(C)	Subtract two extended integers.
<code>esub(e1, e2, eresult)</code>	(Fortran)	
<code>etos(e, s)</code>		Convert extended integer to string.
<code>stoe(s)</code>	(C)	Convert string to extended integer.
<code>stoe(s, e)</code>	(Fortran)	

The extended arithmetic calls manipulate 64-bit integers, also called *extended integers*. You use these calls to manipulate the parameters used by the parallel I/O calls described in the previous section.

Extended integers are signed 64-bit integers with values from  $(2^{63} - 1)$  to  $-2^{63}$  ( $2^{63}$  is approximately  $9.2 \times 10^{18}$ ).

- In Fortran, extended integers are stored in a two-element array of type `integer*4`.
- In C, extended integers are stored in a variable of type `esize_t`, a structure type defined in the header file `<sys/estat.h>`. (For compatibility with the iPSC system, there is also a header file `<estat.h>` that simply includes `<sys/estat.h>`.)

You should always use extended arithmetic calls to operate on an extended integer, rather than access its internal structure.

Some of these calls return extended integers. The C versions of these calls return a value of type *esize\_t*. However, Fortran does not allow functions to return arrays, so the Fortran versions of these calls are subroutines with an additional parameter: the result of the operation on the first two parameters is stored into the third parameter. For example, the following call adds the extended integers *e1* and *e2* and stores the result in *e\_sum*:

```

/* C version */
#include <sys/estat.h>
esize_t e1, e2, e_sum;
e_sum = eadd(e1, e2);

c      Fortran version
      integer e1(2), e2(2), e_sum(2)
      call eadd(e1, e2, e_sum);

```

If you want to add an ordinary integer to an extended integer, you must create your own extended integer from the desired integer value. To create an extended integer, use *stoe()*. This call takes a string whose value is a number, and returns the corresponding numeric value as an extended integer. For example, the following code fragment adds 1 to the value of the extended integer *e1*. It does this by converting the string "1" to an extended integer with *stoe()*, storing the resulting extended integer in *e2*, and then adding *e2* to *e1* (note that in Fortran the string must be declared to be one character larger than the actual string being converted):

```

/* C version */
#include <sys/estat.h>
esize_t e1, e2, e_sum;
char *one = "1";

e2 = stoe(one);
e_sum = eadd(e1, e2);

c      Fortran version
      character*2 one
      parameter (one = '1')
      integer e1(2), e2(2), e_sum(2)

      call stoe(one, e2)
      call eadd(e1, e2, e_sum)

```

The other extended arithmetic calls allow you to subtract, multiply, divide, and find the remainder after division of extended integers. When you use *ediv()* or *emod()*, the divisor and answer must be 4-byte integers, not extended integers. Similarly, when you use *emul()*, the second argument must be a 4-byte integer, not an extended integer.

You can also compare two extended integers; *ecmp()* returns -1, 0, or 1, depending on whether the first extended integer is less than, equal to, or greater than the second.

## Getting Information About PFS File Systems

Synopsis	Description
<b>getpfsinfo</b> ( <i>buf</i> )	Get PFS-specific information about all mounted PFS file systems.
<b>statpfs</b> ( <i>path</i> , <i>fs_buffer</i> , <i>pfs_buffer</i> , <i>pfs_bufsize</i> )	Get PFS-specific and non-PFS-specific information for the file system containing <i>path</i> .
<b>fstatpfs</b> ( <i>fildev</i> , <i>fs_buffer</i> , <i>pfs_buffer</i> , <i>pfs_bufsize</i> )	Get PFS-specific and non-PFS-specific information for the file system containing the file open on <i>fildev</i> .

The functions **getpfsinfo**(), **statpfs**(), and **fstatpfs**() let C programmers get information about PFS file systems. These functions are not available to Fortran programs. See “PFS File Systems and PFS Files” on page 5-3 for more information on the concepts discussed in this section.

## Getting Information About All Mounted PFS File Systems

**getpfsinfo**() gets information about all mounted PFS file systems. It is similar to the standard OSF/1 call **getmntinfo**(), except that instead of returning information in an array of *statfs* structures, it returns information in an array of *pfsmntinfo* structures. It allocates the memory for this array of structures, each of which describes one PFS file system, and stores a pointer to this array into its argument. **getpfsinfo**() returns the number of elements in this array. The *pfsmntinfo* structure, defined in the header file *<pfs/pfs.h>*, contains the following fields:

<i>m_mntonname</i>	Directory on which the PFS file system is mounted.
<i>m_statpfs</i>	<i>statpfs</i> structure that describes the PFS file system.

The *statpfs* structure, also defined in the header file *<pfs/pfs.h>*, describes the PFS-specific attributes of a file system. This is a variable-size structure. It contains the following fields:

<i>p_reclen</i>	Total size of this <i>statpfs</i> structure, in bytes.
<i>p_sunitsize</i>	Stripe unit size for this PFS file system, in bytes.
<i>p_sfactor</i>	Number of stripe directories within this PFS file system.
<i>p_sdirs</i>	List of stripe directories within this PFS file system. The number of pathnames in the list is specified by <i>p_sfactor</i> .

Each pathname in *p\_sdirs* is a structure of type *pathname\_t* (defined in *<pfs/pfs.h>*); you can use the **NEXTPATH()** macro defined in *<pfs/pfs.h>* to examine each pathname in turn.

Here's an example of **getpfsinfo()**:

```
#include <sys/types.h>
#include <nx.h>
#include <pfs/pfs.h>

main() {
    struct pfsmntinfo *pfsinfo;
    struct statpfs    *sattr;
    pathname_t        *sdir;
    int               cnt, i, incr;

    cnt = getpfsinfo(&pfsinfo);

    if(cnt == 0) {
        printf("No PFS file systems mounted\n");
    } else {
        for(i = 0; i < cnt; i++) {
            printf("Mount point: %s\n", pfsinfo->m_mntonname);

            sattr = &(pfsinfo->m_statpfs);
            printf("  Stripe unit size: %d\n",
                sattr->p_sunitsize);
            printf("  Stripe factor: %d\n", sattr->p_sfactor);

            sdir = &(sattr->p_sdirs);
            printf("  Stripe directories:\n");
            for(i = 0; i < sattr->p_sfactor; i++) {
                printf("    %s\n", sdir->name);
                sdir = NEXTPATH(sdir);
            }

            incr = sizeof(pfsinfo->m_mntonname)
                + sattr->p_reclen;
            pfsinfo = (struct pfsmntinfo *)((char *)pfsinfo
                + incr);
        }
    }
}
```

This program prints out the attributes of all mounted PFS file systems, something like the command **shows -t pfs**. Note that you must use the **NEXTPATH()** macro to step through the *p\_sdirs* field of the *statpfs* structure, and you must increment the pointer into the array of *pfsmntinfo* structures by the size of the current *pfsmntinfo* structure (using the value of its *p\_reclen* field).

## Getting PFS Information About a Single File System

**statpfs()** gets information about a file system given the pathname of a file or directory in that file system; **fstatpfs()** gets information about a file system given the file descriptor of an open file in that file system.

These functions get both general and PFS-specific information about the specified file system. They can be used on both PFS and non-PFS file systems, but they return PFS-specific information only for PFS file systems. They are similar to the standard OSF/1 calls **statfs()** and **fstatfs()**, except that instead of returning information in a *statfs* structure, they return information in an *estatfs* structure and a *statpfs* structure.

- The *estatfs* structure, defined in the header file `<pfs/pfs.h>`, describes the basic attributes of the file system. It is just like the *statfs* structure defined in `<sys/mount.h>`, except that some of its fields are of type *esize\_t* (see “Performing Extended Arithmetic” on page 5-37 for information on this type). This is necessary because some of the values returned for PFS file systems are too large to be stored into an ordinary integer.

Some of the more generally useful fields of the *estatfs* structure are:

<i>f_type</i>	The type of the file system, expressed as a constant such as <b>MOUNT_UFS</b> , <b>MOUNT_NFS</b> , or <b>MOUNT_PFS</b> (these constants are defined in <code>&lt;sys/mount.h&gt;</code> ).
<i>f_bavail</i>	Number of free 1024-byte disk blocks in the file system available to ordinary users, expressed as a value of type <i>esize_t</i> .
<i>f_mntonname</i>	Directory on which the file system is mounted, expressed as a string.
<i>f_mntfromname</i>	Device name of the file system, expressed as a string.

See **statpfs()** in the *Paragon™ C System Calls Reference Manual* for a complete description of all fields in the *estatfs* structure.

- The *statpfs* structure is the same *statpfs* structure described for **getpfsinfo()** in the previous section. However, the way it is returned is different: **getpfsinfo()** allocates space for several *statpfs* structures and returns you a pointer to this space, but **statpfs()** and **fstatpfs()** store information in a *statpfs* structure that you provide.

Because the *statpfs* structure is variable-size, you must tell **statpfs()** and **fstatpfs()** how big your *statpfs* structure is; you do this with the third parameter of **statpfs()** and **fstatpfs()** (called *pfs\_bufsize*). Then you must check the *p\_reclen* field in the returned *statpfs* structure to be sure the returned information fit in your provided structure; if it didn't, try again with a larger structure.

Here's an example of `statpfs()`:

```
#include <sys/types.h>
#include <sys/mount.h>
#include <malloc.h>
#include <nx.h>
#include <pfs/pfs.h>

#define SDIRS_INIT_SIZE 1024

main(int argc, char **argv) {
    struct statpfs *statpfsbuf;
    int          bufsize;
    struct estatfs estatbuf;
    pathname_t   *sdir;
    char         blocks[80];
    int          i;

    if(argc != 2)
    {
        printf("Usage: %s <mountpoint>\n", argv[0]);
        exit(1);
    }

    bufsize=sizeof(struct statpfs) + SDIRS_INIT_SIZE;

    statpfsbuf=(struct statpfs *)malloc(bufsize);

    if(statpfs(argv[1], &estatbuf, statpfsbuf, bufsize) < 0)
    {
        nx_perror("statpfs");
        exit(1);
    }

    if(statpfsbuf->p_reclen > bufsize)
    {
        bufsize=statpfsbuf->p_reclen;
        statpfsbuf=(struct statpfs *)realloc(statpfsbuf,
                                             bufsize);

        if(statpfs(argv[1], &estatbuf, statpfsbuf, bufsize)
           < 0)
        {
            nx_perror("statpfs");
            exit(1);
        }
    }
}
```



```

printf("Selected PFS statistics for %s:\n", argv[1]);

/* From estatfs structure */

printf(" File system type: %d\n", estatbuf.f_type);
etos(estatbuf.f_bavail, blocks);
printf(" # of 1K blocks available: %s\n", blocks);
printf(" Mount point: %s\n", estatbuf.f_mntonname);
printf(" Device name: %s\n", estatbuf.f_mntfromname);

/* From statpfs structure */

printf(" Stripe unit size: %d\n",
       statpfsbuf->p_sunitsize);
printf(" Stripe factor: %d\n", statpfsbuf->p_sfactor);

printf(" Stripe directories:\n");
sdir = &(statpfsbuf->p_sdirs);
for (i = 0; i < statpfsbuf->p_sfactor; i++) {
    printf("    %s\n", sdir->name);
    sdir = NEXTPATH(sdir);
}
}

```

This program prints out the attributes of the file system containing the file specified by its first argument. Note that you must allocate enough space for the *statpfs* structure plus the stripe directory pathnames and check the returned *p\_reclen* against the currently-allocated size of the structure (*bufsize*).

This example starts off by allocating an extra **SDIRS\_INIT\_SIZE** bytes (an arbitrary value) for the stripe directory pathnames. If *p\_reclen* is larger than the size of the structure, this example uses **realloc()** to enlarge the structure and calls **statpfs()** again. It then uses the **NEXTPATH()** macro to step through the *p\_sdirs* field of the *statpfs* structure, as discussed earlier for **getpfsinfo()**.

## Controlling Tape Devices

---

### Synopsis

**ioctl**(*fd, request, argp*)

### Description

Perform an operation on an open tape or other device.

---

You can use standard OSF/1 I/O calls or parallel I/O calls to open, read, and write tape devices. To control tape devices, use the standard OSF/1 system call **ioctl**(). The header file `<sys/mtio.h>` defines the tape-specific structures and constants you need.

### NOTE

Only one node at a time can open a tape device, and it must use I/O mode **M\_UNIX** (0).

`<sys/mtio.h>` defines two constants you can use as the second argument of **ioctl**():

**MTIOCTOP** Perform operation on tape.

**MTIOCGET** Get status of tape.

The rest of this section explains the details of using these constants.

## Naming Tape Devices

The Paragon OSF/1 operating system uses the following conventions for naming tape devices:

<code>/dev/ioN/rmtX</code>	Raw cartridge tape, rewinds automatically when closed.
<code>/dev/ioN/nrmtX</code>	Raw cartridge tape, does not rewind automatically when closed.
<code>/dev/ioN/rmtcX</code>	Raw cartridge tape with compression, rewinds automatically when closed.
<code>/dev/ioN/nrmtcX</code>	Raw cartridge tape with compression, does not rewind automatically when closed.

## NOTE

The *rmtc* devices can only be used with tape drives that support data compression.

In each case, *N* is the node number of the I/O node to which the tape device is connected, and *X* is the SCSI ID of the tape device (typically 6). So, for example, to use the cartridge tape device with SCSI ID 6 on the boot node (node 0) and have it rewind automatically when closed, use the pathname */dev/io0/rmt6*. To use the same device but have it *not* rewind automatically when closed, use the pathname */dev/io0/nrmt6*.

## Performing Operations on Tape Devices

When you call `ioctl()` with `MTIOCTOP` as its second argument, you must use a structure of type `mtop` as the third argument. The `mtop` structure is defined as follows:

```
struct mtop {
    short  mt_op;    /* operation to perform */
    short  fill;    /* ignored */
    long   mt_count; /* how many operations to perform */
};
```

This structure tells `ioctl()` what operation to perform. The valid values of the `mt_op` field include the following constants:

<b>MTWEOF</b>	Write <i>mt_count</i> end-of-file marks.
<b>MTFSF</b>	Space the tape forward by <i>mt_count</i> files.
<b>MTBSF</b>	Space the tape backward by <i>mt_count</i> files.
<b>MTFSR</b>	Space the tape forward by <i>mt_count</i> records.
<b>MTBSR</b>	Space the tape backward by <i>mt_count</i> records.
<b>MTREW</b>	Rewind the tape. If the tape has been written to, writes two end-of-file marks before rewinding. (Two end-of-file marks indicate the end of data.)
<b>MTOFFL</b>	Rewind the tape and put the drive offline. If the tape has been written to, writes two end-of-file marks before rewinding.
<b>MTNOP</b>	No operation, sets status only.
<b>MTRETEN</b>	Retension the tape.

**MTERASE** Erase the entire tape.

**MTEOM** Position the tape at end of media (SCSI only).

Closing the tape device after writing to it also writes an end-of-file mark (or two end-of-file marks if the tape was opened in variable-block mode or the tape mode “rewind” is set). If the tape was opened in variable-block mode, the tape head is then positioned between the two end-of-file marks, so that any subsequent write will overwrite the second one.

For example, the following C program rewinds the tape on the device connected to */dev/io0/rmt6*:

```
#include <fcntl.h>
#include <errno.h>
#include <sys/mtio.h>

main() {
    int fd;
    struct mtop s;

    fd = open("/dev/io0/rmt6", O_RDONLY, 0666);
    if(fd == -1) {
        perror("opening /dev/io0/rmt6");
        exit(1);
    }

    s.mt_op = MTREW;
    s.mt_count = 1;
    if (ioctl(fd, MTIOCTOP, &s) == -1) {
        perror("rewinding tape");
        exit(2);
    }
}
```

## Getting Status of Tape Devices

When you call `ioctl()` with `MTIOCGET` as its second argument, you must provide a structure of type `mtget` as the third argument. The `mtget` structure is defined as follows:

```
struct mtget {
    short  mt_type; /* type of magtape device */
    short  mt_dsreg; /* ``drive status'' register */
    short  mt_erreg; /* ``error'' register */
    short  mt_resid; /* residual count */
};
```

**ioctl()** fills in the elements of this structure with information about the device. The value of the *mt\_type* field is always **0x0C** (indicating a generic SCSI device). The values of the *mt\_dsreg* and *mt\_erreg* fields are device-dependent.

For example, the following C program prints the status of the device connected to */dev/io0/rmt6*:

```
#include <fcntl.h>
#include <errno.h>
#include <sys/mtio.h>

main() {
    int fd;
    struct mtget s;

    fd = open("/dev/io0/rmt6", O_RDONLY, 0666);
    if(fd == -1) {
        perror("opening /dev/io0/rmt6");
        exit(1);
    }

    if (ioctl(fd, MTIOCGET, &s) == -1) {
        perror("getting status of tape");
        exit(2);
    }

    printf("mt_type = 0x%x\n", s.mt_type);
    printf("mt_dsreg = 0x%x\n", s.mt_dsreg);
    printf("mt_erreg = 0x%x\n", s.mt_erreg);
    printf("mt_resid = 0x%x\n", s.mt_resid);
}
```

## Synchronization Summary

Table 5-4 lists the I/O modes and summarizes the I/O calls that are synchronizing calls in each one. Table 5-5 lists the most commonly-used I/O calls and summarizes the I/O modes that cause them to become synchronizing calls.

**Table 5-4. Synchronization in Each I/O Mode**

<b>I/O Mode</b>	<b>I/O Calls that Synchronize</b>
<b>M_UNIX</b>	<b>gopen() and setiomode()</b>
<b>M_LOG</b>	<b>gopen(), setiomode() and close()</b>
<b>M_SYNC</b>	All
<b>M_RECORD</b>	<b>gopen(), setiomode(), lseek(), eseek(), and close()</b>
<b>M_GLOBAL</b>	All

**Table 5-5. File I/O Calls that Synchronize**

<b>Call</b>	<b>I/O Modes Causing the Call to Synchronize</b>
<b>close()</b>	<b>M_LOG, M_SYNC, M_RECORD, and M_GLOBAL</b>
<b>cread() and creadv()</b>	<b>M_SYNC and M_GLOBAL</b>
<b>cwrite() and cwritev()</b>	<b>M_SYNC and M_GLOBAL</b>
<b>eseek()</b>	<b>M_SYNC, M_RECORD, and M_GLOBAL</b>
<b>gopen()</b>	All
<b>iread() and ireadv()</b>	<b>M_SYNC and M_GLOBAL</b>
<b>iseof()</b>	<b>M_SYNC and M_GLOBAL</b>
<b>iwrite() and iwritev()</b>	<b>M_SYNC and M_GLOBAL</b>
<b>lseek()</b>	<b>M_SYNC, M_RECORD, and M_GLOBAL</b>
<b>setiomode()</b>	All

# Using Pthreads



6

## Introduction

In Paragon OSF/1, each process consists of a set of *resources*, such as memory objects and open files, and one or more *threads* (short for *threads of control*). Each thread consists of an instruction pointer and a stack.

By default, each process has only one thread. When there is more than one thread in a single process, each thread executes independently, but they share resources. For example, all the threads in a single process share memory; when one thread writes to a variable in memory, it modifies the value of that variable for all threads.

Because threads share memory, you must carefully coordinate access to shared areas of memory. For example, if two threads write to the same area of memory at the same time, the results may be indeterminate. You can use arbitration mechanisms such as *mutexes* (short for “mutual exclusion”) to protect areas of memory from being accessed at the same time. Software that uses these techniques and can be safely executed by two or more threads at the same time is referred to as *thread-safe* or *reentrant*.

## The Pthreads Package

To create and manage threads, you should use the *pthread*s package. *Pthreads* is short for *POSIX threads*; the threads created and managed by this package are also referred to as pthreads.

The current Paragon OSF/1 implementation of pthreads is based on the *POSIX Threads Extension [C language] P1003.4a/D4 (Draft 4), August 1990*; however, it is not strictly conformant to this draft. Also, note that this is not the most recent draft of this extension, so Paragon OSF/1 pthreads programs may not be portable to or from other systems. Future versions of the Paragon OSF/1 pthreads package may be based on later drafts of this extension, and may not be compatible with the current version. In particular, the treatment of thread cancellation and signals may be different in the future.

The pthreads package consists of the following libraries:

- |                      |  |
|----------------------|--|
| <i>libpthreads.a</i> | Contains thread management calls, such as <b>pthread_create()</b> . The calls in this library are discussed in “Using Pthreads Library Calls” on page 6-11.                                    |
| <i>libc_r.a</i>      | Contains reentrant versions of standard C library ( <i>libc.a</i> ) calls, such as <b>printf()</b> . The calls in this library are discussed in “Using Reentrant C Library Calls” on page 6-6. |

The only programming interface to the pthreads package provided in the current release is for the C programming language. No Fortran interface is provided.

## NOTE

Pthreads are *not* the same as, and are *not* compatible with, “cthreads” or “Mach threads.”

These other types of thread are not supported for use in user programs. These types of thread are not compatible with *libc\_r.a* and cannot be recognized or managed by the calls in *libpthreads.a*.

## What's In This Chapter

This chapter introduces the Paragon OSF/1 pthread package and its usage in parallel applications. It includes the following sections:

- Limitations of pthreads.
- Recommended safe operating environment.
- Compiling and linking a pthread application.
- Using reentrant C library calls.
- Using pthreads library calls.
- Interfacing with non-thread-safe code.
- Message passing and pthreads library calls.
- File I/O and pthreads library calls.
- **nx\_nfork()** and **nx\_initve()** and pthreads library calls.



- Signals and pthreads library calls.
- Handling errors.

## Limitations of Pthreads

The pthreads package has the following limitations in the current release:

- Currently, none of the libraries supplied with Paragon OSF/1 (except for *libpthreads.a* and *libc\_r.a*) are thread-safe. In particular, the library *libn.x.a*, which contains all the calls discussed in the other chapters of this manual, is not thread-safe. In a process containing multiple pthreads, any calls to non-thread-safe libraries must be protected so that no two pthreads can call the same library at the same time. See “Interfacing with Non-Thread-Safe Code” on page 6-37 for information on how to do this.
- Any global variables used or set by a non-thread-safe library may also have to be protected. For example, if a non-thread-safe function sets the global variable *errno*, you must be sure to read the value of *errno* before allowing any other pthread to make any call that could change the value of *errno*. See “errno Confusion” on page 6-41 for more information on *errno*.
- In the current implementation, all the pthreads in a process always run on the same processor. Scheduling of pthreads is handled by the kernel, which uses a policy of time sharing with aging. You cannot control, or get information about, pthread scheduling by using pthread library calls.
- Pthreads use kernel resources as well as user-level resources (to be specific, each pthread uses one *kernel thread*). This means that using very large numbers of pthreads can exhaust certain resources within the kernel.
- The *POSIX Threads Extension [C language] P1003.4a/D4 (Draft 4), August 1990* includes an optional feature called “thread priority scheduling.” This feature is not available in the current release. If you attempt to make use of this feature, you will get compilation errors (for use of an unsupported data type), link errors (for use of an unsupported library call), or run-time errors (for use of an unsupported system call). If a run-time error occurs, the call fails with the *errno* value *ENOSYS*.
- The Paragon application development tools currently are not thread-aware and do not have any features to support pthreads.
- There is no Fortran interface to the pthreads package. If you must use pthreads in a Fortran program, you could make the calls to the pthreads library from a C function, which can then be compiled to a *.o* file and linked into the Fortran program. However, this programming model has not been tested.

## Recommended Safe Operating Environment

The previous section described the limitations which *cannot* be exceeded. This section recommends limitations which you *should not* exceed in the current release. Exceeding these limitations may result in unexpected behavior, up to and including system crashes and data loss.

- No process should have more than 6 pthreads at once.

Any pthread can create or terminate another pthread at any time. The system does not impose a limit on the number of active pthreads in a process, on a node, or in the whole system. However, the total active pthreads per process (including the main thread) should be kept at or below 6. Exceeding this limit may result in an emulator exception.

- Only one pthread in a process should use the message-passing calls described in Chapter 3. The message-passing pthread can be the main thread or another pthread, but a pthread other than the main thread will experience higher message latency than the main thread.

This limitation is due to the fact that the message-passing library (*libnx.a*) is not thread-safe. Also, there is no mechanism in current message passing calls to send or receive messages to or from a specific pthread within a process. See “Message Passing and Pthreads Library Calls” on page 6-37 for more information.

If more than one pthread in a process attempts to perform message passing, message-passing performance may degrade, incorrect information may be returned from an **info...()** call, and global operations such as **gsync()** may give unexpected results.

- All global operations (such as **gsync()**) must be performed by the message-passing pthread. This is necessary because all global operations use message-passing to synchronize the nodes. You can synchronize pthreads *within a process* by using a global variable counter as a barrier.
- Only one pthread in a process should use the parallel file I/O calls described in Chapter 5. The I/O pthread can be the main thread or another pthread. See “File I/O and Pthreads Library Calls” on page 6-38 for more information.
- The calls **gopen()** and **setiomode()** use message-passing internally. If the I/O pthread is not also the message-passing pthread, you must make sure that these calls are not used at the same time as any message-passing calls in the message-passing pthread.
- The standard OSF/1 file I/O system calls, such as **read()** and **write()**, can be called from multiple pthreads at the same time *if* they are called from a controlling process and they are only used with files that reside in UFS file systems. Otherwise, only one pthread in a process can use them.
- Applications with multiple pthreads should not be run in gang-scheduled partitions. Gang scheduling of threaded applications has not been thoroughly tested, but has been known to cause server exceptions.

- Do not call `sigwait()` to wait on synchronous signals (those that are generated synchronously as the results of a pthread's faults, such as `SIGBUS` and `SIGSEGV`). Doing this may cause the application or the system to crash. See "Managing Signals" on page 6-34 for more information on `sigwait()`.
- Do not use calls from `libpthreads.a` or `libc_r.a` within a signal handler. Some of these calls use mutexes internally, which may result in deadlock (the handler can find itself waiting on an unavailable mutex lock, while the mutex lock cannot be released until the signal handler has returned).
- Asynchronous cancellation is very destructive and should be avoided. In particular, attempting to cancel a pthread doing file I/O on a PFS or UFS file system can cause the entire application to hang. See "Canceling Pthreads" on page 6-28 for more information on asynchronous cancellation.
- Do not attempt to use the Paragon application development tools on applications with multiple pthreads. In particular, IPD currently is not thread-aware and should not be used to debug or profile an application that uses pthreads.
- Do not terminate a process when other pthreads are progressing. For example, calling `exit()` or returning from `main()` kills all threads and terminates the entire process. If there are any other pthreads in the process, including pthreads generated transparently by library calls, null processes may result. Be sure to terminate all pthreads gracefully before terminating the program. In particular, be sure that all asynchronous and interrupt-driven message and I/O operations (such as `hrecv()` or `iread()`) are complete before the program terminates. It is also a good idea to ensure that `main()` always exits by calling `pthread_exit()`, never by calling `exit()` or by reaching the closing brace of `main()`.
- No Mach calls can be used in a pthreads program. The Mach kernel interface (`libmach.a`) is not supported in the current release; use of Mach features in pthreads programs can cause pthreads internal errors or system crashes.

## Compiling and Linking a Pthread Application

When compiling a program that uses the pthreads package, you should define the symbol `_REENTRANT`; this symbol ensures that thread-safe definitions are used in all included header files. (The compiler switch `-M[no]reentrant` does *not* have any effect on whether or not the resulting code is thread-safe. It only determines whether or not the code can be called recursively.)

When linking a program that uses the pthreads package, you must link in the library `libpthreads.a`, followed by the library `libc_r.a`. The standard C library is also linked by default, but only after searching all libraries specified on the command line. If you use the `-nx` switch, it can appear on the command line either before these two libraries or after them; if you use the `-lnx` switch, it should appear after both these libraries.

For example:

- To compile and link a non-parallel program:

```
% cc -D_REENTRANT -o node node.c -lpthreads -lc_r
```

- To compile and link a controlling process:

```
% cc -D_REENTRANT -o node node.c -lpthreads -lc_r -lnx
```

- To compile and link a parallel application:

```
% cc -D_REENTRANT -o node node.c -lpthreads -lc_r -nx
```

## Using Reentrant C Library Calls

Only the calls in the reentrant C library (*libc\_r.a*) and the calls in the pthreads library (*libpthreads.a*) are guaranteed to be thread-safe in the current release. Any calls to other libraries must be protected so that no two pthreads can call them at the same time. Table 6-1 lists the calls in *libc\_r.a*.

Table 6-1. Calls in Reentrant C Library (*libc\_r.a*) (1 of 2)

abort()	endgrent()	fstatfs()	getsockname()	memcpy()
abs()	endpwent()	fsync()	getsockopt()	memmove()
accept()	endttyent()	ftell()	gettimeofday()	memset()
access()	endusershell()	ftruncate()	gettimer()	mkdir()
acct()	endutent()	ftw()	getttyent_r()*	mknod()
adjtime()	exec_with_loader()	funlockfile()*	getttynam_r()*	mkstemp()
advance()	execlp()	fwrite()	getuid()	mktemp()
alarm()	execvp()	gcvt()	getusershell_r()*	mktime()
alloca()	exit()	getaddressconf()	getutent_r()*	mktimer()
asctime_r()*	fabs()	getc()	getutid_r()*	mmap()
async_daemon()	fchdir()	getchar()	getutline_r()*	modf()
atexit()	fchmod()	getclock()	getw()	mount()
atof()	fchown()	getcwd()	getwc()	mprotect()
atoi()	fclose()	getdirentries()	getwd()	msem_init()
atol()	fcntl()	getdtablesize()	gmtime()	msem_lock()
bcmp()	fcvt_r()*	getegid()	gmtime_r()*	msem_remove()
bcopy()	fdopen()	getenv()	htonl()	msem_unlock()
bind()	feof()	geteuid()	htons()	msgctl()
brk()	ferror()	getfh()	initstate()	msgget()
bzero()	fflush()	getfsent_r()*	initstate_r()*	msgrcv()
calloc()	ffs()	getfsfile_r()*	insque()	msgsnd()
catclose()	fgetc()	getfsspec_r()*	ioctl()	msync()
catgets()	fgets()	getfssstat()	isalnum()	munmap()
catopen()	fileno()	getgid()	isatty()	mvalid()
chdir()	flock()	getgrent_r()*	isdigit()	nfssvc()
chmod()	flockfile()*	getgrgid_r()*	isnan()	nice()
chown()	fopen()	getgrnam_r()*	isnand()	nl_langinfo()
chroot()	fpathconf()	getgroups()	isnanf()	ntohl()
clearerr()	fpgetmask()	gethostid()	ispace()	ntohs()
clock()	fpgetround()	gethostname()	isupper()	open()
close()	fpgetsticky()	getitimer()	isxdigit()	opendir()
closedir()	fprintf()	getlogin_r()*	kill()	pathconf()
connect()	fpsetmask()	getpagesize()	ldexp()	perror()
creat()	fpsetround()	getpeername()	link()	pipe()
ctermid()	fpsetsticky()	getpgrp()	listen()	plock()
ctime_r()*	fputc()	getpid()	localtime_r()*	poll()
cuserid()	fputs()	getppid()	longjmp()	printf()
dbm_close()	fread()	getpriority()	lseek()	profil()
dbm_fetch()	free()	getpwent_r()*	lstat()	ptrace()
dbm_open()	freopen()	getpwnam_r()*	madvise()	putc()
dup()	frexp()	getpwuid_r()*	malloc()	putchar()
dup2()	fscanf()	getrlimit()	memcpy()	puts()
ecvt_r()*	fseek()	getrusage()	memchr()	pututline_r()*
endfsent()	fstat()	gets()	memcmp()	putw()

\* Does not exist in the standard C library (*libc.a*).

Table 6-1. Calls in Reentrant C Library (*libc\_r.a*) (2 of 2)

<b>putc()</b>	<b>semctl()</b>	<b>settyent()</b>	<b>strchr()</b>	<b>ulimit()</b>
<b>quotactl()</b>	<b>semget()</b>	<b>setuid()</b>	<b>strcmp()</b>	<b>umask()</b>
<b>raise()</b>	<b>semop()</b>	<b>setusershell()</b>	<b>strcpy()</b>	<b>umount()</b>
<b>rand()</b>	<b>send()</b>	<b>setutent()</b>	<b>strcspn()</b>	<b>uname()</b>
<b>rand_r()*</b>	<b>sendmsg()</b>	<b>setvbuf()</b>	<b>strdup()</b>	<b>ungetc()</b>
<b>random()</b>	<b>sendto()</b>	<b>shmat()</b>	<b>strerror_r()*</b>	<b>unlink()</b>
<b>random_r()*</b>	<b>setbuf()</b>	<b>shmctl()</b>	<b>strftime()</b>	<b>unlocked_fclose()*</b>
<b>re_comp_r()*</b>	<b>setbuffer()</b>	<b>shmdt()</b>	<b>string()</b>	<b>unlocked_fflush()*</b>
<b>re_exec_r()*</b>	<b>setclock()</b>	<b>shmget()</b>	<b>strlen()</b>	<b>unlocked_fread()*</b>
<b>read()</b>	<b>setfsent()</b>	<b>shutdown()</b>	<b>strncat()</b>	<b>unlocked_fseek()*</b>
<b>readdir()</b>	<b>setgid()</b>	<b>sigaction()</b>	<b>strncmp()</b>	<b>unlocked_fwrote()*</b>
<b>readdir_r()*</b>	<b>setgrent()</b>	<b>sigaddset()</b>	<b>strncpy()</b>	<b>unlocked_getc()*</b>
<b>readlink()</b>	<b>setgroups()</b>	<b>sigdelset()</b>	<b>strpbrk()</b>	<b>unlocked_getchar()*</b>
<b>readv()</b>	<b>sethostid()</b>	<b>sigemptyset()</b>	<b>strrchr()</b>	<b>unlocked_getwc()*</b>
<b>realloc()</b>	<b>sethostname()</b>	<b>sigfillset()</b>	<b>strspn()</b>	<b>unlocked_putc()*</b>
<b>reboot()</b>	<b>setitimer()</b>	<b>sigismember()</b>	<b>strtok_r()*</b>	<b>unlocked_putchar()*</b>
<b>recv()</b>	<b>setjmp()</b>	<b>signal()</b>	<b>strtol()</b>	<b>unlocked_setvbuf()*</b>
<b>recvfrom()</b>	<b>setlinebuf()</b>	<b>sigprocmask()</b>	<b>strtoul()</b>	<b>utimes()</b>
<b>recvmsg()</b>	<b>setlocale()</b>	<b>sigreturn()</b>	<b>swapon()</b>	<b>utmpname()</b>
<b>reltimer()</b>	<b>setlogin()</b>	<b>sigstack()</b>	<b>symlink()</b>	<b>vfork()</b>
<b>remque()</b>	<b>setpgid()</b>	<b>sigsuspend()</b>	<b>sync()</b>	<b>vfprintf()</b>
<b>rename()</b>	<b>setpgrp()</b>	<b>sleep()</b>	<b>sysconf()</b>	<b>vprintf()</b>
<b>revoke()</b>	<b>setpriority()</b>	<b>socket()</b>	<b>table()</b>	<b>vsprintf()</b>
<b>rewind()</b>	<b>setpwent()</b>	<b>socketpair()</b>	<b>tempnam()</b>	<b>wait4()</b>
<b>rewinddir()</b>	<b>setregid()</b>	<b>sprintf()</b>	<b>time()</b>	<b>waitpid()</b>
<b>rindex()</b>	<b>setreuid()</b>	<b>srand()</b>	<b>times()</b>	<b>write()</b>
<b>rmdir()</b>	<b>setrlimit()</b>	<b>srandom()</b>	<b>tmpfile()</b>	<b>writev()</b>
<b>rmknod()</b>	<b>setsid()</b>	<b>srandom_r()*</b>	<b>tmpnam()</b>	
<b>rmtimer()</b>	<b>setsockopt()</b>	<b>sscanf()</b>	<b>tolower()</b>	
<b>sbrk()</b>	<b>setstate()</b>	<b>stat()</b>	<b>truncate()</b>	
<b>scanf()</b>	<b>setstate_r()*</b>	<b>statfs()</b>	<b>ttyname_r()*</b>	
<b>select()</b>	<b>settimeofday()</b>	<b>strcat()</b>	<b>tzset()</b>	

\* Does not exist in the standard C library (*libc.a*).

The calls in *libc\_r.a* can be divided into three groups, according to their names:

- Most of the calls in *libc\_r.a* have the same names as calls in the standard C library (*libc.a*). These calls generally work the same as the equivalent calls in *libc.a*. However, they perform special checks and locks internally to be sure they will work if called by multiple pthreads at the same time. Also, many blocking system calls only block the calling pthread instead of every pthread of the process. The following commonly-used calls have the following effects in programs with multiple pthreads:

<b>exit()</b>	Kills all pthreads of the process and closes all opened files. See “Calling exit()” on page 6-42 for more information on using <b>exit()</b> in programs with multiple threads.
<b>fork()</b>	Copies only the calling pthread to the new process's address space. If a mutex lock is held by another pthread, then the calling pthread in the new process may deadlock. For example, if a process has 2 pthreads and pthread 0 calls <b>fork()</b> when pthread 1 is holding a mutex lock inside a call to <b>printf()</b> , the only pthread in the new process will hang when it calls <b>printf()</b> .
<b>exec()</b>	Kills all pthreads other than the calling pthread, then loads a new program into the process's address space. The result is a new program with one pthread. The calling pthread can create additional pthreads if it wants.
<b>chdir()</b>	Changes the current working directory for all pthreads in the calling process.
<b>sleep()</b>	Puts only the calling pthread to sleep.
<b>wait()</b>	Blocks only the calling pthread; does not return until every pthread of the process being waited for exits. Note that a pthread can <b>wait()</b> for a process created by a different pthread.
<b>perror()</b>	Uses the per-pthread <i>errno</i> (see “ <i>errno</i> Confusion” on page 6-41).

Note that these calls may have different semantics on different platforms:

- *libc\_r.a* also includes some calls whose names end in `_r`. These calls do the same thing as the similarly-named calls in the standard C library, but have different parameters.
  - If the `_r` call is the only version of the call in *libc\_r.a*, you must use the `_r` call to be sure your code is thread-safe. This is the case for most of the `_r` calls.
  - If both an `_r` and a non-`_r` version of the call exist in *libc\_r.a*, both versions are thread-safe, but the `_r` version offers better performance. This is the case for `gmtime()`, `initstate()`, `rand()`, `random()`, `readdir()`, `setstate()`, and `srandom()`.

The `_r` calls are noted with an asterisk in Table 6-1.

- Finally, the calls `flockfile()`, `funlockfile()`, and `unlocked_...()` exist only in *libc\_r.a*:
 

<b>flockfile()</b>	Locks the specified standard I/O stream for exclusive use by the calling pthread.
<b>funlockfile()</b>	Unlocks the specified stream.
<b>unlocked_...()</b>	Perform operations on a stream while it is locked (these are called “unlocked” calls because they do not perform any locking or unlocking of their own).

The `unlocked_...()` calls are not thread-safe by themselves; they must be used together with `flockfile()/funlockfile()`.

These calls offer better I/O performance and more control over I/O from pthreads than the standard thread-safe I/O calls. For example, the thread-safe version of `putc()` locks out all other I/O calls, writes the specified character, then unlocks. If you write a series of characters to a file with `putc()`, this locking and unlocking results in considerable overhead; also, there is nothing to prevent characters written by two different pthreads from becoming intermingled.

You can instead use `flockfile()` to lock out all other operations on the file, a series of `unlocked_putc()` calls to write characters without locking and unlocking, and finally a `funlockfile()` to release the lock. In this case only one pair of lock/unlock operations is performed; your I/O performance will be better, and no other pthread's output can interfere. See the *OSF/1 Programmer's Reference* for more information on these calls.



## Using Pthreads Library Calls

This section tells you how to use the calls in *libpthreads.a* to create and control pthreads in your programs. See the *OSF/1 Programmer's Reference* for more detailed information on each call.

### Pthreads Library Data Types and Symbols

In order to use any calls from the pthreads library, your program must include the file `<pthread.h>`, which defines several types and symbols used by this library. The most important of these are:

***pthread\_t*** Within each process, each active pthread is identified by a unique *pthread ID*, which is a value of type *pthread\_t*. You use a pthread's ID to identify the pthread in all calls that control pthreads.

***pthread\_mutex\_t* and *pthread\_cond\_t*** Each active mutex is identified by a value of type *pthread\_mutex\_t*, and each active condition variable is identified by a value of type *pthread\_cond\_t*.

***pthread\_attr\_t*, *pthread\_mutexattr\_t*, and *pthread\_condattr\_t*** Objects of these types, called *attributes objects*, are used to specify the attributes (characteristics) of pthreads, mutexes, and condition variables. These types are extensible, and can support new features added by later revisions of the pthreads standard while maintaining compatibility with existing programs. Objects of these types are created with default values and can be changed by pthreads library calls.

***pthread\_attr\_default*, *pthread\_mutexattr\_default*, and *pthread\_condattr\_default*** These are external symbols whose values are the default attributes for an object of the appropriate type. If you want to create an object with the default attributes, you can use one of these symbols instead of creating a new attributes object with the default attributes.

***pthread\_key\_t*** This data type supports the per-pthread global data structure in the pthreads library. This enables different functions to access global data that only belongs to a single pthread.

## The Main Thread

Each program initially has a single thread—the flow of control that starts at the beginning of the function `main()`. This thread is referred to as the *main thread*.

Any other pthreads in the program are created by the main thread, either directly or indirectly. But threads do not have a parent-child relationship, as processes do, so the main thread does not have any special relationship with or control over other pthreads in the process.

However, the C library treats the function `main()` specially, in a way that can affect other threads in the process:

### NOTE

If the function `main()` returns (either by executing a `return` statement or by reaching the closing brace of the function), the C library generates an implicit call to `_exit()`, which kills all pthreads in the process and terminates the process.

This means that you must either have `main()` wait for all other pthreads before returning, or make sure that `main()` always terminates itself by calling `pthread_exit()` rather than calling `exit()` or returning.

## Managing Pthread Execution

---

Synopsis	Description
<pre>int pthread_create(     pthread_t *thread,     pthread_attr_t attr,     void *(*routine)(void *arg),     void *arg );</pre>	Creates a pthread.
<pre>pthread_t pthread_self(void);</pre>	Returns the ID of the calling pthread.
<pre>int pthread_equal(     pthread_t thread1,     pthread_t thread2 );</pre>	Compares two pthread identifiers.
<pre>void pthread_yield(void);</pre>	Allows the scheduler to run another pthread instead of the current one.
<pre>void pthread_exit(     void *status );</pre>	Terminates the calling pthread.
<pre>int pthread_join(     pthread_t thread,     void **status );</pre>	Waits for a pthread to terminate.
<pre>int pthread_detach(     pthread_t *thread );</pre>	Detaches a pthread.

---

To create a pthread, call **pthread\_create()**. This call has the following parameters:

<i>thread</i>	Pointer to a variable of type <i>pthread_t</i> that receives the pthread ID of the newly-created pthread.
<i>attr</i>	An object of type <i>pthread_attr_t</i> that describes the desired attributes of the new pthread. This can be the default pthread attribute object <i>pthread_attr_default</i> , or a user-created pthread attribute object (see “Managing Pthread Attributes” on page 6-15).
<i>routine</i>	Pointer to the initial function to be executed by the new pthread. This function is assumed to return <i>void *</i> and to have one argument of type <i>void *</i> .
<i>arg</i>	Value of type <i>void *</i> to be passed to the initial function as its argument.

A pthread can determine its own pthread ID by calling **pthread\_self()**, and a pthread ID can be compared against another pthread ID by calling **pthread\_equal()**. Note that *pthread\_t* is an “opaque” type, and you should not use standard C operators on it.

The **pthread\_yield()** call will decrease the priority of the calling pthread and give up the node's processor to other pthreads that have higher priorities than the calling pthread. The kernel decides which thread to run next, based on its time sharing and aging policies. Eventually, the calling pthread will be scheduled to run again when other pthreads become lower priority pthreads. A pthread should call **pthread\_yield()** to give up the processor when it is making no progress or has no work to do.

A pthread terminates when it calls **pthread\_exit()** or returns from its initial function. However, the termination of a pthread does not release all the resources associated with the pthread. To release a terminated pthread's resources, a different pthread must call **pthread\_join()** or **pthread\_detach()**:

- **pthread\_join()** blocks until the specified pthread terminates, then releases the specified pthread's resources and returns the exit status of the specified pthread to its caller. The exit status is the value specified in the pthread's **pthread\_exit()** call, or the return value of its initial function if it did not call **pthread\_exit()**.
- **pthread\_detach()** tells the pthreads library to release the specified pthread's resources, then returns immediately to its caller. Later, when the specified pthread terminates, the library releases the pthread's resources and discards the pthread's exit status.

Any pthread that creates other pthreads should call **pthread\_join()** or **pthread\_detach()** for each pthread it created before it terminates itself.

## Managing Pthread Attributes

---

Synopsis	Description
<code>int pthread_attr_create( pthread_attr_t *attr );</code>	Creates a pthread attributes object.
<code>int pthread_attr_setstacksize( pthread_attr_t *attr, long stacksize );</code>	Sets the value of the stack size attribute of a pthread attributes object.
<code>int pthread_attr_delete( pthread_attr_t *attr );</code>	Deletes a pthread attributes object.
<code>int pthread_attr_getstacksize( pthread_attr_t attr );</code>	Returns the value of the stack size attribute of a pthread attributes object.

---

The only pthread attribute that is currently modifiable is stack size. (A pthread's priority and scheduling policy are managed by the kernel and cannot be inspected or changed.) To set a pthread's stack size, use the following procedure:

1. Call `pthread_attr_create()` to create a pthread attributes object (an object of type `pthread_attr_t`).
2. Call `pthread_attr_setstacksize()` to set the stack size in that object.
3. Use the modified pthread attributes object in the call to `pthread_create()` that creates the pthread.
4. Call `pthread_attr_delete()` to remove the pthread attributes object.

Once a pthread has been created, the size of its stack is fixed and can't be changed.

To use the default stack size, you can simply use the default pthread attribute object `pthread_attr_default` instead of creating your own pthread attributes object.

You can use `pthread_attr_getstacksize()` to find out the current stack size in a pthread attributes object.

## NOTE

Whenever possible, use the same stack size for all pthreads.

Each pthread is built on a lower-level construct called a *kernel thread*. When you create a pthread, the pthread library tries to re-use a kernel thread from a pool of existing kernel threads. This means that creating a new pthread is more expensive if there are no existing kernel threads inside the pthreads library that can be reused. A major cause for the kernel being unable to recycle kernel threads is using a different stack size for new pthreads; this should be avoided.

## Managing Mutexes

Synopsis	Description
<pre>int pthread_mutex_init(     pthread_mutex_t *mutex,     pthread_mutexattr_t attr );</pre>	Creates a mutex.
<pre>int pthread_mutex_lock(     pthread_mutex_t *mutex );</pre>	Locks a mutex.
<pre>int pthread_mutex_trylock(     pthread_mutex_t *mutex );</pre>	Tries once to lock a mutex.
<pre>int pthread_mutex_unlock(     pthread_mutex_t *mutex );</pre>	Unlocks a mutex.
<pre>int pthread_mutex_destroy(     pthread_mutex_t *mutex );</pre>	Deletes a mutex.

A pthread *mutex* is a binary semaphore with two states: *locked* and *unlocked*. When a mutex is created, its initial state is unlocked. Only one pthread at a time can lock a mutex. When a pthread successfully locks a mutex, it becomes the mutex's *owner*. Any other pthread that attempts to lock the mutex will block until the owner unlocks the mutex. Mutexes cannot be used recursively: if the owner attempts to lock the mutex again, the attempt fails.

You should use mutex locks to serialize pthread access to a block of code that accesses a nonshareable resource, such as a file or a non-thread-safe library. A pthread that is waiting on a mutex lock will not use any of the node's processor time.

To create and initialize a mutex, call `pthread_mutex_init()`. This call creates a new mutex with the attributes specified by *attr* (typically the default mutex attributes object `pthread_mutexattr_default`) and stores the new mutex's ID into the variable pointed to by *mutex*. A newly-created mutex is unlocked.

To lock a mutex, call `pthread_mutex_lock()`. The call to `pthread_mutex_lock()` will block the calling pthread until the mutex lock is available. A pthread waiting on a mutex lock will be scheduled out and another pthread will be scheduled to run. When the calling pthread is again scheduled to run because no higher-priority pthread can run, it checks the availability of the mutex lock again and is scheduled out again if the mutex lock is still unavailable.

Note that there is no guarantee that a pthread waiting on a `pthread_mutex_lock()` will eventually get the lock. If you do not want to block until the lock is available, call `pthread_mutex_trylock()`. This call tries once to lock the specified mutex. If the attempt succeeds, the call returns 1 immediately; but if the mutex is already locked, the call returns 0 immediately.

When a pthread is finished using the resource controlled by the mutex, it should release the lock by calling `pthread_mutex_unlock()`. This allows any other pthread that has been waiting to lock the mutex to proceed.

When all pthreads have finished using the mutex, you should remove it and release all resources associated with it by calling `pthread_mutex_destroy()`. You cannot destroy a mutex that is currently locked. Attempting to lock or unlock a mutex that has been successfully destroyed will result in undefined behavior.

## Managing Mutex Attributes

---

Synopsis	Description
<code>int pthread_mutexattr_create( pthread_mutexattr_t *attr );</code>	Creates a mutex attributes object.
<code>int pthread_mutexattr_delete( pthread_mutexattr_t *attr );</code>	Deletes a mutex attributes object.

---

No mutex attributes are currently defined. You can either use the default mutex attributes object, `pthread_mutexattr_default`, or create a mutex attribute object for use in `pthread_mutex_init()` by calling `pthread_mutexattr_create()`. A user-created mutex attributes object should be released by calling `pthread_mutexattr_delete()`.

## An Example Pthreads Program

The following program demonstrates some principles of using pthreads and mutexes. It creates a user-specified number of pthreads, each of which prints its node number, ptype, and pthread ID and the message "Done."

```
#include <pthread.h>
#include <stdlib.h>
#include <nx.h>

#define MAXTHREAD      6                /* thread maximum limit */

/* pthread resources */
pthread_t              thread[MAXTHREAD]; /* per-thread pthread ID */
pthread_mutex_t       mutex;             /* mutex to protect global
                                         variable "thread_alive" */

/* global variables that only the main thread writes to */
int    max_thread = MAXTHREAD;          /* maximum thread number */
int    my_node;                          /* my node number */
int    my_ptype;                          /* my ptype */

/* shared global variable that is modified by all threads */
int    thread_alive;                     /* count of living threads */

/* forward declarations */
void thread_fun(int thread_id);          /* initial function for
                                         new threads */

main(int argc, char *argv[])
{
    int    index;                          /* loop index */
    int    my_thread = 0;                  /* main thread is indexed 0 */

    my_node = mynode();
    my_ptype = myptype();

    if(argc != 2) {
        if(my_node == 0) {
            printf("Usage: %s <nthreads>\n", argv[0]);
        }
        exit(1);
    }
}
```



```
max_thread = atoi(argv[1]);

if(max_thread > MAXTHREAD) {
    if(my_node == 0) {
        printf("Error: %d threads requested, must be %d or less\n",
            max_thread, MAXTHREAD);
    }
    exit(1);
}

/* The main thread is the last thread alive, so don't count itself. */
thread_alive = max_thread - 1;

/* create and initialize a mutex to control access to "thread_alive" */
if(pthread_mutex_init(&mutex, pthread_mutexattr_default) == -1) {
    perror("pthread_mutex_init Error");
}

/*
 * Spawn threads and remember each thread's pthread ID.
 * The main thread is indexed as thread 0.
 */
thread[my_thread] = pthread_self();
for(index = 1; index < max_thread; index++) {
    if(pthread_create(&thread[index], pthread_attr_default,
        (void *)thread_fun, (void *)index) == -1) {
        perror("pthread_create Error");
        exit(2);
    }
}

/* loop until all other threads are finished. */
while(thread_alive != 0) {
    pthread_yield();
}

/*
 * Ignore other threads' exit status (can also be done right after
 * pthread_create()).
 */
for(index = 1; index < max_thread; index++) {
    pthread_detach(&thread[index]);
}

printf("( %3d, %3d, %3d) Done\n", my_node, my_ptype, my_thread);
}
```

```

/*****
 * thread_fun() -- This is the initial function for new threads
 *****/
void thread_fun(int my_thread)
{
    printf("( %3d, %3d, %3d) Done\n", my_node, my_ptype, my_thread);

    /*
     * use mutex to protect global variable "thread_alive"
     */
    if(pthread_mutex_lock(&mutex) == -1) {
        perror("pthread_mutex_lock Error");
    }
    thread_alive--;
    if(pthread_mutex_unlock(&mutex) == -1) {
        perror("pthread_mutex_unlock Error");
    }

    /* terminate (status is ignored) */
    pthread_exit(NULL);
}

```

Assuming this program is called *pthread.c*, use the following command to compile it:

```
% cc -D_REENTRANT -o pthreads pthreads.c -lpthreads -lc_r -nx
```

To run this program with three pthreads per process on two nodes of your default partition, use the following command:

```
% pthreads 3 -sz 2
```

The results of this application run:

```

(  0,  0,  2) Done
(  1,  0,  1) Done
(  0,  0,  1) Done
(  0,  0,  0) Done
(  1,  0,  2) Done
(  1,  0,  0) Done

```

Note that the results may appear in a different order on each run.

## Using Condition Variables to Synchronize Pthreads

---

Synopsis	Description
<pre>int pthread_cond_init(     pthread_cond_t *cond,     pthread_condattr_t attr );</pre>	Creates a condition variable.
<pre>int pthread_cond_wait(     pthread_cond_t *cond,     pthread_mutex_t *mutex );</pre>	Waits on a condition variable.
<pre>int pthread_cond_timedwait(     pthread_cond_t *cond,     pthread_mutex_t *mutex,     struct timespec *abstime );</pre>	Waits on a condition variable for a specified period of time.
<pre>int pthread_cond_signal(     pthread_cond_t *cond );</pre>	Wakes up a pthread that is waiting on a condition variable.
<pre>int pthread_cond_broadcast(     pthread_cond_t *cond );</pre>	Wakes up all pthreads that are waiting on a condition variable.
<pre>int pthread_cond_destroy(     pthread_cond_t *cond );</pre>	Destroys a condition variable.

---

The pthreads package provides *condition variables* to help you synchronize the execution of pthreads. You can also synchronize pthreads by looping on the value of a variable, as shown in the previous example, but using condition variables is more efficient. Condition variables use the following objects:

- A *mutex* (an object of type *pthread\_mutex\_t*, as discussed under “Managing Mutexes” on page 6-16) that is used to protect the pthread condition variable and the global predicate variable.
- A *condition variable* (an object of type *pthread\_cond\_t*) that links all pthreads waiting on a particular condition.
- A *global predicate variable* that indicates the current state of the condition. It could be a global integer variable.

To create and initialize a condition variable, call **pthread\_cond\_init()**. This call creates a new condition variable with the attributes specified by *attr* (typically the default condition attributes object *pthread\_condattr\_default*) and stores the new condition variable's ID into the variable pointed to by *cond*. The list of pthreads that are waiting on the new condition variable is initially empty.

To wait for a condition, call **pthread\_cond\_wait()**. This call unlocks the specified mutex and blocks until the specified condition is signaled by another pthread. When the condition is signaled, the call re-locks the mutex and returns to the caller. **pthread\_cond\_timedwait()** is similar, but if the specified amount of time passes before the condition is signaled, the call re-locks the mutex and returns an error condition. You must successfully lock the specified mutex before calling **pthread\_cond\_wait()** or **pthread\_cond\_timedwait()**, and you should unlock the mutex after the call to **pthread\_cond\_wait()** or **pthread\_cond\_timedwait()** returns.

To signal a condition, call **pthread\_cond\_signal()** or **pthread\_cond\_broadcast()**. **pthread\_cond\_signal()** signals the specified condition to one of the pthreads that is waiting for it (if more than one pthread is waiting for the condition to be signaled, the kernel selects one of them arbitrarily). **pthread\_cond\_broadcast()** signals the specified condition to all of the pthreads that are waiting for it. If no other pthread is waiting on the condition, these calls have no effect. No mutex lock is required for these calls (but a mutex can be used to prevent certain race conditions; see the example on page 6-25).

If a pthread calls **pthread\_cond\_wait()** *after* the specified condition has been signaled, that pthread could wait forever. To prevent this problem, use a *global predicate variable*. This can be any variable that is visible to all pthreads. You use it as follows:

- Before calling **pthread\_cond\_signal()** or **pthread\_cond\_broadcast()**, a pthread should set the value of the condition's global predicate value to indicate that the condition has occurred. The global predicate value should be protected by a mutex if there is any possibility that more than one pthread could try to set it at once.
- Before calling **pthread\_cond\_wait()** or **pthread\_cond\_timedwait()**, a pthread should check the current value of the condition's global predicate variable. If the condition has already occurred, the pthread should proceed without calling **pthread\_cond\_wait()** or **pthread\_cond\_timedwait()**.
- After a successful call to **pthread\_cond\_wait()** or **pthread\_cond\_timedwait()**, a pthread should check the global predicate value to be sure it has the expected value. If the global predicate variable does not indicate that the condition has occurred, the pthread should call **pthread\_cond\_wait()** or **pthread\_cond\_timedwait()** again.

The first example shown under "Examples of Condition Variables" on page 6-24 gives an example of this technique.

When all pthreads have finished using a condition variable, you should remove it and release all resources associated with it by calling **pthread\_cond\_destroy()**. You cannot destroy a condition variable that is currently being waited on.

## Managing Condition Attributes

---

### Synopsis

```
int pthread_condattr_create(  
    pthread_condattr_t *attr );
```

```
int pthread_condattr_delete(  
    pthread_condattr_t *attr );
```

### Description

Creates a condition variable attributes object.

Deletes a condition variable attributes object.

---

No condition attributes are currently defined. You can either use the default condition attributes object, *pthread\_condattr\_default*, or create a condition attribute object for use in `pthread_cond_init()` by calling `pthread_condattr_create()`. A user-created condition attributes object should be released by calling `pthread_condattr_delete()`.

## Examples of Condition Variables

The following example uses a mutex to protect the global predicate variable `cond_true`, which is used to prevent the signaling pthread from calling `pthread_cond_signal()` until the waiting pthread has called `pthread_cond_wait()`. Note that the call to `pthread_cond_signal()` is within the mutex lock; this is not necessary, but can prevent certain race conditions.

The pthread waiting for the condition executes the following code:

```
if(pthread_mutex_lock(&mutex) == -1) {
    perror("pthread_mutex_lock Error");
}

/*
 * If the expected condition already exists, don't call
 * pthread_cond_wait(), since the condition signal will not be
 * sent if no threads are waiting for this condition.
 *
 * Recheck the state of cond_true after calling
 * pthread_cond_wait() to ensure that the received condition
 * signal is for this expected state change.
 */
while(!cond_true) {
    /*
     * mutex will be unlocked in pthread_cond_wait() when
     * calling thread is ready to wait.
     */
    if(pthread_cond_wait(&cond, &mutex) == -1) {
        perror("pthread_cond_wait Error");
        break;
    }
    /*
     * mutex will be locked again when the calling thread
     * is awakened.
     */
}

if(pthread_mutex_unlock(&mutex) == -1) {
    perror("pthread_mutex_unlock Error");
}
```

The pthread signaling the condition executes the following code:

```

if(pthread_mutex_lock(&mutex) == -1) {
    perror("pthread_mutex_lock Error");
}

/* This global variable needs a mutex's protection. */
++cond_true;

/*
 * The pthread_cond_signal() call does not use a mutex internally.
 * The mutex protection will guarantee that every thread can
 * catch the expected condition signal once it calls
 * pthread_cond_wait(). This will prevent the endless block of
 * the calling thread.
 */
if(pthread_cond_signal(&cond) == -1) {
    perror("pthread_cond_signal Error");
}

if(pthread_mutex_unlock(&mutex) == -1) {
    perror("pthread_mutex_unlock Error");
}

```

Here's another example, which uses `pthread_cond_broadcast()` to wake up all waiting pthreads.

```

/*
 * Simulate a thread-level gsync() to synchronize all active
 * threads at a barrier. A counter in this example can only be
 * used once.
 */

long          cond_gsync; /* counter of threads arrived */
long          max_thread; /* number of active threads */
pthread_mutex_t mutex;
pthread_cond_t cond;

void thread_gsync(long *cond_gsync)
{
    if(pthread_mutex_lock(&mutex) == -1) {
        perror("pthread_mutex_lock Error");
    }

    /* Increase the count of threads that have called
     * thread_gsync() */
    *cond_gsync++;
}

```

```
if(*cond_gsync == max_thread) {
    /*
     * If I'm the last thread to call thread_gsync(),
     * wake up all threads waiting on this condition.
     */
    if(pthread_cond_broadcast(&cond) == -1) {
        perror("pthread_cond_broadcast Error");
    }
} else {
    while(*cond_gsync != max_thread) {
        /*
         * Other threads haven't called thread_gsync() yet,
         * so wait for them in pthread_cond_wait().
         */
        if(pthread_cond_wait(&cond, &mutex) == -1) {
            perror("pthread_cond_wait Error");
            break;
        }
    }
}

if(pthread_mutex_unlock(&mutex) == -1) {
    perror("pthread_mutex_unlock Error");
}
}

main()
{
    /*
     * Initialize counter and the condition that counter
     * will meet.
     */
    max_thread = atoi(argv[1]); /* threads to create */
    cond_gsync = 0;           /* init counter */

    if(pthread_cond_init(&cond,
                        pthread_condattr_default) == -1) {
        perror("pthread_cond_init Error");
    }

    if(pthread_mutex_init(&mutex,
                        pthread_mutexattr_default) == -1) {
        perror("pthread_mutex_init Error");
    }
    .
    .
    .
}
```



```

    /* Create more pthreads */
    .
    .
    .
}

/* Every thread calls thread_gsync() once */
thread_gsync(&cond_gsync);
/* Every thread passes this barrier at the same time */

```

Here's an example using `pthread_cond_timedwait()`:

```

#include <sys/timers.h>

long    interval = 10;          /* 10 seconds interval */

struct timespec abs_time;

/* Get the current time */
getclock(TIMEOFDAY, &abs_time);

/*
 * Can use another member of structure timespec to specify the
 * waiting interval in nanoseconds. But the resolution cannot be
 * smaller than the interval between updates of the system clock.
 *
 * The wait time should not be so small that the
 * absolute time specified is smaller than the
 * time spent inside the pthread_cond_timedwait() call.
 */
abs_time.tv_sec = abs_time.tv_sec + interval;

if(pthread_mutex_lock(&mutex) == -1) {
    perror("pthread_mutex_lock Error");
}

while(!condition) {
    if(pthread_cond_timedwait(&cond, &mutex, &abs_time) == -1) {
        /* EAGAIN is the timeout error code. */
        if(errno != EAGAIN) {
            perror("pthread_cond_timedwait Error");
        }
        break;
    }
}

if(pthread_mutex_unlock(&mutex) == -1) {
    perror("pthread_mutex_unlock Error");
}

```

## Canceling Pthreads

Synopsis	Description
<code>int pthread_cancel(pthread_t thread);</code>	Requests cancellation of a pthread.
<code>int pthread_setcancel(int state);</code>	Enables or disables the general cancelability of the calling pthread.
<code>int pthread_setsynccancel(int state);</code>	Enables or disables the asynchronous cancelability of the calling pthread.
<code>void pthread_testcancel(void);</code>	Creates a cancellation point in the calling pthread.

The pthreads package includes a *pthread cancellation mechanism* that allows a pthread to terminate the execution of other pthreads. The call `pthread_cancel()` requests cancellation of the specified pthread; however, the specified pthread may terminate later or not at all, depending on its *cancelability states*.

### Cancelability States

Each pthread has two *cancelability states* that determine how it reacts to cancellation requests. Each of the two states can be set to the value `CANCEL_ON` (enabled) or `CANCEL_OFF` (disabled).

- If *general cancelability* determines whether or not the pthread can be canceled:
  - If general cancelability is enabled, cancellation requests are accepted. Cancellation may or may not occur immediately, depending on the *asynchronous cancelability* state.
  - If general cancelability is disabled, cancellation requests are queued until general cancelability is enabled again.

General cancelability is enabled by default; a pthread can change its general cancelability state by calling `pthread_setcancel()`.

- When general cancelability is enabled, a second cancelability state called *asynchronous cancelability* determines how quickly the cancellation occurs:
  - If asynchronous cancelability is enabled, when a cancellation request is received the pthread begins termination immediately.
  - If asynchronous cancelability is disabled, when a cancellation request is received the pthread does not begin termination until it reaches a *cancellation point*. The default cancellation points are calls to `pthread_cond_wait()`, `pthread_cond_timedwait()`, `pthread_join()`, and `pthread_setcancel(CANCEL_ON)`. A pthread can also create an explicit cancellation point by calling `pthread_testcancel()`, which otherwise does nothing.

Asynchronous cancelability is disabled by default; a pthread can change its asynchronous cancelability state by calling `pthread_setsynccancel()`.

## NOTE

Asynchronous cancelability should not be enabled in the current release.

Asynchronous cancellation of certain pthreads, particularly pthreads performing file I/O, can cause the entire application to hang.

## NOTE

You must be careful not to cancel a pthread that is holding a mutex lock.

Canceling a pthread that is holding a mutex lock leaves the mutex locked with no way to unlock it, possibly resulting in deadlock. For example, a pthread calling `printf()` will get a mutex lock inside the reentrant C library. A cancellation of this pthread during the call to `printf()` will cause all other pthreads calling `printf()` to deadlock.

Functions such as `printf()`, which can cause deadlock if they are canceled, are called *not safe to cancel*.

## NOTE

Most library functions are not safe to cancel.

In particular, all of the calls in *libn.a* are not safe to cancel. The list of functions that *is* safe to cancel can be found in the `pthread_setsynccancel()` manpage in the *OSF/1 Programmer's Reference*.

## Cancellation Examples

Here's an example of changing a pthread's cancelability states:

```
/* flip the general cancelability of the calling thread */
if(pthread_setcancel(CANCEL_ON) == -1) {
    perror("pthread_setcancel Error");
}

if(pthread_setcancel(CANCEL_OFF) == -1) {
    perror("pthread_setcancel Error");
}

/* flip the asynchronous cancelability of the calling thread */
if(pthread_setsynccancel(CANCEL_ON) == -1) {
    perror("pthread_setsynccancel Error");
}

if(pthread_setsynccancel(CANCEL_OFF) == -1) {
    perror("pthread_setsynccancel Error");
}
```

Here's an example of delivering and accepting cancellations:

```
pthread_t  thread_id; /* value from pthread_create() call */
.
.
.
/*
 * Cancel another thread whose pthread ID is "thread_id".
 */
if(pthread_cancel(thread_id) == -1) {
    perror("pthread_cancel Error\n");
}
.
.
.
/*
 * If a cancellation request is already posted, this call will
 * not return.
 */
pthread_testcancel();
/* Execution continues if no posted cancellation request */
```

## Pthreads Cleanup Routines

---

**Synopsis**

```
void pthread_cleanup_pop(  
    int execute );
```

```
void pthread_cleanup_push(  
    void (*routine)(void *arg),  
    void *arg );
```

**Description**

Removes a routine from the top of the cleanup stack of the calling pthread and optionally executes it.

Pushes a routine onto the cleanup stack of the calling pthread.

---

Pthreads may have resources that must be released before the pthread terminates. Each pthread can create a list of cleanup routines, called the *cleanup stack*, to release those resources. The routines on the cleanup stack are called, in order from top to bottom, when the pthread terminates for any of the following reasons:

- Calling `pthread_exit()`.
- Returning from its initial function.
- Being cancelled by another pthread.

To place a function on the cleanup stack, call `pthread_cleanup_push()`; to remove the top function from the cleanup stack, call `pthread_cleanup_pop()`. You can optionally execute the function as it is popped. Every call to `pthread_cleanup_push()` must be matched with a `pthread_cleanup_pop()` call in the same *lexical scope* (that is, within the same set of “{ ... }” braces).

If general cancelability is enabled, whenever a pthread allocates a resource it should push a function that deallocates that resource onto the cleanup stack; when the pthread is finished with the resource it should deallocate it by popping the function off the cleanup stack and executing it. This ensures that all resources are accounted for if the pthread is cancelled.

## Managing Pthread Keys

---

### Synopsis

```
int pthread_keycreate(
    pthread_key_t *key,
    void (*destructor)(void *value) );
```

```
int pthread_setspecific(
    pthread_key_t key,
    void *value );
```

```
int pthread_getspecific(
    pthread_key_t key,
    void **value );
```

---

### Description

Creates a key to be used with pthread-specific data.

Binds a pthread-specific value to a key.

Returns the value bound to a key.

The pthreads package provides *pthread-specific data objects* to associate information with individual pthreads. Each pthread-specific data object is controlled by a *key* (an object of type *pthread\_key\_t*). A pthread creates a new key by calling `pthread_keycreate()`, associates the key with a pthread-specific data object by calling `pthread_setspecific()`, and then retrieves the data associated with the key by calling `pthread_getspecific()`. See the *OSF/1 Programmer's Reference* for more information on these calls.

## Executing a Routine Once

---

### Synopsis

```
int pthread_once(
    pthread_once_t *once_block,
    void(*routine)());
```

### Description

Calls an initialization routine.

---

The `pthread_once()` call executes the specified routine the first time it is called (from any pthread), and does nothing every subsequent time. The parameter `once_block` must be declared as `static`. For example:

```
static pthread_once_t  init_once;

void lib_util_init() {
    /* perform some initialization that can only be done once */
}
.
.
.
/*
 * Every pthread calls pthread_once(), but only the first one
 * executes lib_util_init().
 */
if(pthread_once(&init_once, lib_util_init) == -1) {
    perror("pthread_once Error");
}
```

## Managing Signals

---

### Synopsis

```
int sigwait(
    sigset_t *set);
```

### Description

Suspends the calling pthread until one of a specified set of signals is received.

---

The `sigwait()` call is used to turn asynchronous signals into synchronous notifications. Before calling `sigwait()`, you must create a *signal set*, using the standard signal calls `sigemptyset()`, `sigfillset()`, `sigaddset()`, and `sigdelset()`, and then block the signals in that set from being delivered. When you call `sigwait()` with that signal set, the calling pthread is suspended until one or more of the signals in the set is received by the process containing the pthread. If one of the specified signals was received (and blocked) before the call to `sigwait()`, the call returns immediately. `sigwait()` returns the signal number of the signal that was received.



The `sigwait()` call only works for *asynchronous* signals (those that are generated externally from the pthread, such as those generated by `kill()` in other processes or by the user pressing `<Ctrl-C>`). Contrast this with `sigaction()`, which only works for *synchronous* signals (those that are generated as the result of the pthread's faults, such as `SIGBUS`). If both `sigaction()` and `sigwait()` are used on the same signal, the results are unspecified.

## NOTE

In a parallel application, sending an asynchronous signal to an application's controlling process affects the controlling process (as specified by the controlling process's signal mask), and also causes the signal to be broadcast to the compute processes. In an application linked with `-nx`, the controlling process's signal mask is always the default.

See "Signals and Pthreads Library Calls" on page 6-39 for more information on signals in applications with multiple pthreads.

Here's an example that uses `sigwait()` to deal with the asynchronous signal `SIGQUIT`. This example uses a parent process to generate the asynchronous signal by calling `kill()`.

```

long    sig;
long    ret;

sig = SIGQUIT;

pid = fork();

if(pid == -1) {
    perror("fork()");
    exit(1);
} else if(pid != 0) {
    /* parent process */
    sleep(2);
    /*
     * Deliver the signal SIGQUIT to child process.
     */
    if(kill(pid, sig) == -1) {
        perror("kill ");
        exit(1);
    }
    exit(0);
}

```

```
/* child process */
if(sigemptyset(&set) != 0) {
    perror("sigemptyset");
}

/* Add the signal SIGQUIT to the signal mask */
if(sigaddset(&set, sig) != 0) {
    perror("sigaddset");
}

/* Block the signal SIGQUIT from delivery */
if(sigprocmask(SIG_BLOCK, &set, NULL) != 0) {
    perror("sigprocmask()");
}

/*
 * During the next 10 seconds, the posted signal from the
 * parent process becomes a pending signal.
 */
sleep(10);

/*
 * sigwait() blocks the calling thread until the specified signal
 * arrives, then unblocks and returns the value SIGQUIT.
 */
if((ret = sigwait(&set)) == -1) {
    perror("sigwait()");
} else {
    printf("Received signal %d, expected %d\n", ret, sig);
}

/*
 * The thread can decide what to do with this signal.
 */

/*
 * There is no destructive default action of core dump on the
 * posted signal SIGQUIT.
 */
```

## Interfacing with Non-Thread-Safe Code

Whenever you call a non-thread-safe library from a process with multiple pthreads, you must make sure that no two pthreads call the same library at the same time. There are two ways to do this:

- Make sure that only one pthread ever calls the library.
- Use mutexes to protect all calls to the library.

Here's an example of the second technique:

```
pthread_mutex_lock(&mutex);
non_thread_safe_call();
pthread_mutex_unlock(&mutex);
```

Note that the *same* mutex must be used by all pthreads for any calls from the same library. If all calls to the non-thread-safe library are surrounded by a lock and unlock of the same mutex, as shown here, any pthread that calls the library while another pthread is currently calling it will block until the other pthread returns and unlocks the mutex. See “Managing Mutexes” on page 6-16 for more information on mutexes.

## Message Passing and Pthreads Library Calls

Paragon OSF/1 message-passing is done on a process-by-process basis. All message-sending calls specify the recipient by node and process type; there is no way to specify a particular pthread within that process (all the threads in a process have the same process type). Similarly, when a message arrives at a process, there is nothing to prevent confusion among pthreads; for example, a pthread could probe for a message, find a pending message of the specified type, and then attempt to receive it—only to find that another pthread has already received it. For this reason, you should make sure that only one pthread in each process uses message-passing calls.

You should also keep the following special considerations in mind when using message-passing and pthreads in the same application:

- Blocking calls, such as `csend()`, only block the calling pthread, not the entire process. While the calling pthread is blocked, other pthreads can continue to run. The pthread that is blocked releases processor resources until the `csend()` returns.
- When the message-passing pthread uses one of the global calls (those described under “Global Operations” on page 3-27), the call blocks until the message-passing pthread on every other node makes the same call. If one of those message-passing pthreads is blocked (for example, by a mutex lock), the operation will hang all the message-passing pthreads in the application.

- An **hsend()/hrecv()** handler cannot use calls from *libpthreads.a* or *libc\_r.a* (note that *libc\_r.a* includes almost the entire C library). Using these calls within a handler can give unexpected results or cause the handler to hang. If you need to use any of these calls, you can have the handler use **csend()** to send a message to the main message-passing pthread to carry out the desired operation.
- Because an **hsend()/hrecv()** handler cannot use calls from *libc\_r.a*, if a call to *libn.a* within a handler fails, it could also hang the handler. This occurs because the failed library call will call **printf()** to print out the error message. To avoid this problem, use only underscore calls within **hsend()/hrecv()** handlers.
- An **hsend()/hrecv()** handler also should not use the **info...()** calls. Because the handler executes concurrently with the main message-passing pthread, the **info...()** calls may return values representing messages received by the main message-passing pthread. The main message-passing pthread can use **masktrap()** to protect critical regions from the handler.
- If an **hsend()/hrecv()** handler performs any message passing, you must put **masktrap()** calls around any message-passing calls in the main message-passing pthread that could be called while the handler is active. Otherwise, any **info...()** calls in the handler could reflect the value of a message received by the main message-passing pthread.

In addition, any **info...()** call in the main program must be within the same set of **masktrap()** calls as the message-receiving call to which it applies. Otherwise, the **info...()** call in the main message-passing pthread could reflect the value of a message received by the handler.

## File I/O and Pthreads Library Calls

In general, opened files are per-process resources. A pthread can open a file, a second pthread can use the open file descriptor to write or read, and a third pthread can use the same file descriptor in an **lseek()** call. The movement of file pointers is visible to all pthreads, so if multiple pthreads are accessing the same file they must coordinate their actions with mutexes or condition variables. However, blocking calls such as **read()** and **cwrite()** only block the calling pthread, not the entire process.

If two pthreads doing file I/O read and write concurrently, they can each read and write their own data independently. If you are performing I/O to a file in a synchronized PFS I/O mode (see “Using I/O Modes” on page 5-13), the synchronization information is stored with the file descriptor; each file is synchronized independently.

See “Recommended Safe Operating Environment” on page 6-4 for limitations on using I/O calls from multiple pthreads.

## nx\_nfork() and nx\_initve() and Pthreads Library Calls

In a controlling process with multiple pthreads:

- nx\_nfork()** Copies only the calling pthread to the new process on each node. Can only be called from one pthread.
- nx\_initve()** If the user's shell is the Bourne shell (**sh**), **nx\_initve()** performs a **fork()** internally. As described under "Using Reentrant C Library Calls" on page 6-6, **fork()** copies only the calling pthread to the new process. This means that if there are multiple pthreads in the calling process before the call to **nx\_initve()**, after the **nx\_initve()** all pthreads except the calling pthread will appear to cease to exist. (If the user's shell is **ksh** or **csk**, this problem does not exist.)
- nx\_initve()** can be called at most once in a process. This means that at most one pthread in a process can call it.

## Signals and Pthreads Library Calls

The following special considerations apply to signals in programs with multiple pthreads.

### Signal Types

There are two types of signals: synchronous and asynchronous:

- *Synchronous* signals are caused by a pthread's own actions, such as when a pthread divides by zero (causing a **SIGFPE** signal) or attempts to access memory outside its address space (causing a **SIGSEGV** signal).
- *Asynchronous* signals are caused by something external to the pthread, such as another process calling **kill()** or the user pressing <Ctrl-\\> on the keyboard (causing a **SIGQUIT** signal).

If a pthread causes a synchronous signal, the handler routine executes in the context of that pthread only. If an asynchronous signal is delivered to a process, the handler routine executes in the context of the main thread.

## Signals are a Per-Process Resource

Signals are generally managed as per-process objects in pthreads programs. Signal masks, signal handlers, and signal sending and receiving are all oriented toward the process, not toward a particular pthread. This means that signals affect the entire process. Of particular interest:

- A **SIGSTOP** signal stops all pthreads of the receiving process.
- A **SIGCONT** signal continues all pthreads of the receiving process.
- If one pthread of a program with multiple pthreads causes a **SIGSEGV** or **SIGBUS**, the entire process (not just the faulting pthread) receives the signal. If this signal has not been handled, all pthreads are killed and the program core dumps.

Which pthread is interrupted to execute a registered handler for a signal may be specific to the type of signal, but in general the entire process receives the signal.

It is important to be aware that a pthread program's signal mask has a per-process visibility. In other words, all pthreads share the same mask. If one pthread changes its mask (for example, by calling **sigprocmask()**) the change affects all pthreads. The thread-safe **sigwait()** requires manipulation of the signal mask, as does **sigaction()** and other common signal-management routines.

Along with the signal mask, signal handlers are also process-wide objects. A signal handler can be registered for the process (for example, by calling **sigaction()** or **sigwait()**) by any pthread. Because the handlers are process-wide objects, a second pthread registering a handler for a given signal will override the handler registered by the first pthread.

In general, blocking calls only block the calling pthread. This is the case with **sigsuspend()** as it is with **wait()**, **sleep()** etc. Note also, that if multiple pthreads are blocking on **sigsuspend()** for a given signal, all pthreads will continue when that signal arrives. This differs from **sigwait()** which unblocks only one of the pthreads.

**sigwait()** creates a hidden pthread which manipulates the process's signal mask and registers a signal handler for each signal **sigwait()** has been asked to wait for. Because of this, use of other signal management calls (especially **sigaction()**) on the signals being waited for, would be hazardous. Care must be taken when changing a signal mask so the state of a **sigwait()**ed signal's bit is not changed.

## Dealing with Signals

A way to deal with signals in a pthread application is the following:

- Use `sigaction()` to catch the synchronous signals. `sigaction()` only works with synchronous signals.
- Use `sigprocmask()` to block the asynchronous signals, then `sigwait()` to receive the signal as a notification. `sigwait()` only works with asynchronous signals.

Do not use `sigwait()` and `sigaction()` on the same signal.

## Handling Errors

The handling of error situations in a program with multiple pthreads should be robust and graceful. It should protect the pthreads that did not cause the error from being interrupted or terminated. It also should give information on which pthread caused the error and coordinate a proper shutdown of all pthreads if the error is fatal. If the error cannot be recovered from by the pthread causing the error, and other pthreads are depending on this pthread to progress, it might be best to terminate those pthreads right away and shut down the rest of the pthreads later.

## errno Confusion

In a single-threaded program, the global variable `errno` is set to an error value when a system or library call fails. An immediate call to `perror()` or `nx_perror()` will read the value of `errno` and print out the error message corresponding to its current value. User-written code may also set `errno` to take advantage of this standard error-handling mechanism.

In a program with multiple pthreads, the global `errno` variable is replaced by a *per-pthread* `errno`. When a failing system or library call sets the per-pthread `errno`, the other pthreads' `errno`s are not affected. The selection of per-pthread or per-process `errno` is determined at compile time by the preprocessor symbol `_REENTRANT`:

- If `_REENTRANT` is defined at the point the file `<errno.h>` is included, the symbol `errno` refers to the per-pthread `errno`.
- Otherwise, the symbol `errno` refers to the global (per-process) `errno`.

The symbol `_REENTRANT` is defined in `<pthread.h>`. However, because some programs include other header files before `<pthread.h>`, you should always use the switch `-D_REENTRANT` on the command line when compiling a program that uses multiple pthreads. This symbol ensures that the correct versions of call prototypes and preprocessor symbols are pulled in from header files.

In the current release, the libraries provided with Paragon OSF/1 are not consistent in their use of the two different *errno*s:

- *libpthreads.a* sets and references only the per-pthread *errno*.
- *libc\_r.a* sets both the global and the per-pthread *errno*, but only references the per-pthread *errno*.
- All other libraries set and reference only the global *errno*.

This creates an inconsistency of different *errno*s between pthreads programs and most libraries. For example, *libnx.a* uses the global *errno*. This means that if a call to *libnx.a* fails in a program compiled with `_REENTRANT`, the *errno* value it returns is not usable by the calling pthread. The library call sets the global variable *errno*, but the calling pthread can see only its own pthread-specific *errno* (whose value does *not* reflect the error in the *libnx.a* call that just failed).

In general, this means that code compiled with `_REENTRANT` cannot use *errno* values returned by non-thread-safe libraries. However, because some non-thread-safe libraries make calls to the standard C library, some *errno* values *are* usable. For example, `cread()` (in the non-thread-safe library *libnx.a*) calls `read()`. If you link with *libc\_r.a*, any error that occurs in `read()` will be reflected in the per-pthread *errno* and will be visible to the calling pthread. But any error that occurs in the `cread()` call itself (before or after the call to `read()`) will *not* be visible to the calling pthread.

## perror() and nx\_perror()

The `perror()` and `nx_perror()` calls print an error message based on the following versions of *errno*:

- The `perror()` call in *libc.a* uses the global *errno*.
- The `perror()` call in *libc\_r.a* uses the per-pthread *errno*.
- The `nx_perror()` call in *libnx.a* uses the global *errno*.

Because `nx_perror()` is not a thread-safe call, you should be sure to protect it with a mutex so that no two pthreads can call it at once.

## Calling exit()

Calling `exit()` when an error occurs terminates the entire process and closes any opened files. For this reason, it's a bad idea to call `exit()` on an error returned from a system call inside a pthread. Instead, you should call `pthread_exit()` to terminate the failing pthread and return a value indicating failure to the pthread that calls `pthread_join()`. The pthread that calls `pthread_join()` should use this information to shut down all other pthreads properly.



## Use of Underscore Versions of Paragon System Calls

The standard versions of most Paragon system calls in *libnx.a* terminate the calling process when an error occurs and send a message to standard error describing the error. For example, `isend()` will call `nx_perror()` to print out the error message, then call `exit()` to terminate the process. This implies that if a pthread causes an error in an `isend()` call, this error will kill the rest of the pthreads in the process. If this error occurs inside a `hsend()` and `hrecv()` handler routine, it will hang the handler routine.

For this reason, in programs with multiple pthreads you should always use the underscore versions of these calls instead. For example, calling `_isend()` will return -1 and set the global `errno` in case of error, instead of terminating the process. The calling pthread can then use this information to shut down the rest of the pthreads cleanly. See “Handling Errors” on page 4-42 for more information on underscore calls.

## Catch Signals Causing Core Dump by Default

The default action for the signals `SIGFPE` (floating point exception) and `SIGSEGV` (segmentation violation) is to core dump, terminate the process, and terminate the application. This will also kill all pthreads in the application.

When this occurs, you want to be able to figure out which pthread was responsible for this problem. For synchronous signals, the best way to do this is to install a signal handler to catch them and print out the pthread ID when the signal is received. For asynchronous signals, use `sigwait()` to catch them and then terminate all pthreads gracefully.

## When One Pthread Hangs

When debugging a program with multiple pthreads, always keep track of every active pthread, as much as possible, to detect the hang of a single pthread. Knowing which pthread has hung will help you determine the cause of a program hang.



# Designing a Parallel Application



7

## Introduction

This chapter describes some general design guidelines to follow when writing parallel applications. However, the best way to become skilled in parallel programming is to do it. With that in mind, this chapter presents three examples of parallel applications. Each example is intended to illustrate a different aspect of parallel design technique.

- The first example is a nearly-perfectly-parallel application that evaluates a definite integral to calculate  $\pi$ . This example illustrates how a sequential application can be ported to a parallel system with minimal effort. Much of the sequential algorithm can be maintained. The parallel design consists of separating the user interface from the core computation and then distributing that core computation onto the nodes.
- The next example is the multiplication of a matrix by a vector. In addition to the numerical technique, this example illustrates the use of parallel file I/O by assuming a matrix that is too large to reside entirely in memory.
- The third example solves a classic computer science problem called the N-Queens problem. Given a chess board with  $N \times N$  grid locations, where can you place  $N$  queens so that no queen is under attack? This example illustrates a technique called control decomposition. This technique also appears in more complicated real-world applications such as electronic design rule checking.

## The Paragon™ OSF/1 Programming Model

As described in Chapter 1, the Intel supercomputer is a distributed-memory parallel computer with a high-speed interconnect network. The following characteristics of the system should be kept in mind when designing or porting applications:

- The system is made up of an ensemble of processor/memory pairs called *nodes*. The nodes do not share memory. They present a single system image (for example, a process running on one node can send a signal to a process running on another node), but the nodes operate independently of each other.
- All the nodes are fully connected. They communicate with each other and the host by passing messages.
- Each node executes its own program. In many applications, it turns out that each node executes the same program on a different set of input data. There may be some conditional code that identifies one or more nodes that perform special actions.

These characteristics influence the design of parallel applications, as described in the remainder of this chapter.

## Parallel Programming Techniques

Parallel applications have varying degrees of parallelism. A *perfectly-parallel* application is one that requires no internode communication. In a perfectly-parallel application, if you double the number of nodes, you halve the computation time.

Most applications involve a mix of computation and internode communication; in these applications, increasing the number of nodes reduces the computation time, but can never yield a “perfect” speedup. The more time a program spends communicating instead of computing, the less speedup you get by adding nodes.

In order to get the best possible speed from a parallel program, you must design it so that each node spends as much time as possible computing, and as little time as possible communicating (or waiting for communication). Here are some techniques that can help you to do this:

- Separate the user interface from the computational parts of the code.
- Distribute the computation among the nodes so that their computational load is evenly balanced.
- Write your application so that you can run it on more nodes, thus improving performance, without having to recode.
- Design your internode communication such that the nodes spend as little time in communication (or waiting for communication) as possible.

The following sections tell you more about these techniques.

## Separating the User Interface from the Computation

To have each node do as much computation, and as little non-computational work, as possible, you should analyze the algorithm and separate the user interface from the computational kernel. You can designate one of the nodes to handle the user interface, or put the user interface in the application's *controlling process* (see "The Controlling Process" on page 4-21 for information on this process). In either case, the part of the program that handles the user interface and the part of the program that does the computation communicate by passing messages.

In the  $\pi$  example, node 0 requests the number of integration intervals from the user. It then sends that number to the other nodes, and all the nodes do the calculation.

## Balancing the Load

You should keep all the nodes busy and have them finish at the same time, because if some nodes have to wait for others to finish, they're wasting cycles doing nothing. Analyze your application and distribute the computation among the nodes so that their computational load is evenly balanced.

The process of distributing a problem among the nodes is referred to as *problem decomposition*, or just *decomposition*. There are two kinds of decomposition: *domain decomposition* and *control decomposition*.


## Domain Decomposition

In domain decomposition, the input data (the domain) is partitioned and assigned to different processors. How you divide and distribute the data among the nodes can have a significant effect on the efficiency of your application.

For example, consider an application that performs image enhancement (see Figure 7-1). Because some parts of the image may be more detailed than others, they will require more processing. The shaded portion of Figure 7-1 shows the work done by node 0. If you divide the image sequentially among the nodes, as shown in the top half of Figure 7-1, some nodes may get a partition that requires a lot of work and other nodes may get a partition that requires little or no work. In the top half of Figure 7-1, node 0 gets a lot of work and node 7 gets no work at all. This is inefficient.

You can often achieve better load balancing by dividing the image into smaller partitions and then distributing the partitions sequentially among the nodes, as shown in the bottom half of Figure 7-1. This is analogous to dealing out the partitions like cards in a deck; it spreads out the work more evenly among the nodes. As the bottom half of Figure 7-1 shows, each node gets some slices that require a lot of work, some slices that require a moderate amount of work, and some slices that require no work. This is more balanced and efficient for this type of problem, and may be appropriate for your problem as well.

Poor load balancing: Nodes 0 through 3 get most of the work. Nodes 4 through 7 have little or nothing to do.



Good load balancing: The partitions in the domain are dealt out to the nodes like cards from a deck. Now, each node has approximately the same amount of work.




Figure 7-1. Using Domain Decomposition to Achieve Load Balancing

## Control Decomposition

Control decomposition, on the other hand, divides the *tasks* to be performed rather than the *data*. For many problems, this is a more natural decomposition.

For example, consider a tree-search used in a game-playing algorithm. Assume that you're at some mid-level of the tree. You could approach the problem as a domain decomposition and divide the current branches among the nodes. Each node would then follow its branch down to the leaves and then return the leaves as an answer. The leaves in this case are the possible moves. Depending on the current state of the game, some of the branches may be quite involved and require a great deal of processing. Other branches may be simple. The result is that some nodes finish before others. This is a poor problem decomposition.

Approaching this problem as a control decomposition achieves better load balancing. In a control decomposition, you think of the branches not as data partitions but rather as tasks that need to be performed.

To manage these tasks, you have to introduce a little bureaucracy. Assign one node as a manager node. This manager node then gives tasks to idle nodes. When the node finishes a task, it reports its answer and requests another task. It's this "reporting for duty" that characterizes a control decomposition.

The manager node must, of course, do some initial setup. For example, it may follow the tree down until the number of branches exceeds the number of available nodes by some predetermined factor.

This method produces the best results when the tasks assigned near the end of the problem are about the same size. For example, if one of the last tasks assigned was a very long task, the other nodes may be idle while that last node finishes.

The N-Queens example (presented later in this chapter) shows control decomposition.

## Making the Program Independent of the Number of Nodes

You should write your application so that you can run it on more nodes, thus improving performance, without having to recode.

This method also turns out to be the most natural one to use when porting an existing sequential application. After you've separated the user interface from the core computation, you still have a sequential algorithm, but you can think of it as the special case of an application that runs on one node. Once you have done this, you can parallelize the computation part for an arbitrary number of nodes.

The  $\pi$  example illustrates this technique. The number of nodes appears only as the variable *nodes*.

## Designing Your Communication Strategy

You should design your internode communication such that the nodes spend as little time in communication as possible. This may involve running some tests to determine an optimal message length. Often, you can decrease the number of messages by increasing the size of each message. You may also be able to improve communication performance by using asynchronous message-passing calls, as described under “Asynchronous Send and Receive” on page 3-10.

## Using Global Operations

You should use the global operations, described under “Global Operations” on page 3-27, when possible. That section described a simple example of a global sum. Using `gdsun()` results in a significant improvement over having one node perform the global sum by explicitly collecting all the partial sums. Also, after the execution of the `gdsun()`, the global sum is available on each node.

The matrix\*vector example in this chapter uses another global operation called `gcolx()`. In that example, a large vector is distributed over the nodes. `gcolx()` collects the components from each node and constructs the complete vector on each node. As with `gdsun()`, the answer is available on each node.

## Using Alternate Node Topologies

The nodes in the Intel supercomputer are connected in either a hypercube or a mesh network. However, because of the specialized message-passing hardware in both architectures, communication with distant nodes is nearly as fast as communication with neighboring nodes. This means that you do not have to structure your application's communications as a hypercube or mesh; you can choose an alternate topology that makes more sense for your program. This can make your program easier to write and understand, at a tiny cost in performance.

When you use an alternate node topology, you embed your node topology (a *virtual topology*) into the nodes' actual network topology (the *physical topology*). One example of a virtual topology is the ring. This topology is useful in certain types of many-body calculations. The technique consists of partitioning the particles into groups and assigning each group to a different node. A node then calculates the state of its group. This state information is then passed to another node which



calculates the state of its own particles and takes into account the state received from the previous node. The state information moves from node to node around a ring. You can implement a ring topology by writing a function like this one:

```
succ(int n)
{
    int maxnode;
    maxnode = numnodes() - 1;

    if ( (n >= 0) && (n < maxnode) )
        return(n+1);
    else if (n == maxnode)
        return(0);
    else
        return(-1);
}
```

Given a valid node ID ( $n$ ), this function returns the node ID of the successor of node  $n$  in a ring embedded in a partition with `numnodes()` nodes. Else it returns -1. (The predecessor function is similar.) A node can send a message to process type 0 on its successor node with the following `csend()` call:

```
csend(MSGTYPE, buf, sizeof(buf), succ(mynd()), 0);
```

## Example Application: Calculating pi

This application uses an  $n$ -point quadrature rule to evaluate the following definite integral:

$$\pi = \int_0^1 \frac{4}{(1+x^2)} dx$$

Admittedly, using the power of an Intel supercomputer for such a simple application is overkill, but the application demonstrates concepts that are just as valid for more challenging problems.

Here is a sequential program (written in Fortran) that evaluates the above integral. The source for this program is available on the Intel supercomputer in */usr/share/examples/fortran/pi/piserial.f*. Note that the user interface consists only of a *read* statement that solicits the number of intervals.

```

      program piserial
      double precision h,sum,x,pi,f,a
      integer n

      c Define the function
      f(a) = 4.0d0/(1.d0 + a*a)

      c Input the number of intervals.
      1  print *, ' Enter number of intervals:'
      read(5,*,end=100) n

      c Calculate the scaling factor
      h = 1.d0/n

      c Integrate.  The value of x used to calculate the slice is
      c the value at the midpoint of the integration slice.
      sum = 0.d0
      do 10 i = 1,n
      x = h * (dble(i) - 0.5d0)
      sum = sum + f(x)
10  continue
      pi = h * sum

      c Output the answer
      print *, ' The value of pi for',n,' intervals is',pi
      goto 1

      c
      c Terminate
      100  stop
      end

```

In the parallel version of this program, each node performs a portion of the integration. The decomposition is a domain decomposition that “deals out” the work, as illustrated in Figure 7-2. For example, if you choose 16 nodes and 512 points, each node gets 32 points. The first point goes to node 0, the second point goes to node 1, and so on through the 16th point, which goes to node 15. The 17th point goes to node 0, the 18th point goes to node 1, and so on until all the points have been dealt out. (It is not strictly necessary to deal out the work in this way, because the integration work is evenly balanced. However, since the data is calculated by each node, it is just as easy to deal out as not, and this example deals out the data to give you an example of this technique.)




Figure 7-2. The Decomposition Used for the pi Example

Here is the parallel version of the program. The source for this program is available on the Intel supercomputer in */usr/share/examples/fortran/pi/pinode.f*; differences from the serial version are shown here in **boldface**.

```

    program pinode
    include 'fnx.h'
    double precision h,sum,x,pi,f,a,tmp
    integer n
    integer nodes, iam, intsiz

    data intsiz / 4 /

c Define the function
    f(a) = 4.0d0/(1.d0 + a*a)

c Do some bookkeeping
    iam = mynode()
    nodes = numnodes()

1   if(iam .eq. 0) then
c   Input the number of intervals.
    print *, ' Enter number of intervals:'
    read(5,*,end=100) n
    call csend(300,n,intsiz,-1,0)
    else
    call crecv(300,n,intsiz)
    endif

c Calculate the scaling factor
    h = 1.d0/n

c Integrate. The value of x used to calculate the slice is
c the value at the midpoint of the integration slice.
    sum = 0.d0
    do 10 i = iam+1,n,nodes
        x = h * (dble(i) - 0.5d0)
        sum = sum + f(x)
10   continue
    pi = h * sum
    call gdsun(pi,1,tmp)

    if(iam .eq. 0) then
c Output the answer
    print *, ' The value of pi for',n,' intervals is',pi
    endif

    goto 1
c ,

```

```

c Terminate all nodes
100  i = kill(0, 9)
      end

```

Note that the parallel version is not much longer than the sequential version. Note also that the decomposition takes place entirely in the `do` statement. The sequential version is:

```
do 10 i = 1, n
```

while the parallel version is:

```
do 10 i = iam+1, n, nodes
```

If you run the application on more nodes, you don't have to change one line of the node program!

In the parallel version, only node 0 interacts with the user. The other nodes do only calculation. If the `print` and `read` statements were not surrounded with `if(iam .eq. 0)then ... endif` statements, then when you ran the program on 100 nodes you would have to input the number of intervals 100 times and see the answer 100 times!

## Example Application: Matrix\*Vector Multiplication

The following example computes the matrix-vector product  $y = Ax$ , where  $A$  is an  $n \times n$  matrix and  $x$  and  $y$  are vectors with  $n$  components. In addition to the numerical technique, this example illustrates the use of the parallel file I/O calls.

The matrix  $A$  is assumed to be too large to fit in the node's memory, requiring an "out-of-core" multiplication. For simplicity,  $n$ , the number of rows in the matrix, is assumed to be divisible by  $p$ , the number of nodes in the application. The number of rows per node,  $n/p$ , is referred to as  $m$ .

The problem decomposition is again a domain decomposition. Each node collects all of  $x$ , but then takes only a portion of  $A$  (specifically  $m$  rows) to form its portion of the product vector. There is no attempt to "deal out" the rows of  $A$ .

The vector  $x$  is initially divided among the nodes. (This example assumes that each node has obtained its portion of  $x$  before this routine is called.) Each node contains  $m$  components of  $x$ . Node 0 has components 1 through  $m$ ; node 1 has components  $m + 1$  through  $2*m$ , etc. (In general, node  $Z$  has components  $(Z-1)*m$  through  $Z*m$ .) The answer, the vector  $y$ , will be stored in the same way.

The matrix  $A$ , which is too large to fit in a single node's memory, is also divided among the nodes. It is initially stored in a file called *matrix*. The elements of the matrix are stored in the file by rows, as follows:

```
A(1,1), A(1,2), ... A(1,n), A(2,1), A(2,2), ... A(2,n), ... A(n,1), A(n,2), ... A(n,n)
```

Each row of the matrix  $A$  has  $n$  elements of length  $REALSIZE$  bytes, and so each row takes up  $n*REALSIZE$  bytes in the file. Each node is responsible for  $m$  rows in the matrix; it reads its portion of the matrix from the file by first moving the file pointer to  $mynode()*m*n*REALSIZE$  bytes from the beginning of the file, then reading  $m$  rows of  $n*REALSIZE$  bytes each beginning at that point.

Here is the code that collects  $x$ , reads the node's portion of  $A$ , and performs the multiplication:

```

      subroutine matvmul(m, n, x, y, xtotal, arow)
      integer REALSIZE
      parameter(REALSIZE = 4)
      integer ncnt, fileptr, xlens(128)
      integer m, n
      real x(m), y(m), xtotal(n), arow(n)
c
c  m is n/p where n is the dimension of A
c  and p is numnodes()
c
c  Collect all of x on each node.
      do 3 i = 1, numnodes()
          xlens(i) = m*REALSIZE
3      continue
      call gcolx(x, xlens, xtotal)
c
c  Open the file and seek to the appropriate location
      open(unit=10, file = 'matrix',
+         form = 'unformatted')
      fileptr = lseek(10, mynode()*m*n*REALSIZE, 0)
c
c  Read the rows and use the BLAS call sdot() to do
c  the multiplication.
      do 10 i = 1, m
          call cread(10, arow, n*REALSIZE)
          y(i) = sdot(n, arow, 1, xtotal, 1)
10     continue
      .
      .
      .

```

This subroutine takes the following parameters:

- |     |   |
|-----|---|
| $m$ | The size of each node's portion of the matrix $A$ and the vector $x$ ( $n/p$ ).                               |
| $n$ | The number of rows and columns in the entire matrix $A$ and the number of elements in the entire vector $x$ . |
| $x$ | This node's portion of the vector $x$ ( $m$ elements).  |

<i>y</i>	This node's portion of the result vector <i>y</i> ( <i>m</i> elements).
<i>xtotal</i>	A temporary array used to hold the entire vector <i>x</i> ( <i>n</i> elements).
<i>arow</i>	A temporary array used to hold one row of the matrix <i>A</i> ( <i>n</i> elements).

The subroutine first calls `gcolx()` to collect the nodes' portions of *x* together into the array *xtotal*. It then opens the file containing *A*, moves the file pointer to the beginning of the section of the file that belongs to this node, and then reads *m* rows from the file. After reading each row, it uses the BLAS (Basic Linear Algebra Subroutines) routine `sdot()` to perform the dot product between the current row and the vector *x*, storing the result (a scalar) into the appropriate element of the vector *y*.

## NOTE

You must use the `-lkmath` switch on the `if77` command line to link in the library that contains `sdot()`.

See the *Paragon™ Fortran System Calls Reference Manual* for more information on `gcolx()`; see Chapter 5 for information about parallel file I/O; see the *CLASSPACK Basic Math Library User's Guide* or *CLASSPACK Basic Math Library/C User's Guide* for more information on `sdot()`.

## Example Application: The N-Queens Problem

This application collects all the board configurations that solve the N-Queens problem. This problem is: "Given an *N* x *N* chess board, where can you place *N* queens so that no queen can capture any other?" In chess, queens attack in straight lines along the X, Y, and diagonal directions.

The N-Queens problem is typical of problems for which there is no analytical solution. Instead, there exists a large set of candidate solutions. You test each solution and accept those that pass.

The difficulty lies in the enormous size of the candidate set. For example, an 8 x 8 chess board has 64 squares. The total number of possible positions for 8 queens can be represented as the *combination* of *n*=64 things taken *m*=8 at a time. The formula for the number of combinations is:

$$n! / ( m! * (n-m)! )$$

which evaluates to  $2^{32}$  possibilities. Even on a state-of-the-art sequential computer, it would take several hours to check every one of those combinations.

Even before you begin thinking about an algorithm, however, you can eliminate a large number of possibilities. For example, any solution that has more than one queen in the same column is invalid. This reduces the number of possibilities to  $8^8$  or  $2^{24}$ .

This section shows how to use an Intel supercomputer to evaluate those  $2^{24}$  possibilities. You can arrange the possibilities into a tree. The technique involves following a tree down until it either reaches a dead end (an invalid state) or until it reaches a leaf (a valid solution). Figure 7-3 illustrates such a tree. To make the figure simpler, the chess board is shown as  $4 \times 4$ . Instead of  $2^{24}$  possibilities, you have  $2^8$ .

The root of the tree (the zero level) is the null board — no queens present. The next level (the first level) consists of states where a queen is in each of the positions that make up the first column. In Figure 7-3, there are four of those. In an  $8 \times 8$  board, there would be eight.

The next level (the second level) consists of states with two queens on the board, one in the first column and one in the second. In Figure 7-3, there are four of those under each second level state. Notice, however, that some states are already invalid. There is no need to follow the tree any further down this branch. In Figure 7-3, the two leftmost states in the second level are invalid. The second state in the first level has three dead ends in its second level.

You can see how the algorithm is going. Some paths are going to finish early because they reach dead ends. Others are going to take longer and reach the solutions at the leaves. This is a problem for control decomposition.

Manager/worker decomposition (a type of control decomposition) is a useful way of achieving balanced computational loads when the application consists of a large number of tasks that are of varying length. Because there is no way of determining up front what the length of the task is, the method consists of dividing the application into a large number of tasks (more than the number of nodes) and then assigning tasks to individual nodes as the node becomes available.

One way of generating the task is for the manager node to follow the tree down until the number of states is larger than the number of available nodes. As a further enhancement, the manager node may even enlist the aid of the other nodes when doing this initial processing.

Then, the manager node assigns a state to a node. The node follows that state down the tree and collects all the possible solutions. When the node finishes, it reports its solutions, if any, and requests more work. In the case of a  $4 \times 4$  board, the tree is shallow and there are only two solutions. An  $8 \times 8$  board results in 92 solutions.

The directory `/usr/share/examples/c/nqueens` contains a C version of the  $8 \times 8$ -Queens problem. The example is written in C because the N-Queens algorithm makes use of recursion.

In this example, a task is represented as a partially-filled board (only the first few columns contain queens) given to one of the nodes. The example as described here runs on four nodes. Node 0 is the manager, and nodes 1 through 3 are the workers. The manager is assigned a certain number of columns (in this example, two) and creates partial boards by placing queens on the board, one for each column it is assigned. When the manager controls two columns of an  $8 \times 8$  board, it creates 64 partial boards.






Figure 7-3. The N-Queens Solution Tree for a 4 x 4 Board

Also, in this example, the manager does not create the boards intelligently. For example, the manager will create a board with two queens in the same row. If a worker gets a partial board that contains invalid queens (such as two queens in the same row), the worker immediately throws the board away and requests another.

The manager creates boards by counting in a radix equal to the number of rows in the board. Each digit in the resulting number represents a column with the least significant digit being column 0. The value of the digit is the row position of the queen. Hence, 00 represents two queens in row 0, and 01 represents one queen in row 0 of column 0 and another queen in row 1 of column 1.

The workers signal their availability by sending a "ready" message to the manager. This is a zero length message of type **READY**. When the manager receives a **READY** message, it determines who sent it, then sends a partial board to that node as a message of type **TASK**. The manager keeps doing this until it has no more partial tasks to assign. Finally, the manager waits until all workers are idle (that is, it receives a **READY** message from every worker) and then sends a final message with the special value **FINISHED** to all workers.

Here are the key lines that implement the manager control.

```

/* This is the manager part */
  if (!iam) {          /* If I am node 0 */
    printf("\n\n\n");
    printf("\nSTARTING ... \n");

/* Manager keeps a count of how many workers are available
and sends out boards to a worker when the worker identifies
itself as READY. The manager uses the routine get_board() to get
a new board. There are no more new boards when this routine
returns DONE.*/

    while ( get_board(board) != DONE ) {
      crecv(READY, NULL, 0);
      nodenbr = infonode();
      msgcount++; /* Count how many nodes are ready */
      csend(TASK, board, sizeof(twoD), nodenbr, 0);
      msgcount--; /* When a node gets a task, it is no longer
ready for another. Hence, decrease
msgcount */
    }

/* Wait for all workers to be free (the msgcount must be equal
to the number of worker nodes) */

    while(msgcount != nodes-1) {
      crecv(READY, NULL, 0);
      msgcount++;
    }

```

```

/* Send the FINISHED message to all nodes and then say goodbye */

    board[0][0] = FINISHED;
    csend(TASK,board,sizeof(twoD),-1,0);
    goodbye();
}

```

The manager does not know if a worker has found a solution or not, and the workers do not know how many initial boards there are. When a worker receives a partial board, it first checks for the special value **FINISHED**, and calls **goodbye()** if it finds this value. (The **goodbye()** routine prints a summary message in the output file, closes the file, and exits.) Next, the worker checks that the queens already on the board are valid. If they are, the worker finds all the solutions that exist with that partial board by recursively calling **move\_to\_right()**. When the worker finds a solution, it writes the solution to a file called *queens.out*. This file was opened by all nodes in mode **M\_LOG (1)**, which is the mode in which all nodes have a common file pointer and access the file on a first-come first-served basis.

Here are the key lines that implement the worker control.

```

else {
/* This is the worker part. */

/* Each node enters an infinite loop where it receives a partial
board and checks whether that partial board contains valid
queens. If the board contains a FINISHED message, the node
cleans up and exits by calling goodbye(). If the board contains
invalid queens, the node considers itself done with the task.
Otherwise, it tries to place a queen in the next column by calling
move_to_right(). This routine will find all possible solutions
given the initial board. */

    for(;;) {
        csend(READY,0,0,0,0);
        crecv(TASK,board,sizeof(board));
        if(board[0][0] == FINISHED) {
            goodbye();
        }
        if ( chk_board(board) ) {
            move_to_right(board,0, MCOLS);
        }
    }
} /* end of else */

```

There are many opportunities for optimizing this algorithm. For example, you could write the manager in such a way that it only gave workers boards that had the potential of containing one or more solutions. In addition, the manager could mark positions on the board that are invalid due to the presence of the initial queens, and the worker would not have to check those.

The file *queens.out* contains copies of all the 92 solutions for the 8-Queens problem. Each board is preceded by a header that identifies the node that found the solution and the number of solutions found so far by the node. Finally, the total number of solutions is printed. The tail of the file looks as follows:

```

      .
      .
      .
Node 1 found solution 30

```

```

  0 1 2 3 4 5 6 7
0 - - - Q - - - -
1 - - - - - Q - -
2 - - - - - - - Q
3 - - Q - - - - -
4 Q - - - - - - -
5 - - - - - - Q -
6 - - - - Q - - -
7 - Q - - - - - -

```

```

Node 2 found solution 31

```

```

  0 1 2 3 4 5 6 7
0 - - - - Q - - -
1 - - - - - Q - -
2 - - - Q - - - -
3 Q - - - - - - -
4 - - Q - - - - -
5 - - - - - - Q -
6 - - - - - Q - -
7 - Q - - - - - -

```

```

Node 3 found solution 31

```

```

  0 1 2 3 4 5 6 7
0 - - - - Q - - -
1 - - Q - - - - -
2 - - - - - - - Q
3 - - - Q - - - -
4 - - - - - - Q -
5 Q - - - - - - -
6 - - - - - Q - -
7 - Q - - - - - -

```

```

Total solutions = 92

```

If you want to investigate another manager/worker application, look at the *triangle* program in */usr/share/examples/c/triangle*. Its operation is described in a *README* file.

# Improving Performance

8

## Introduction

This chapter presents some techniques you can use to improve the performance of your parallel applications. It includes the following sections:

- Single Node Performance.
- Multi-Node Performance.
- I/O Performance.

In general, however, the best thing you can do to improve performance is to choose an efficient numerical method and algorithm for solving your problem. A good numerical method and an efficient algorithm will always give better performance than a poor method and algorithm. This is true even if the good method is implemented in a high-level language and the poor method is implemented in hand-coded assembly language.

Another general performance technique is to use the Paragon system's profiling and performance analysis tools to help pinpoint the parts of your application that could benefit the most from optimization. See the *Paragon™ Application Tools User's Guide* for information on the available tools.

## Single Node Performance

This section discusses things you can do to increase the speed of calculation (MFLOPS or GFLOPS) on each node. Many of these are general performance-improvement techniques that you can use on any computer; some of them are specific to the i860® microprocessor. Techniques discussed in this section include:

- Use profiling tools.
- Avoid repeated use of system calls.
- Avoid virtual memory paging.
- Use compiler optimizations.
- Increase problem size.
- Access contiguous memory locations.
- Use caching wisely.
- Use optimized libraries.
- Use assembly language subroutines.
- Avoid error checking (C language only).

## Use Profiling Tools

The Paragon system comes with the **prof** and **gprof** profilers, and their graphical versions **xprof** and **xgprof**. You can use these tools to help track down the parts of your application that are consuming the most time and concentrate your optimization efforts there. See the *Paragon™ Application Tools User's Guide* for more information on these tools.

## Avoid Repeated Use of System Calls

Don't make a system call twice if once will do. This is an obvious performance improvement technique, but unfortunately it is missing from many applications. For example, a process may need its node number and process type to do message passing. Avoid using **mynode()** and **myptype()** each time you need those numbers. Instead, invoke each once and store their values in variables.

## Avoid Virtual Memory Paging

The Paragon OSF/1 operating system provides *virtual memory*, which lets you use more memory than is physically available on the node. When a program tries to allocate a memory space that is larger than the node's available free memory, one or more 8K-byte *virtual memory pages* that haven't been referenced recently are *paged out*. This means that their contents are written to disk and replaced with the new data. Later, when the program references data in the paged-out memory section, a different section is paged out and the old data is *paged in* (read back from disk) in its place.

Although virtual memory makes it possible for the system to support multiple users and very large programs, you should try to avoid it when you can. Accessing pages of virtual memory that are not currently paged in is much slower than physical memory and generates a lot of disk activity. Try to reduce the memory used by your application until it fits in physical memory, including dynamically-allocated buffers and system message buffers (see "Understand Message-Passing Flow Control" on page 8-13 for information on the sizes of system message buffers).

You can use the `vm_stat` command to get information about your application's memory usage. See the *Paragon™ Commands Reference Manual* for information on this command.

Once you have reduced your application so that it fits in physical memory, you may be able to use the `-plk` switch to lock parts of your application into physical memory. This reduces paging and improves message-passing latency, but has certain consequences; see "Process Locking" on page 8-15 for more information.

## Use Compiler Optimizations

When you compile a program, you can use *compiler optimization switches* to tell the compiler what techniques to use to optimize your code. Optimization can produce a compiled program that does the same work in less time by making better use of the processor's special features. However, optimization can sometimes produce a program that runs more slowly or produces wrong answers, so it must be used carefully.

The compiler optimization switches you can use and compiler-specific code changes you can make depend on which language you program in and the revision level of the compiler. See the *Paragon™ Fortran Compiler User's Guide* or *Paragon™ C Compiler User's Guide* for complete information on compiler optimizations for your specific compiler. However, here are a few general hints:

- Experiment with the **-O** switch, which controls the level of compiler optimization:
  - Level 0 performs no optimization.
  - Levels 1 and 2 perform straightforward optimizations that should always result in improvement.
  - Levels 3 and 4 attempt to make use of the i860® microprocessor's *pipelining* and *dual-instruction* modes to improve performance; whether or not they improve your program's performance, and by how much, depends on the characteristics of your program.

In some cases, different parts of the program should be compiled with different optimization levels.

- Try the **-Mvect** switch to invoke the *vectorizer*. The vectorizer attempts to rearrange your code to allow more efficient use of pipelining. You can get better results out of the vectorizer if your innermost loops have the following characteristics:
  - The loop index increments the first dimension in Fortran, or the last dimension in C.
  - Arrays are accessed with unit stride.
  - The number of iterations within the loop is not too small.

Also, **if** tests and subroutine calls should be avoided within the loop. See the *Paragon™ Fortran Compiler User's Guide* or *Paragon™ C Compiler User's Guide* for specific examples of code changes you can make.

- If your application does not depend on strict IEEE semantics for mathematical operations, try the **-Knoieee** switch. This switch provides much faster mathematical operations than those provided by the default IEEE math library, but may result in slightly decreased accuracy and different behavior in exceptional circumstances (operations on 0 or infinity and NaNs).
- Use the **-Mnostride0** switch, unless your program accesses arrays with zero stride (that is, incrementing the array pointer by 0 in each loop iteration). There are some important compiler optimizations that are only possible if the compiler knows the code does not do this.



## Increase Problem Size

Once you have optimized a program's single-node performance, you may find that running the program on a larger number of nodes with the same data set gives a lower per-node performance. This can occur because the per-node vector size has gone down, reducing the efficiency of the i860 microprocessor's pipelines. To avoid this problem, you can increase the problem size as you increase the number of nodes, or you could even write two different inner loops—one optimized for short vectors, the other optimized for longer vectors.

## Access Contiguous Memory Locations

Whenever you access memory, try to access contiguous memory locations. In particular, whenever your program reads or changes the value of an array element in memory, try to be sure that the next array element it reads or changes is adjacent to the previous one in memory. This is important because the i860 microprocessor accesses memory in 4K-byte *physical memory pages*. Once you have read from a physical page, another read from the same page takes only one more cycle, but a read from a different page takes 10 to 14 more cycles. Every cycle spent switching from one page to another is a cycle that can't be used for calculation.

To keep your memory accesses within a physical page, you can use some of the following techniques:

- Group a series of memory reads out of the same array.
- Do consecutive references across the rows (C) or down the columns (Fortran) of matrices. (In C the rightmost index of an array varies the fastest, while in Fortran the leftmost index varies fastest. This means that, for example, if you distribute the elements of a two-dimensional array among the nodes, you should give out rows in C and columns in Fortran.)
- "Strip-mine" loops, so that you do several accesses to the same array at a time. For example, you should read several elements from vector *A*, then several from *B*, instead of reading *A[1]*, *B[1]*, *A[2]*, *B[2]*, and so on.

You should also try to group reads and writes. Once you have read from a page, a following write takes about 6 more cycles than a following read. Switching from write to read also takes about 6 cycles.

## Use Caching Wisely

The i860 microprocessor has a 16K-byte *data cache* for recently-accessed memory locations. Whenever you read or change a bit in memory, a 32-byte area of memory containing that bit is copied into the cache. (This 32-byte area is called a *cache line* and always begins on a 32-byte boundary.) When you access memory that is already in the cache, the access is very fast. However,

whenever a new cache line is copied into the cache, cache lines that have not been accessed recently are written back to memory (if necessary) and removed from the cache to make room. Try to arrange your code so that all operands for a loop can be accommodated in the cache at the same time.

The i860 microprocessor also has an *instruction cache* (4K bytes on the i860 XR, 16K bytes on the i860 XP) which is used to hold program instructions once they have been fetched and decoded from memory. You can use this cache in two ways:

- Try to keep your loops small. If the code for an entire loop fits in the instruction cache, the loop can execute very quickly.
- In an *if/else* block, try to put the code that is used more often in the *if* part and the code that is used less often in the *else* part. The instruction cache works in a “lookahead mode,” and when pre-fetching instructions will fetch the code immediately following the *if*. If the *else* branch is executed instead, the *if* branch code must be discarded from the cache.

Both these techniques can be used in high-level languages as well as assembly code.

Note that the data cache, the instruction cache, the physical memory page, and the virtual memory page are separate functions that have different sizes and different effects. In general, cache management is handled by the compiler, but you should try to arrange your code to make the compiler's job easier.

## Use Optimized Libraries

Several optimized libraries of math and utility functions are available with Paragon OSF/1. These libraries have been carefully hand-tuned to give the best possible performance; you can save time and increase efficiency by using routines from these libraries rather than writing the equivalent code yourself. The available libraries include:

- The Basic Math Library (*libkmath.a*). This library is a standard part of Paragon OSF/1; it includes optimized BLAS (Basic Linear Algebra Subroutines) and FFT (Fast Fourier Transform) routines. See the *CLASSPACK Basic Math Library User's Guide* or *CLASSPACK Basic Math Library/C User's Guide* for more information.
- The Signal Processing Library (*libsignal.a*). This library is an optional product; it includes optimized vector and signal-processing routines. See the *CLASSPACK Signal Processing Library User's Guide* or *CLASSPACK Signal Processing Library/C User's Guide* for more information.

Note that these are *single-node libraries*; they improve the numeric performance of each node of your program, but do not affect its multi-node performance.

## Use Assembly Language Subroutines

Re-writing key routines in the i860 microprocessor's assembly language can sometimes bring significant performance benefits. See the *Paragon™ i860® 64-Bit Microprocessor Assembler Reference Manual* for information on the assembler.

## Avoid Error Checking (C Language Only)

In C, there are two versions of most calls in *libnx.a*: the standard version and the *underscore version* (for example, the underscore version of `crecv()` is `_crecv()`). When you call the standard version, the call checks for certain error conditions before it returns; if an error is detected, the call terminates your program with an error message. The underscore version works the same as the non-underscore version, but if an error occurs, the call simply returns the value -1 and sets the external variable *errno* to a value that describes the error. This is useful if you want to handle an error yourself and not let the system do it. But if you are confident that your program is working, you may choose to use the underscore version and *not* check the return value, thereby improving performance. (If an error *does* occur, unexpected and difficult-to-debug behavior will result, so use this technique with caution.)

## Multi-Node Performance

This section discusses things you can do to increase the efficiency of applications running on multiple nodes, including:

- Use dynamic memory allocation for large arrays.
- Avoid serializing calls.
- Use ParaGraph.
- Maintain data locality.
- Overlap computation and communication.
- Avoid message buffering.
- Align application buffers.
- Understand message-passing flow control.

## Use Dynamic Memory Allocation for Large Arrays

You should always use *dynamic memory allocation* for large arrays. Dynamic memory allocation means allocating the memory for the array at run time, using the **ALLOCATE** statement in Fortran or the **malloc()** call in C. The alternative, *static* memory allocation, means declaring the array in the program source.

Dynamic allocation is important even on one node, but becomes more and more important as the number of nodes increases. The larger the array and the more nodes, the more performance can be improved by using dynamic memory allocation. If the array or number of nodes is large enough, the application may not run at all unless you use dynamic memory allocation.

For example, the following program fragment uses static memory allocation. It simply creates two 4M-byte arrays of **real\*4** (one in a common block, the other not) and initializes each element of each array to the element number.

```

parameter huge_size = 1024*1024
real*4 huge(huge_size), huge_common(huge_size)
common /giant/ huge_common
integer i

do 10 i=1,huge_size
    huge(i) = i
    huge_common(i) = i
10  continue

```

The equivalent code with dynamic memory allocation is as follows (changes are shown in boldface). With 16M bytes of memory on each node, this version runs as much as *ten times as fast* as the previous version on one node; it runs as much as *fifteen times as fast* on eight nodes. The more nodes, the greater the speedup.

```

parameter huge_size = 1024*1024
real*4 huge(huge_size), huge_common(huge_size)
pointer(p, huge)
common, allocatable /giant/ huge_common
integer i

allocate(huge, /giant/)

do 10 i=1,huge_size
    huge(i) = i
    huge_common(i) = i
10  continue

deallocate(huge, /giant/)

```

Note that a common block must be declared **ALLOCATABLE** before it is allocated with an **ALLOCATE** statement. A variable or array that is not part of a common block must be declared as a pointer-based variable with a **POINTER** statement before it is allocated with **ALLOCATE**; the corresponding pointer variable, in this case **p**, does not have to be used. See the *Paragon™ Fortran Language Reference Manual* for more information on these statements.

The reason statically allocated arrays cause your program to run slowly is that, since they are compiled into the program, the initial contents of the array must be obtained from the executable program on disk. Whenever a process on a compute node reads or changes a byte in a memory page of statically allocated data that it hasn't touched before, the data for that page may have to be *paged in*. (See "Avoid Virtual Memory Paging" on page 8-3 for an introduction to virtual memory paging.) A message requesting the initial contents of that page is sent to the node in the service partition where the compute process's parent process is running. The entire page—8K bytes—is then sent back to the compute process. This occurs even if the statically allocated data is uninitialized (all zeroes).

Sending these pages across the mesh takes time. Even worse, if many compute processes all want pages at the same time, the parent process's node can become overwhelmed, slowing the application drastically. The effect is magnified if the statically-allocated array is so large that parts of the operating system have to be paged out to make room for it; in this case, pages of the operating system have to go out at the same time pages of static data are coming in.

When you use **ALLOCATE** or **malloc()** to dynamically allocate an array, the memory is not associated with the program on disk. Instead, each node has its own copy of the array, and it doesn't have to be paged in. When a compute process touches a page of dynamically-allocated data it hasn't touched before, the page is simply allocated from the available node memory—no messages are sent. This greatly reduces traffic on the mesh and increases the performance of your application.

## Avoid Serializing Calls

Avoid using serializing calls repeatedly or on many nodes. A *serializing* call is one that relies on a single resource which can only service one request at a time (typically a daemon or server on the boot node). Using a serializing call once takes little time, but if many nodes in a large application call it at the same time the boot node can only service these requests one at a time. Each node must wait until the boot node services its request, which can cause the entire application to run slowly.

Many calls that perform I/O or make use of the file system, such as **stat()**, **chdir()**, and **chmod()**, are serializing calls, because they must communicate with the root file system server on the boot node. **getrusage()** is also a serializing call, because it sends a message to all the I/O nodes to get information on the caller's I/O activity. You can detect the presence of serializing calls in your program by profiling it. If common operations, especially I/O operations, are taking much more time than expected, they may be serializing calls.

Whenever possible, avoid overuse of serializing calls by having only one node make the call. For example, instead of having every node process call **stat()**, have one node call **stat()** and then use **gisum()** to distribute the information to the other nodes. Also, instead of having every node process

call `chdir()` when it starts up, have the controlling process call `chdir()` before creating the node processes. You can avoid serialization in I/O by using the I/O mode `M_RECORD` or `M_GLOBAL`, as described under “Use the Appropriate I/O Mode” on page 8-24.

## Use ParaGraph

The Paragon system comes with the ParaGraph performance visualization tool. You can use this graphical tool to help analyze your application's message passing behavior and determine where to concentrate your optimization efforts. See the *Paragon™ Application Tools User's Guide* for more information on ParaGraph.

## Maintain Data Locality

Wherever possible, try to distribute the data to the nodes so that each node has all the data it needs, and does not have to get any data from other nodes. Where it is not possible to keep all related data on one node, try to keep the data as close as possible to the nodes that need it.

For example, suppose you are writing a simulation where the value of each data point in a plane is computed from the values of nearby data points. To parallelize this simulation, you would divide up the plane into segments and assign each segment to a node. However, each node must communicate with other nodes to get data for points that are just past the edge of its data segment. Since the Paragon system has a mesh architecture, you would typically divide up the plane in a two-dimensional decomposition (rectangles). You would then assign the rectangles to nodes in such a way that neighboring segments are the responsibility of neighboring nodes. (Use the `nx_app_rect()` call to determine the “shape” of your application.) This will reduce message traffic and ensure that each message reaches its destination as quickly as possible.

## Overlap Computation and Communication

The Paragon OSF/1 message-passing calls are available in both *synchronous* versions (call names beginning with `c`) and *asynchronous* versions (call names beginning with `i`). Synchronous calls do not return until the message-passing operation is complete; asynchronous calls return immediately, giving you a message ID that you can use to check when the operation is complete.

Although the synchronous calls are easier to use and have slightly lower overhead, you should use the asynchronous calls whenever the results of the call are not needed immediately. Using asynchronous calls can let your application do useful computation in the time when it would otherwise just be waiting for a message to arrive. During this time, the node's message coprocessor can process the communication without interrupting the main processor.

## Avoid Message Buffering

Try to avoid message buffering whenever possible. For example, assume that you have the same process running on two nodes and that these processes must exchange information. Each process must issue a receive and a send. If a `csend()` call is executed before its corresponding `crecv()`, the message is sent and buffered in a system buffer (if there is not enough space in system buffers, the sender blocks). When the `crecv()` is executed, the message is copied from the system buffer into the application buffer. (For detailed information on message buffering, see "Understand Message-Passing Flow Control" on page 8-13.)

Your code runs more efficiently if you can avoid the system buffer and copy the message directly into the application buffer. You can do this by using an `irecv()` (the asynchronous receive) and posting the receive before the corresponding `csend()`. Remember, however, that because the nodes do not run in lock step, coding the `irecv()` before its corresponding `csend()` does not guarantee that the `irecv()` is executed before its `csend()` (even if the same program runs on every node). You can make sure that the `irecv()` is executed before the `csend()` by using zero-length messages to synchronize the nodes, as shown in the following example.

For example, consider the following C routine, *shadow*. This routine might appear in an application that needs to have the nodes exchange the rows of a matrix (a Fortran version would probably exchange columns instead).

A typical application for *shadow* might be a Gauss-Seidel iteration or any technique based on nearest neighbor interactions. The application processes a two dimensional array called `s[ ][ ]` and exchanges rows between nodes. The first index in the array represents a row, and it is passed as a pointer. Each node contains a horizontal partition of the array with *range* rows. It has a top buffer (`s[0]`) and a bottom buffer (`s[range+1]`) containing the boundary rows from other nodes.

```
void shadow(
    long topnode;
    long botnode;
    int (*s)[MAX_LATTICE]
    int range)
{
    long topid, botid, syncbotid, synctopid;

    /* Node sends upper boundary row s[1] to the bottom buffer
    s[range+1] of the node controlling the upper partition,
    (topnode).

    Node sends lower boundary row s[range] to the top buffer s[0]
    of the node controlling the lower partition, botnode */

    topid = irecv(TOP, s[0], sizeof(s[0]));
    botid = irecv(BOT, s[range+1], sizeof(s[range+1]));
```

```

/* The following code ensures that the csend()s corresponding
to the above irecv()s are not executed until all the irecv()s
have been posted. */

syncidtop = irecv(SYNCTOP, 0, 0);
syncidbot = irecv(SYNCBOT, 0, 0);
csend(SYNCTOP, 0, 0, topnode, 0);
csend(SYNCBOT, 0, 0, botnode, 0);
msgwait(syncidtop);
msgwait(syncidbot);

/* End of synchronization code. */

csend(BOT, s[1], sizeof(s[1]), topnode, 0);
csend(TOP, s[range], sizeof(s[range]), botnode, 0);

msgwait(topid);
msgwait(botid);
}

```

Note that when the data sends are performed the asynchronous receives have already been executed. This is ensured by the zero-length synchronization messages. The data then goes directly into the application buffer (unless it is paged out, as discussed later in this section).

Another way of achieving synchronization is to issue a **gsync()** after the **irecv()**'s. However, **gsync()** can be expensive—it synchronizes *all* the nodes, when all that's really necessary is to synchronize senders and receivers. A good rule of thumb is to synchronize only what is really necessary.

## Align Application Buffers

Try to ensure that send and receive buffers are properly aligned and sized whenever possible. Although the message-passing system calls will work with any size or alignment of buffers, the hardware works best with well-aligned buffers. The software may have to copy messages that are in misaligned buffers to new, aligned buffers, which decreases performance. There are several degrees of alignment. All other things being equal:

1. The best performance can be achieved by aligning the send or receive buffer on a 4K-byte boundary (this means that the buffer's address is an even multiple of 4K). This corresponds with the i860 microprocessor's 4K-byte physical memory page.
2. Good performance (only slightly worse) occurs if the buffer is aligned on a 32-byte boundary, which corresponds with the microprocessor's cache line, but crosses a 4K-byte memory page boundary.
3. The next-best performance (not nearly as good) occurs if the buffer is aligned on an 8-byte boundary, which corresponds with the microprocessor's FIFO size.



4. Some performance improvement can be seen if the buffer is aligned on a 4-byte boundary.
5. The worst performance comes when the buffer is not even aligned on a 4-byte boundary (that is, its address is not a multiple of 4).

To be sure the buffer is well-aligned, you should use `malloc()` to allocate it rather than allocating it statically. (This call is available in both Fortran and C.) Buffers allocated with the standard `malloc()` or its derivatives will always be aligned on a 32-byte boundary.

You can arrange for the pointer to a message buffer to be on a 4K-byte boundary by using pointer arithmetic. This technique can be used even if the buffer is statically allocated. You do this by declaring your buffer to be 4095 (4K-1) bytes longer than needed, adding 4095 to the buffer pointer, and then ANDing the pointer with the NOT of 4095. For example, assume that you need a 50K-byte buffer called `buf`. You can cause the buffer pointer `bufp` to be on a 4K-byte boundary by declaring `buf` to be 54K-1 bytes and doing the following pointer arithmetic:

```
char *bufp, buf[55295];
    .
    .
    .
/* bufp points to the nearest 4K-byte boundary in the buffer */
bufp = (char *) (((int)buf + 4095) & ~4095);
```

Finally, when you set up C structures that you intend to use as messages, try to minimize *padding* (open areas within the structure inserted by the compiler to make the following structure element properly aligned). To do this, you should be aware of the sizes and alignments of the different data types; in general, you can minimize padding by placing the larger data types first. Reducing padding reduces the number of empty bytes sent with the message.

## Understand Message-Passing Flow Control

Whenever you send a message, the sending process and the receiving process use *message-passing flow control* to make sure that the message is safely stored in memory when it arrives at the destination node. This flow control guarantees that messages flow from their source to destination without blocking on the mesh. If you understand flow control, you can select the message sizes, message-passing configuration parameters, and message synchronization techniques that give the best possible message-passing performance for your application.

Note, though, that the most important thing you can do to improve message-passing flow control is to avoid message buffering (as discussed under “Avoid Message Buffering” on page 8-11). You should only read this section if you cannot avoid message buffering and need to improve its efficiency, or if you really want to understand all the nitty-gritty details of message-passing flow control. If neither of these applies to you, skip to “Recommendations” on page 8-21.

## Overview of Message-Passing Flow Control

Here's an overview of what happens when you send a message. (This is a simplified view; the low-level details of message-passing flow control are proprietary and subject to change.)

1. The sending process checks to see if the memory page containing the message to be sent is currently in physical memory. If not, it is paged in. (See "Avoid Virtual Memory Paging" on page 8-3 for information on paging.)
2. The sender checks to see how much memory it thinks is available in system message buffers on the receiver and sends the appropriate number of bytes (see "System Message Buffers" on page 8-16 for details). If the whole message has been sent at this point, the send is complete; otherwise, it waits for a request from the receiver for the next part of the message. This waiting may or may not block the sending process, depending on whether the sending call was synchronous or asynchronous.
3. When the message (or the first part of the message) arrives at the receiving node, the node's operating system checks to see if there is a receive posted for a message of that type by the specified receiving process. ("Posted" means that the process has an outstanding message-receiving call that has not yet been fulfilled.)
  - A. If there is a receive posted and the application buffer (the buffer specified in the receive call) is currently paged in, the message is stored directly into the application buffer.
  - B. If there is a receive posted and the specified application buffer is not paged in, the message is stored in a system buffer. Then the application buffer is paged in and the message is copied into it.
  - C. If there is no receive posted, the message is stored in a system buffer and the receiver waits until a receive is posted. When the receive is posted, the specified application buffer is paged in (if necessary) and the message is copied into it.
4. If the whole message has been received at this point, the receive is complete; otherwise, it sends a request to the sender for the next part of the message and waits. (This waiting may or may not block the receiving process, depending on whether the receiving call was synchronous or asynchronous.) The request also includes the current free space in system message buffers on the receiver, which is used to calculate how big the next part should be. Go back to step 1 and continue until the message has been completely sent and completely received.

**Special case:** if the sender thinks the message is too large to send all at once, and there is no receive posted, but there is actually enough space in system buffers on the receiver to accommodate the entire message, the receiver stores the first part of the message in a system buffer and immediately sends a special request to the sender saying "send the whole rest of the message." In this case, the entire message can be sent before the receive is posted. Otherwise, only the first part of the message is sent and the rest of the message waits on the sender until the receive is posted.

## Process Locking

The message-passing flow control procedures check to make sure that application buffers are paged into physical memory, and copy information from one buffer to another if they are not. You can avoid these steps by using the **-plk** switch on the application command line. This switch locks parts of each process into physical memory, like the OSF/1 system call **plock()** (see the *OSF/1 Programmer's Reference* for information on **plock()**). This locking is also referred to as *wiring*.

The **-plk** switch locks the following parts of your application into physical memory:

- The entire *data segment* (the part of memory that contains global variables) is locked. This occurs when the program is loaded.
- If you use an application buffer that is located on the *stack* (the part of memory that contains local variables) or on the *heap* (the part of memory that is allocated by **malloc()** or **ALLOCATE**), the area from the beginning of the stack/heap to the end of the buffer is locked. This occurs the first time you use the buffer in a message-sending or -receiving call.

All areas of memory not mentioned in this list, including the code segment (the part of memory that contains executable instructions), are not locked and are still subject to paging. Note that locking is done a page at a time: to lock a single byte, the system must lock the entire 8K-byte virtual memory page containing that byte.

The **-plk** switch greatly reduces the effect of virtual memory on your application and improves message-passing latency. However, it has the following consequences:

- Your application must fit in physical memory. If it does not, any operation that results in the allocation or locking of more memory may fail unexpectedly, possibly terminating the application.

Ideally, the operating system and the application's code and data should all fit into the node's physical memory at once. However, the code segment is subject to paging even when **-plk** is in effect, so the application may still work if there is enough physical memory left over after subtracting the size of the operating system and the total amount of locked data. The definition of "enough" depends on the application's pattern of access to its code in memory and how much of the code needs to be present in memory at once. (See the *Paragon™ System Software Release Notes for the Paragon™ XP/S System* for information on how much memory is needed by the operating system.)

- The physical memory available for other processes is reduced by the size of your application's locked data for the life of your application.

When **-plk** is in effect, none of the locked data will be removed from physical memory until the application terminates. Even if your application is not actually executing (for example, because it is "rolled out" by gang scheduling), it still retains control of this memory, and the application that is currently executing cannot use the memory that is locked by your application. This can cause the other application to run very slowly or "thrash."

To prevent this “thrashing,” your system administrator can configure the system so that **-plk** cannot be used in gang-scheduled or standard-scheduled partitions. If **-plk** is allowed at your site, it should be used with extreme care because of its impact on other users.

**-plk** also conditions message-passing flow control to run more efficiently by assuming that all message buffers are locked into memory. For example, suppose a message is too large to send all at once. With **-plk**, after the first part of the message arrives, the receiver can request the entire rest of the message—no matter how big it is—as soon as the receive is posted. Since the application buffer cannot be paged out, the receiver can be sure it will be there to receive the rest of the message when it arrives. Without **-plk**, the application buffer could be paged out while the second part is on its way, so the second and subsequent parts of the message must be smaller than the available system message buffer space. This means that many more exchanges might be required before the message is completely received.

## Packetization

Messages from one node to another may be broken into smaller messages, called *packets*, before they are placed on the mesh. Using packets results in a slight additional overhead on large messages, but it gives much better overall message bandwidth because it allows several large messages to be interleaved on the same wire at the same time.

The maximum size of each packet is referred to as *packet\_size*, and is 1792 bytes by default (this size does not include the header appended to each packet). If a message is larger than *packet\_size*, it is sent in several pieces, each at most *packet\_size* bytes long. You can change the packet size with the **-pkt** switch on the application command line.

## System Message Buffers

In each node process, an area of memory is set aside for *system message buffers*. These buffers are used to store messages that arrive at the node before the receiving process is ready to receive them. For example, if a sending process calls **csend()** before the receiving process has called the corresponding **crecv()**, the message goes into a system buffer in the receiving process. Then, when the receiving process does call **crecv()**, the message is copied from the system buffer to the buffer specified in the **crecv()** call, which is referred to as the *application message buffer*.

The size and behavior of the system message buffers are controlled by several parameters that you can set on the application command line. The following list describes these parameters and their effects.

- The total amount of memory allocated to system message buffers in each process is referred to as *message\_buffer*. The *message\_buffer* is always *wired* into physical memory, which means that it can never be paged out (see “Avoid Virtual Memory Paging” on page 8-3 for information on paging). This is necessary to ensure that all messages that arrive at the node can be stored somewhere, even if the rest of the application is paged out. The default value of *message\_buffer* is 1152K bytes. You can change this size with the **-mbf** switch on the application command line

- The *message\_buffer* is divided into an area for messages from any process, and a series of areas dedicated to messages from particular processes. The number of dedicated areas is referred to as *correspondents*, and the size of each dedicated area is referred to as *memory\_each*. When a message is received, it is stored in the open area if there is room; the *memory\_each* areas are used only when the open area is full. The default value for *correspondents* is *numnodes()*; you can change it with the **-noc** switch. The default value of *memory\_each* is determined by the current values of *correspondents*, *message\_buffer*, and *packet\_size*; you can change it with the **-mea** switch.
- Each node process maintains a value called the *send\_avail* for each other process in the application. The *send\_avail* is the maximum amount of memory that the process can depend on to be available in its *memory\_each* segment in that other process. (The *send\_avail* may be smaller than the actual amount of memory available, but is never larger.) When a process sends a message to another process, it decreases the *send\_avail* for that process by the message size; when the message has been “consumed” (completely placed in the application buffer on the receiving process), the receiver tells the sender and the sender increases its *send\_avail* for the receiving process accordingly. The initial value of *send\_avail* is *memory\_each*; the *send\_avail* value is maintained dynamically by each process and cannot be set on the command line.
- When a process has a large message to send, it uses its *send\_avail* value for the receiving process to determine how much of the message to send at first. Two parameters called *send\_threshold* and *send\_count* control this behavior:
  1. If the *send\_avail* value is equal to or greater than the size of the message, the sender sends the whole message at once.
  2. Otherwise, if the *send\_avail* value is equal to or greater than the *send\_threshold*, the sender sends the first *send\_count* bytes of the message and waits for an acknowledgment from the receiver that they have been consumed before proceeding.
  3. Otherwise, if the *send\_avail* value is equal to or greater than the *packet\_size*, the sender sends the first packet of the message and waits for an acknowledgment from the receiver that it has been consumed before proceeding.
  4. If the *send\_avail* value is less than *packet\_size*, the send blocks until the receiver tells it that some messages have been consumed and *send\_avail* can be increased. (See the discussion of *give\_threshold* later in this section for more information on how this occurs.) This blocking may or may not block the sending process, depending on whether it used a synchronous send (such as **csend()**) or an asynchronous send (such as **isend()**).

Note that deadlock can occur when *send\_avail* is less than *packet\_size*. For example, suppose that node A and node B's system message buffers are both full. Under normal circumstances, eventually a receive would be posted and the buffered messages would be

consumed. But if the two nodes try to exchange messages with synchronous calls, they deadlock: A blocks waiting for more space to become available on B, and B blocks waiting for more space to become available on A.

The default value for *send\_threshold* and *send\_count* is half of *memory\_each*; you can change these two parameters with the **-sth** and **-sct** switches respectively.

- When a message is consumed, the receiver normally informs the sender that the space occupied by the message is available for new messages by “piggy-backing” information on other messages going to the sender. However, if there are no such messages, the sender can get out of date and stop sending messages because it thinks there is no free memory left for it on the receiver. In this case, a parameter called the *give\_threshold* comes into play. If a receiver knows that a sender thinks it has less than *give\_threshold* bytes of memory free, but there is really more memory available, it sends a special message to the sender telling it how much memory is really available. The default value for *give\_threshold* is *packet\_size*; you can change it with the **-gth** switch.

## Message-Passing Configuration Switches

The switches that control the message-passing configuration parameters discussed earlier in this section are referred to as the *mp\_switches*. Although the default values of these parameters have been chosen to give good results for “typical” applications, you may be able to improve your application’s message-passing performance by using different values.

You use the *mp\_switches* on the command line of a parallel application. These switches override the default values of the specified parameters for that run of the application; they do not have any effect on other runs or other applications.

If the application was linked with the **-nx** switch, the *mp\_switches* are automatically interpreted and removed from the command line before the application starts up. An application linked with **-lnx** controls its own execution with system calls, as discussed under “Managing Applications” on page 4-2. Such an application may or may not obey the *mp\_switches*, depending on how it was programmed.

The values used with the *mp\_switches* (except **-plk** and **-noc**) are integer numbers of bytes. The default, maximum, and minimum values for these switches are described under “Default, Maximum, and Minimum Values” on page 8-20. The value you specify may be rounded up or down to ensure correct operation, as described under “Dependencies and Rounding” on page 8-21.

## Summary of the Message-Passing Configuration Switches

The list of available *mp\_switches* is as follows:

<b>-plk</b>	Locks parts of the application into memory (see “Process Locking” on page 8-15 for more information).
<b>-pkt <i>packet_size</i></b>	Sets the size of each packet.
<b>-mbf <i>message_buffer</i></b>	Sets the total amount of memory allocated to message buffers in each process.
<b>-noc <i>correspondents</i></b>	Sets the total number of other processes from which each process expects to receive messages.
<b>-mex <i>memory_export</i></b>	Used in setting the maximum value for <i>memory_each</i> ; otherwise ignored.
<b>-mea <i>memory_each</i></b>	Sets the amount of memory allocated to buffering messages from each correspondent.
<b>-sth <i>send_threshold</i></b>	Sets the threshold for sending multiple packets.
<b>-sct <i>send_count</i></b>	Sets the number of bytes to send right away when the available memory is above <i>send_threshold</i> .
<b>-gth <i>give_threshold</i></b>	Sets the threshold for “give me more messages” message.

See “Process Locking” on page 8-15, “Packetization” on page 8-16, and “System Message Buffers” on page 8-16 for more detailed information.

### Default, Maximum, and Minimum Values

The default, maximum, and minimum values for the *mp\_switches* are shown in Table 8-1.

**Table 8-1. Message-Passing Configuration Switches**

Switch	Parameter	Default	Maximum	Minimum
<b>-plk</b>	none	unlocked	n/a	n/a
<b>-pkt</b>	<i>packet_size</i>	1792 or $((\text{memory\_each} / 2) - \text{sizeof}(\text{xmsg\_t})^1)$ , whichever is less	1792	$\text{sizeof}(\text{xmsg\_t})$
<b>-mbf</b>	<i>message_buffer</i>	1MB + 128KB	32MB + $(10 * \text{full\_packet\_size}^2)$	$(8 * \text{sizeof}(\text{xmsg\_t})) * (\text{correspondents} + 2) + (20 * \text{sizeof}(\text{xmsg\_t}))$
<b>-mex</b>	<i>memory_export</i>	<i>message_buffer</i> - 128KB	<i>message_buffer</i> - 128KB	$2 * (\text{correspondents} + 2) * (2 * \text{full\_packet\_size})$
<b>-noc</b>	<i>correspondents</i>	$\text{numnodes}()$	none	none
<b>-mea</b>	<i>memory_each</i>	$(10 * \text{full\_packet\_size})$ or maximum <i>memory_each</i> , whichever is less	1MB - 31 or $(\text{memory\_export} / 2) / (\text{correspondents} + 2)$ , whichever is less	$2 * \text{full\_packet\_size}$
<b>-sth</b>	<i>send_threshold</i>	$\text{memory\_each} / 2$	$\text{memory\_each} - 1$	none
<b>-sct</b>	<i>send_count</i>	$\text{memory\_each} / 2$	<i>memory_each</i>	<i>packet_size</i>
<b>-gth</b>	<i>give_threshold</i>	<i>packet_size</i>	$\text{memory\_each} / 2$	<i>packet_size</i>

1. *xmsg\_t* is a type defined in `<mcmsg/mcmsg_xmsg.h>` that defines the message header sent along with each packet. The size of this type is currently 64 bytes.

2.  $\text{full\_packet\_size} = \text{packet\_size} + \text{sizeof}(\text{xmsg\_t})$ .



## Dependencies and Rounding

As you can see from Table 8-1, the values for some of the *mp\_switches* depend on the current values of other switches in a circular manner (for example, the default for *packet\_size* depends on the value of *memory\_each*, while the default for *memory\_each* depends on the value of *packet\_size*). These dependencies are resolved using the following procedure:

1. Set *packet\_size*: If **-pkt** is specified, round the specified value up to a multiple of `sizeof(xmsg_t)`. Otherwise, use the default value.
2. Set *message\_buffer*: If **-mbf** is specified, round the specified value up to a multiple of *full\_packet\_size*. Otherwise, use the default value.
3. Set *memory\_export*: If **-mex** is specified, use the specified value. Otherwise, use the default value.
4. Set *memory\_each*: If **-mea** is specified, round the specified value down to a multiple of `sizeof(xmsg_t)`. Otherwise, round the default value down to a multiple of `sizeof(xmsg_t)`.
5. Check that *memory\_each* will hold at least two packets: If  $(memory\_each/2) - sizeof(xmsg\_t)$  is less than *packet\_size*, reset *packet\_size* to the value  $((memory\_each/2) - sizeof(xmsg\_t))$ , round the resulting value down to a multiple of `sizeof(xmsg_t)`, then return to step 2. Otherwise, continue to step 6.
6. Set *send\_threshold*: If **-sth** is specified, round the specified value down to a multiple of *packet\_size*. Otherwise, round the default value down to a multiple of *packet\_size*.
7. Set *send\_count*: If **-sct** is specified, round the specified value down to a multiple of *packet\_size*. Otherwise, round the default value down to a multiple of *packet\_size*.
8. Set *give\_threshold*: If **-gth** is specified, round the specified value down to a multiple of *packet\_size*. Otherwise, round the default value down to a multiple of *packet\_size*.

## Recommendations

Because of the way message-passing flow control works, you should try to do all the following to achieve the best possible message-passing performance:

- Avoid paging, by keeping the application's memory requirements within available physical memory. Once you have done this, use the **-plk** switch if this is allowed at your site.
- Avoid blocking, by using asynchronous calls.
- Avoid system message buffering, by posting receives before the message is sent. (See "Avoid Message Buffering" on page 8-11 for tips on how to do this.)

It is important to make all three of these changes if possible. For example, even if you always post receives before the corresponding send occurs, system message buffering will still be necessary if the application buffer is paged out (or has never been paged in) when the message arrives.

If you cannot avoid system message buffering, you may be able to improve message-passing performance by increasing the *message\_buffer* parameter (**-mbf**). This parameter determines the total amount of memory allocated to message buffers in each process; the other parameters determine how this memory is divided up. When you change the value of *message\_buffer*, the defaults for the other parameters are automatically scaled to match the current *message\_buffer* size. Increasing the *message\_buffer* can increase the efficiency of message passing, but it also increases the memory usage of your application, which may cause paging and slow the application down. Once you have determined the optimal *message\_buffer* size for your application, you can change the other parameters to fine-tune the usage of memory within the *message\_buffer* and optimize message-passing performance.

The performance of some applications that use system message buffering can also be improved by reducing the *correspondents* parameter (**-noc**). This is particularly likely to help if your application slows down or hangs when you run it on more nodes. The **-noc** switch sets the “number of correspondents” for each process, which is the number of other processes from which the process receives messages. This number is used to determine how the memory allocated to buffering messages is divided up; more correspondents means that less memory is available for buffering messages from each correspondent. If you don't use **-noc**, the default for *correspondents* is **numnodes()**; that is, it is assumed that each process may receive messages from one process on each node. If you know that each process does not receive messages from every other node, using **-noc** to decrease the value of *correspondents* increases the *memory\_each* buffer size, which can result in more efficient message passing (especially if the number of nodes is large). However, if the total number of other processes from which any process receives messages during the life of the application exceeds the value of *correspondents*, the application may run more slowly.

Note that certain global operations, such as sending to node -1 (which broadcasts a message to all nodes in the application) or calling **gdsum()**, can send messages to intermediate nodes. For example, sending to node -1 does not simply send one message to every other node; instead, it sends a message to several other nodes, which each send messages to several other nodes, and so on in a “message tree.” This method is more efficient, but it means that if you use any global operations, the actual number of correspondents will be greater than the number of nodes from which each node receives explicit messages (by approximately the log of the number of nodes in the application).

## I/O Performance

If your application performs I/O to files, you can use the following techniques to improve its I/O performance. Note: the term *request size* refers to the number of bytes specified in a single **read** or **write** operation. Techniques discussed in this section include:

- Use PFS file systems.
- Use **gopen()** instead of **open()**.
- Use parallel I/O calls.
- Use asynchronous calls.
- Use the appropriate I/O mode.
- Align I/O buffers with virtual memory pages.
- Read or write whole file system blocks.
- Make good use of file striping.

See Chapter 5 for more information on these techniques.

### Use PFS File Systems

Always store large data files in file systems of type *PFS (Parallel File System)*. These file systems are optimized for large I/O requests (request sizes of 64K bytes or more) and simultaneous access by multiple nodes, and files in them can be larger than 2G bytes in size.

### Use **gopen()** Instead of **open()**

If all nodes in an application open the same file, you should always use **gopen()** rather than **open()**. If all nodes call **open()**, each node sends an “open file” message to the same I/O node at the same time, which can swamp the I/O node with messages. But when all nodes call **gopen()**, only one node communicates with the I/O node; the open file descriptor is then broadcast to the other nodes in the application through efficient global communication techniques. If you *must* use **open()**, try to keep all the nodes from calling it at the same time (do *not* precede the **open()** with a **gsync()**).

## Use Parallel I/O Calls

If you program in Fortran, you should always use Paragon OSF/1 *parallel I/O calls*, such as **cread()**, to access your files. These calls give much better performance than the standard Fortran file I/O statements, such as **READ**.

If you program in C, you will not see any I/O performance increase from using parallel I/O calls, such as **cread()**, rather than standard UNIX I/O calls, such as **read()** (although **cread()** gives better performance than **fread()**, which is the C equivalent of Fortran's **READ**). However, you may be able to improve *computational* performance by using asynchronous I/O calls.

## Use Asynchronous Calls

The parallel I/O calls are available in both *synchronous* versions (call names beginning with **c**) and *asynchronous* versions (call names beginning with **i**). Synchronous calls do not return until the I/O operation is complete; asynchronous calls return immediately, giving you an I/O ID that you can use to check when the operation is complete.

Although the synchronous calls are easier to use and have slightly lower overhead, you should use the asynchronous calls whenever the results of the call are not needed immediately. Using asynchronous calls can let your application do useful computation in the time when it would otherwise just be waiting for a large I/O operation to complete.

## Use the Appropriate I/O Mode

When you use Paragon OSF/1 parallel I/O calls, you can choose from five I/O modes (**M\_UNIX**, **M\_LOG**, **M\_SYNC**, **M\_RECORD**, and **M\_GLOBAL**), each of which is optimized for a particular pattern of file I/O. Be sure to use the correct I/O mode for your application's usage. In particular:

- Don't use **M\_UNIX**, the default I/O mode, unless your application depends on its semantics.
- If all nodes read the same data from the same file at the same time, use **M\_GLOBAL**.
- If all nodes read or write the same file, but each node is accessing a different part of the file, use **M\_RECORD** if at all possible. This mode provides much higher multi-node performance than the other modes, all of which force reads and writes from different nodes to the same file to be performed in strict sequential order (this is required to preserve standard UNIX I/O semantics, but slows the application down).
- If **M\_RECORD** cannot be used because the I/O request size is not constant across all compute nodes, use **M\_SYNC** instead.

## Align I/O Buffers with Virtual Memory Pages

Try to ensure that memory buffers used in I/O calls are aligned on an 8K-byte boundary whenever possible, to align with the operating system's virtual memory page size. This alignment is particularly important in scatter/gather operations with large request sizes to multiple I/O nodes. If you do not specify properly-aligned buffers, the software must copy the data to new, aligned buffers, which decreases performance.

To be sure the buffer is well-aligned, you should use `malloc()` to allocate it rather than allocating it statically. (This call is available in both Fortran and C.) Buffers allocated with the standard `malloc()` or its derivatives will always be aligned on a 32-byte boundary. See "Align Application Buffers" on page 8-12 for more information on aligning buffers.

## Read or Write Whole File System Blocks

Disk space is allocated and managed in units called *file system blocks*. The size of each block in a file system is determined when the file system is created. For best performance, PFS file systems should have a file system block size of 64K bytes.

The file system block size is important because files always begin at a block boundary and data is most efficiently transferred to and from the physical disk in integer numbers of blocks. Furthermore, if a block is modified (but not entirely overwritten) by a write operation, the block may have to be read, modified in memory, and then written back.

Because of this, you will get the best I/O performance if each read or write request begins on a block boundary (a multiple of the block size from the beginning of the file) and the request size is a multiple of the file system block size.

To determine the block size of a file system, you can use the `statfs()` or `fstatfs()` call (see the *OSF/1 Programmer's Reference* for information on these calls), or ask your system administrator.

## Make Good Use of File Striping

Files in PFS file systems are distributed, or *striped*, across several directories called *stripe directories*. The number of stripe directories in a PFS file system is called the *stripe factor*, and the amount of data from each file that is stored in each directory is called the *stripe unit*. The product of the stripe factor and the stripe unit is called the *full stripe size*. A PFS file system's stripe factor and stripe unit are set by the system administrator when the PFS file system is mounted.

Each stripe directory is typically on a separate disk, and each disk is typically controlled by a separate I/O node; you get the best I/O performance when you keep all the I/O nodes busy at once. You can use file striping to help you do this, with two different methods:

1. Use a request size equal to an integer multiple of the full stripe size, and make the starting address of each request the beginning of a full stripe. With this method, each I/O request goes to all the I/O nodes at once. This method can be used on any number of nodes.
2. Use a request size equal to the stripe unit size, make the starting address of each request the beginning of a stripe unit, and choose the starting address of each node's requests so that the nodes' requests are evenly distributed among the I/O nodes. With this method, each I/O request goes to just one I/O node, but the application's I/O requests are distributed among the I/O nodes. This method should be used only if the number of compute nodes is greater than or equal to the number of I/O nodes, preferably an integer multiple of the number of I/O nodes.

These two methods are illustrated in Figure 8-1. Note that method 1 uses fewer, larger requests and method 2 uses more, smaller requests. Method 1 is generally more efficient, but method 2 may give better performance for some situations (depending on the number of compute nodes, the number of I/O nodes, the amount of memory on each I/O node, and the size and frequency of requests). If possible, you should try both methods and use whichever is more efficient. The example program in */usr/share/examples/c/strip*e demonstrates the two methods, and you can use it to help you determine which method is best for your application. (Note that you will see more consistent results from one run to the next if the data size is large—8M bytes per node or more.)

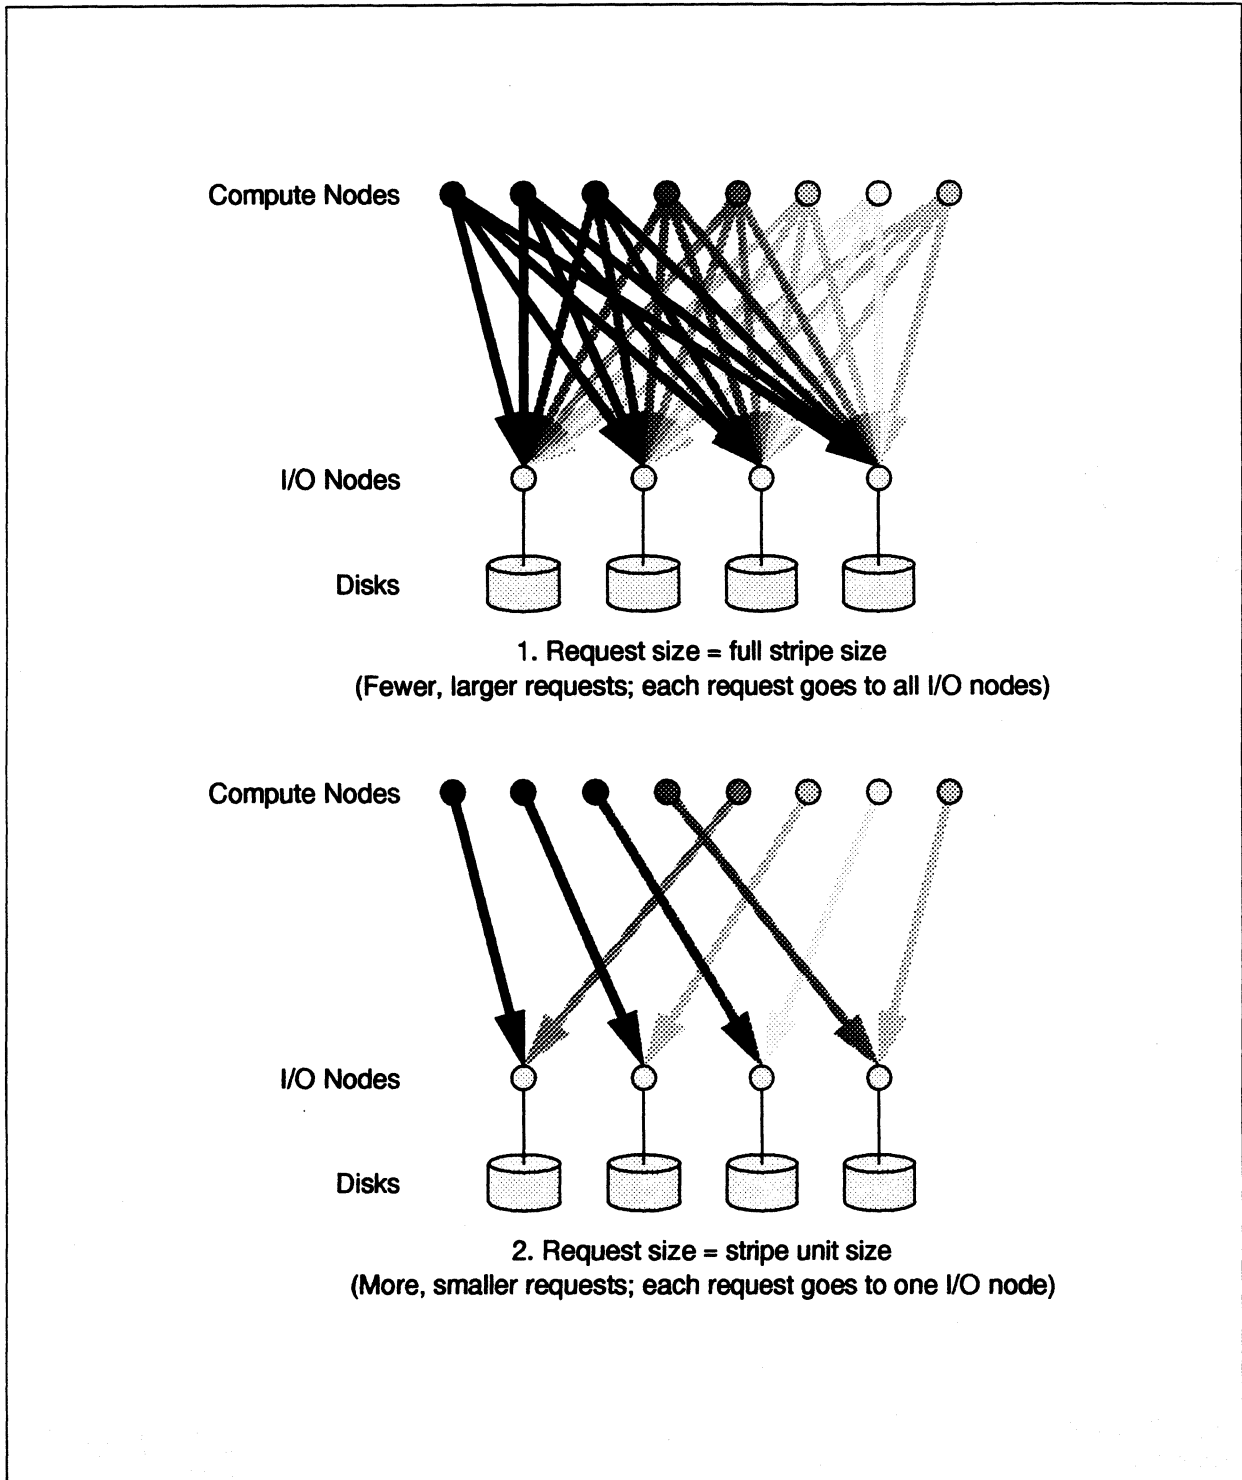
You should always use the I/O mode **M\_RECORD** when using these methods. **M\_RECORD** is the most efficient I/O mode for this type of I/O, and automatically enforces the distribution of data among the I/O nodes. If you use **M\_RECORD**, no file pointer calculation or seeking is required. For example:

```
fd = gopen(file, O_WRONLY|O_CREAT|O_TRUNC, M_RECORD, 0666);

while(data < end) {
    cwrite(fd, data, request_size);
    data += request_size;
}
```

Using this code, if *request\_size* is equal to the full stripe size, each compute node automatically accesses all I/O nodes on each write (method 1). Alternatively, if *request\_size* is equal to the stripe unit, each compute node automatically accesses exactly one I/O node on each write (method 2).

To determine the stripe factor and stripe unit of a PFS file system, you can use the **showfs** command (described under “Displaying File System Attributes” on page 5-5) or the **statpfs()** or **fstatpfs()** call (available only in C; described under “Getting Information About PFS File Systems” on page 5-39). The example program in */usr/share/examples/c/strip*e shows you how you can do this with **fstatpfs()**.



**Figure 8-1. Two Methods of Improving I/O Performance with M\_RECORD**





# Summary of Commands and System Calls

A

This appendix summarizes the commands and system calls of Paragon™ OSF/1. The complete syntax of each command and call is provided, along with a brief description of each. The C and Fortran versions of the calls are discussed in separate sections.

This appendix discusses only the commands and calls that are specific to Paragon OSF/1. For information on the standard commands and calls of OSF/1, see the *OSF/1 Command Reference* and *OSF/1 Programmer's Reference*.

## Command Summary

This section summarizes the commands discussed in Chapter 2 and Chapter 5. See the *Paragon™ Commands Reference Manual* for more information on these commands.

## Compiling and Linking Applications

Table A-1. Commands for Compiling and Linking Applications

Command Synopsis	Description
<b>cc -nx</b> [ switches ] sourcefile...	Compile a Paragon OSF/1 application written in C on an Intel supercomputer.
<b>f77 -nx</b> [ switches ] sourcefile...	Compile a Paragon OSF/1 application written in Fortran on an Intel supercomputer.
<b>icc -nx</b> [ switches ] sourcefile...	Compile a Paragon OSF/1 application written in C on an Intel supercomputer or cross-development workstation.
<b>if77 -nx</b> [ switches ] sourcefile...	Compile a Paragon OSF/1 application written in Fortran on an Intel supercomputer or cross-development workstation.

## Running Applications

Table A-2. Commands for Running Applications

Command Synopsis	Description
<pre>application [ -sz size   -sz hXw   -nd hXw:n ] [ -pri priority ] [ -pt ptype ] [ -on nodespec ] [ -pn partition ] [ -pkt packet_size ] [ -noc correspondents ] [ -mbf memory_buffer ] [ -mex memory_export ] [ -mea memory_each ] [ -sth send_threshold ] [ -sct send_count ] [ -gth give_threshold ] [ -plk ] [ application_args ] [ \; file [ -pt ptype ] [ -on nodespec ] [ application_args ] ] ...</pre>	Execute a Paragon OSF/1 application.

## Managing Partitions

Table A-3. Commands for Managing Partitions

Command Synopsis	Description
<pre>mkpart [ -sz size   -sz hXw   -nd nodespec ] [ -ss   [ [ -sps   -rq time ] [ -epl priority ] ] ] [ -mod mode ] partition</pre>	Create a partition.
<pre>rmpart [ -f ] [ -r ] partition</pre>	Remove a partition.
<pre>showpart [ -f ] [ partition ]</pre>	Show the characteristics of a partition.
<pre>lspart [ -r ] [ partition ]</pre>	List the subpartitions of a partition.
<pre>pspart [ -r ] [ partition ]</pre>	List the applications in a partition.
<pre>chpart [ -epl priority ] [ -g group ] [ -mod mode ] [ -nm name ] [ -o owner[ . group ] ] [ -rq time   -sps ] partition</pre>	Change certain partition characteristics.

## Parallel File System Commands

Table A-4. Parallel File System Commands

Command Synopsis	Description
<b>showfs</b> [ <b>-k</b> ] [ <b>-t type</b> ] [ <i>filesystem</i> \ <i>directory</i> ]	Display file system attributes.
<b>lsize</b> [ <b>-a</b> ] <i>size file ...</i>	Change the size of a file or files.

## Miscellaneous Commands

Note: the commands shown in Table A-5 are not documented in this manual.

Table A-5. Miscellaneous Commands

Command Synopsis	Description
<b>fsplit</b> [ <i>filename</i> ]	Split one file containing several Fortran program units into several files containing one program unit each. (See the <i>Paragon™ Commands Reference Manual</i> for more information.)
<b>pmake</b> [ <b>-bcdeFikmnNpqrsStuUvw</b> ] [ <b>-C dir</b> ] [ <b>-f file</b> ] [ <b>-I dir</b> ] [ <b>-j [ jobs ]</b> ] [ <b>-l [ load ]</b> ] [ <b>-o file</b> ] [ <b>-P partition</b> ] [ <b>-W file</b> ] [ <i>macro_definition ...</i> ] [ <i>target ...</i> ]	Parallel make utility that maintains up-to-date versions of target files and performs shell programs in parallel. (See the <i>Paragon™ Application Tools User's Guide</i> for more information.)
<b>sat</b> [ <b>-bchxV</b> ] [ <b>-d dir</b> ] [ <b>-l log</b> ] [ <b>-m mins</b> ] [ <b>-o output</b> ] [ <b>-p partition</b> ] [ <b>-r reps</b> ] [ <i>test ...</i> ]	Run the Paragon system acceptance test. (See the <i>Paragon™ System Acceptance Test User's Guide</i> for more information.)

## C System Call Summary

This section summarizes the C versions of the system calls discussed in Chapter 3, Chapter 4, Chapter 5, and Chapter 6. See the *Paragon™ C System Calls Reference Manual* for more information on these calls.

### Process Characteristics

Table A-6. C Calls for Process Characteristics

Synopsis	Description
long <b>mynode</b> (void);	Obtain the calling process's node number.
long <b>numnodes</b> (void);	Obtain the number of nodes allocated to the current application.
long <b>myptype</b> (void);	Obtain the calling process's process type.
void <b>setptype</b> ( long <i>ptype</i> );	Set the calling process's process type (only permitted if the process type is currently <b>INVALID_PTYPE</b> ).
long <b>myhost</b> (void);	Obtain the controlling process's node number.

## Synchronous Send and Receive

Table A-7. C Calls for Synchronous Send and Receive

Synopsis	Description
<pre>void csend(   long type,   char *buf,   long count,   long node,   long ptype );</pre>	Send a message, waiting for completion.
<pre>void crecv(   long typesel,   char *buf,   long count );</pre>	Receive a message, waiting for completion.
<pre>long csendrecv(   long type,   char *sbuf,   long scount,   long node,   long ptype,   long typesel,   char *rbuf,   long rcount );</pre>	Send a message and post a receive for the reply. Wait for completion.
<pre>void gsendx(   long type,   char *buf,   long count,   long nodes[],   long nodecount );</pre>	Send a message to a list of nodes, waiting for completion.

## Asynchronous Send and Receive

Table A-8. C Calls for Asynchronous Send and Receive

Synopsis	Description
<pre>long isend(     long type,     char *buf,     long count,     long node,     long ptype );</pre>	Send a message without waiting for completion.
<pre>long irecv(     long typesel,     char *buf,     long count );</pre>	Receive a message without waiting for completion.
<pre>long isendrecv(     long type,     char *sbuf,     long scount,     long node,     long ptype,     long typesel,     char *rbuf,     long rcount );</pre>	Send a message and post a receive for the reply without waiting for completion.
<pre>long msgdone(     long mid );</pre>	Determine whether a send or receive operation has completed.
<pre>void msgwait(     long mid );</pre>	Wait for completion of a send or receive operation.
<pre>void msgignore(     long mid );</pre>	Release a message ID as soon as a send or receive operation completes.
<pre>long msgmerge(     long mid1,     long mid2 );</pre>	Merge two message IDs into a single ID that can be used to wait for completion of both operations.

## Probing for Pending Messages

Table A-9. C Calls for Probing for Pending Messages

Synopsis	Description
<code>void cprobe( long typesel );</code>	Wait for a message of a selected type to arrive.
<code>long iprobe( long typesel );</code>	Determine whether a message of a selected type is pending.

## Getting Information About Pending or Received Messages

Table A-10. C Calls for Getting Information About Pending or Received Messages

Synopsis	Description
<code>long infocount(void);</code>	Return size in bytes of a pending or received message.
<code>long infonode(void);</code>	Return node number of the node that sent a pending or received message.
<code>long infoftype(void);</code>	Return process type of the process that sent a pending or received message.
<code>long infotype(void);</code>	Return message type of a pending or received message.

## Treating a Message as an Interrupt

Table A-11. C Calls for Treating a Message as an Interrupt

Synopsis	Description
<pre>void hsend(     long type,     char *buf,     long count,     long node,     long ptype,     void (*handler) () );</pre>	Send a message and set up a handler procedure to be called when the send completes.
<pre>void hrecv(     long typesel,     char *buf,     long count,     void (*handler) () );</pre>	Receive a message and set up a handler procedure to be called when the receive completes.
<pre>void hsendrecv(     long type,     char *sbuf,     long scount,     long node,     long ptype,     long typesel,     char *rbuf,     long rcount,     void (*handler) () );</pre>	Send a message and post a receive for the reply. Set up a handler procedure to be called when the reply arrives.
<pre>long masktrap(     long state );</pre>	Enable or disable interrupts for message handlers. Required to prevent corruption of global variables.
<pre>void hsendx(     long type,     char *buf,     long count,     long node,     long ptype,     void (*xhandler) (),     long hparam );</pre>	Send a message and set up an extended handler procedure to be called with the value <i>hparam</i> when the send completes. Allows handler sharing.



## Extended Receive and Probe

Table A-12. C Calls for Extended Receive and Probe

Synopsis	Description
<pre>void crecvx(     long typesel,     char *buf,     long count,     long nodesel,     long ptypesel,     long info[] );</pre>	<p>Receive a message of a specified type from a specified sending node and process type, together with information about the message. Wait for completion.</p>
<pre>long irecvx(     long typesel,     char *buf,     long count,     long nodesel,     long ptypesel,     long info[] );</pre>	<p>Receive a message of a specified type from a specified sending node and process type, together with information about the message. Do not wait for completion.</p>
<pre>void hrecvx(     long typesel,     char *buf,     long count,     long nodesel,     long ptypesel,     void (*xhandler) (),     long hparam );</pre>	<p>Receive a message of a specified type from a specified sending node and process type. Set up an extended handler procedure to be called with information about the message and the value <i>hparam</i> when the receive completes.</p>
<pre>void cprobex(     long typesel,     long nodesel,     long ptypesel,     long info[] );</pre>	<p>Wait for a message of a specified type from a specified sending node and process type. Return information about the message.</p>
<pre>long iprobex(     long typesel,     long nodesel,     long ptypesel,     long info[] );</pre>	<p>Determine whether a message of a specified type from a specified sending node and process type is pending. If it is, return information about the message.</p>

## Global Operations

Table A-13. C Calls for Global Operations (1 of 3)

Synopsis	Description
<pre>void gcol(   char x[],   long xlen,   char y[],   long ylen,   long *ncnt );</pre>	Concatenation.
<pre>void gcolx(   char x[],   long xlen[],   char y[] );</pre>	Concatenation for contributions of known length.
<pre>void gdhigh(   double x[],   long n,   double work[] );</pre>	Vector double precision MAX.
<pre>void gdlow(   double x[],   long n,   double work[] );</pre>	Vector double precision MIN.
<pre>void gdprod(   double x[],   long n,   double work[] );</pre>	Vector double precision MULTIPLY.
<pre>void gdsun(   double x[],   long n,   double work[] );</pre>	Vector double precision SUM.
<pre>void giand(   long x[],   long n,   long work[] );</pre>	Vector integer bitwise AND.
<pre>void gihigh(   long x[],   long n,   long work[] );</pre>	Vector integer MAX.

Table A-13. C Calls for Global Operations (2 of 3)

Synopsis	Description
void <b>gilow</b> ( long <i>x</i> [], long <i>n</i> , long <i>work</i> [] );	Vector integer MIN.
void <b>gior</b> ( long <i>x</i> [], long <i>n</i> , long <i>work</i> [] );	Vector integer bitwise OR.
void <b>giprod</b> ( long <i>x</i> [], long <i>n</i> , long <i>work</i> [] );	Vector integer MULTIPLY.
void <b>gisum</b> ( long <i>x</i> [], long <i>n</i> , long <i>work</i> [] );	Vector integer SUM.
void <b>gland</b> ( long <i>x</i> [], long <i>n</i> , long <i>work</i> [] );	Vector logical AND.
void <b>glor</b> ( long <i>x</i> [], long <i>n</i> , long <i>work</i> [] );	Vector logical inclusive OR.
void <b>gopf</b> ( char <i>x</i> [], long <i>xlen</i> , char <i>work</i> [], long ( <i>*function</i> )() );	Arbitrary commutative function.
void <b>gshigh</b> ( float <i>x</i> [], long <i>n</i> , float <i>work</i> [] );	Vector real MAX.
void <b>gslow</b> ( float <i>x</i> [], long <i>n</i> , float <i>work</i> [] );	Vector real MIN.

Table A-13. C Calls for Global Operations (3 of 3)

Synopsis	Description
void <b>gsprod</b> ( float <i>x</i> [], long <i>n</i> , float <i>work</i> []);	Vector real MULTIPLY.
void <b>gssum</b> ( float <i>x</i> [], long <i>n</i> , float <i>work</i> []);	Vector real SUM.
void <b>gsync</b> (void);	Global synchronization.

## Controlling Application Execution

Table A-14. C Calls for Controlling Application Execution (1 of 2)

Synopsis	Description
long <b>nx_initve</b> ( char * <i>partition</i> , long <i>size</i> , char * <i>account</i> , long * <i>argc</i> , char * <i>argv</i> []);	Create a new application.
long <b>nx_initve_rect</b> ( char * <i>partition</i> , long <i>anchor</i> , long <i>rows</i> , long <i>cols</i> , char * <i>account</i> , long * <i>argc</i> , char * <i>argv</i> []);	Create a new application with a rectangular shape.
long <b>nx_pri</b> ( long <i>pgroup</i> , long <i>priority</i> );	Set the priority of an application.
long <b>nx_nfork</b> ( long <i>node_list</i> [], long <i>numnodes</i> , long <i>ptype</i> , long <i>pid_list</i> []);	Copy the current process onto some or all nodes of an application.

Table A-14. C Calls for Controlling Application Execution (2 of 2)

Synopsis	Description
<pre>long nx_load(     long node_list[],     long numnodes,     long ptype,     long pid_list[],     char *pathname );</pre>	Execute a stored program on some or all nodes of an application.
<pre>long nx_loadve(     long node_list[],     long numnodes,     long ptype,     long pid_list[],     char *pathname,     char *argv[],     char *envp[] );</pre>	Execute a stored program on some or all nodes of an application, with specified argument list and environment.
<pre>long nx_waitall(void);</pre>	Wait for all application processes.

## Getting Information About Applications

Table A-15. C Calls for Getting Information About Applications

Synopsis	Description
<pre>long nx_app_rect(     long *rows,     long *cols );</pre>	Obtain the height and width of the rectangle of nodes allocated to the current application.
<pre>int nx_app_nodes(     pid_t pgroup,     nx_nodes_t *node_list,     unsigned long *list_size );</pre>	List the nodes allocated to an application.
<pre>int nx_pspart(     char *partition,     nx_pspart_t **pspart_list,     unsigned long *list_size );</pre>	Obtain information about all applications and active subpartitions in a partition.

## Partition Management

Table A-16. C Calls for Partition Management (1 of 2)

Synopsis	Description
<pre>long nx_mkpart(   char *partition,   long size,   long type );</pre>	Create a partition with a particular number of nodes.
<pre>long nx_mkpart_rect(   char *partition,   long rows,   long cols,   long type );</pre>	Create a partition with a particular height and width.
<pre>long nx_mkpart_map(   char *partition,   long numnodes,   long node_list[],   long type );</pre>	Create a partition with a specific set of nodes.
<pre>long nx_rmpart(   char *partition,   long force,   long recursive );</pre>	Remove a partition.
<pre>int nx_part_attr(   char *partition,   nx_part_info_t *attributes );</pre>	Get a partition's attributes.
<pre>int nx_part_nodes(   char *partition,   nx_nodes_t *node_list,   unsigned long *list_size );</pre>	List the root node numbers for the nodes of a partition.
<pre>long nx_chpart_name(   char *partition,   char *name );</pre>	Change a partition's name.
<pre>long nx_chpart_mod(   char *partition,   long mode );</pre>	Change a partition's protection modes.
<pre>long nx_chpart_epl(   char *partition,   long priority );</pre>	Change a partition's effective priority limit.
<pre>long nx_chpart_rq(   char *partition,   long rollin_quantum );</pre>	Change a partition's rollin quantum.

Table A-16. C Calls for Partition Management (2 of 2)

Synopsis	Description
<pre>long nx_chpart_owner(   char *partition,   long owner,   long group );</pre>	Change a partition's owner and group.
<pre>long nx_chpart_sched(   char *partition,   int sched_type );</pre>	Change a partition's scheduling type.

## Finding Unusable Nodes

Table A-17. C Calls for Finding Unusable Nodes

Synopsis	Description
<pre>int nx_empty_nodes(   nx_nodes_t *node_list,   unsigned long *list_size );</pre>	List the nodes that are empty slots.
<pre>int nx_failed_nodes(   nx_nodes_t *node_list,   unsigned long *list_size );</pre>	List the nodes that failed to boot.

## Handling Errors

Table A-18. C Calls for Handling Errors

Synopsis	Description
<pre>_call( ... );</pre>	Special version of <i>call</i> that returns error value to caller.
<pre>void nx_perror(   char *string );</pre>	Print an error message corresponding to the current value of <i>errno</i> .

## Floating-Point Control

Table A-19. C Calls for Floating-Point Control

Synopsis	Description
<code>int isnan( double <i>dsrc</i> );</code>	Determine if a <b>double</b> value is Not-a-Number.
<code>int isnand( double <i>dsrc</i> );</code>	Determine if a <b>double</b> value is Not-a-Number.
<code>int isnanf( float <i>fsrc</i> );</code>	Determine if a <b>float</b> value is Not-a-Number.
<code>fp_rnd <b>fpgetround</b>(void);</code>	Get the floating-point rounding mode for the calling process.
<code>fp_rnd <b>fpsetround</b>( fp_rnd <i>rnd_dir</i> );</code>	Set the floating-point rounding mode for the calling process.
<code>fp_except <b>fpgetmask</b>(void);</code>	Get the floating-point exception mask for the calling process.
<code>fp_except <b>fpsetmask</b>( fp_except <i>mask</i> );</code>	Set the floating-point exception mask for the calling process.
<code>fp_except <b>fpgetsticky</b>(void);</code>	Get the floating-point exception sticky flags for the calling process.
<code>fp_except <b>fpsetsticky</b>( fp_except <i>sticky</i> );</code>	Set the floating-point exception sticky flags for the calling process.

## Miscellaneous Calls

Table A-20. Miscellaneous C Calls

Synopsis	Description
<code>void <b>flick</b>(void);</code>	Temporarily relinquish the CPU to another process.
<code>void <b>led</b>( long <i>state</i> );</code>	Turn the node's green LED on or off.
<code>double <b>dclock</b>(void);</code>	Return time in seconds since booting the system.



## iPSC® and Touchstone DELTA Compatibility

Table A-21. C Calls for iPSC® and Touchstone DELTA Compatibility (1 of 2)

Synopsis	Description
void <b>flushmsg</b> ( long <i>typesel</i> , long <i>nodesel</i> , long <i>ptypesel</i> );	Flush specified messages from the system.
long <b>ginv</b> ( long <i>j</i> );	Return the position of an element in the binary-reflected gray code sequence. Inverse of <b>gray</b> ( ).
long <b>gray</b> ( long <i>j</i> );	Return the binary-reflected gray code for an integer.
void <b>hwclock</b> ( esize_t <i>*hwtime</i> );	Place the current value of the hardware counter into a 64-bit unsigned integer variable.
long <b>infopid</b> (void);	Return the process type of the process that sent a pending or received message.
void <b>killcube</b> ( long <i>node</i> , long <i>ptype</i> );	Terminate and clear node process(es).
void <b>killproc</b> ( long <i>node</i> , long <i>ptype</i> );	Terminate a node process.
void <b>load</b> ( char <i>*filename</i> , long <i>node</i> , long <i>ptype</i> );	Load a node process.
unsigned long <b>mclock</b> (void);	Return the time in milliseconds.
void <b>msgcancel</b> ( long <i>mid</i> );	Cancel an asynchronous send or receive operation.
long <b>mypart</b> ( long <i>*rows</i> , long <i>*cols</i> );	Obtain the height and width of the rectangle of nodes allocated to the current application.
long <b>mypid</b> (void);	Return the process type of the calling process.

**Table A-21. C Calls for iPSC® and Touchstone DELTA Compatibility (2 of 2)**

Synopsis	Description
<code>long nodedim(void);</code>	Return the dimension of the current application (the number of nodes allocated to the application is $2^{\text{dimension}}$ ).
<code>long restrictvol(   long <i>fildes</i>,   long <i>nvol</i>,   long <i>vollist</i>[] );</code>	Return 0 (does nothing; provided for compatibility only).

## I/O Modes

**Table A-22. C Calls for I/O Modes**

Synopsis	Description
<code>int gopen(   const char *<i>path</i>,   int <i>oflag</i>,   int <i>iomode</i> [ ,   mode_t <i>mode</i> ] );</code>	Open a file on all nodes and set its I/O mode.
<code>void setiomode(   int <i>fildes</i>,   int <i>iomode</i> );</code>	Set the I/O mode for a file.
<code>long iomode(   int <i>fildes</i> );</code>	Return the current I/O mode for a file.

## Reading and Writing Files in Parallel

Table A-23. C Calls for Reading and Writing Files in Parallel

Synopsis	Description
void <b>cread</b> ( int <i>fildev</i> , char * <i>buffer</i> , unsigned int <i>nbytes</i> );	Read from a file, waiting for completion.
void <b>cwrite</b> ( int <i>fildev</i> , char * <i>buffer</i> , unsigned int <i>nbytes</i> );	Write to a file, waiting for completion.
void <b>creadv</b> ( int <i>fildev</i> , struct iovec * <i>iov</i> , int <i>iovcnt</i> );	Read from a file to irregularly-scattered buffers, waiting for completion.
void <b>cwritev</b> ( int <i>fildev</i> , struct iovec * <i>iov</i> , int <i>iovcnt</i> );	Write to a file from irregularly-scattered buffers, waiting for completion.
long <b>iread</b> ( int <i>fildev</i> , char * <i>buffer</i> , unsigned int <i>nbytes</i> );	Read from a file without waiting for completion.
long <b>iwrite</b> ( int <i>fildev</i> , char * <i>buffer</i> , unsigned int <i>nbytes</i> );	Write to a file without waiting for completion.
long <b>ireadv</b> ( int <i>fildev</i> , struct iovec * <i>iov</i> , int <i>iovcnt</i> );	Read from a file to irregularly-scattered buffers, without waiting for completion.
long <b>iwritev</b> ( int <i>fildev</i> , struct iovec * <i>iov</i> , int <i>iovcnt</i> );	Write to a file from irregularly-scattered buffers, without waiting for completion.
long <b>iodone</b> ( long <i>id</i> );	Determine whether an asynchronous I/O operation is complete. If complete, release the I/O ID.
void <b>iowait</b> ( long <i>id</i> );	Wait for completion of an asynchronous I/O operation and release the I/O ID.

## Detecting End-of-File and Moving the File Pointer

**Table A-24. C Calls for Detecting End-of-File and Moving the File Pointer**

Synopsis	Description
long lseof( int <i>fildev</i> );	Test for end-of-file.
off_t lseek( int <i>fildev</i> , off_t <i>offset</i> , int <i>whence</i> );	Move the read/write file pointer.

## Increasing the Size of a File

**Table A-25. C Calls for Increasing the Size of a File**

Synopsis	Description
long lsize( int <i>fildev</i> , off_t <i>offset</i> , int <i>whence</i> );	Increase size of a file.

## Extended File Manipulation

Table A-26. C Calls for Extended File Manipulation

Synopsis	Description
<pre>esize_t esseek(     int <i>fildev</i>,     esize_t <i>offset</i>,     int <i>whence</i> );</pre>	Move file pointer in extended file.
<pre>esize_t esize(     int <i>fildev</i>,     esize_t <i>offset</i>,     int <i>whence</i> );</pre>	Increase size of extended file.
<pre>long estat(     char *<i>path</i>,     struct estat *<i>buffer</i> );</pre>	Get status of extended file from pathname.
<pre>long lestat(     char *<i>path</i>,     struct estat *<i>buffer</i> );</pre>	Get status of extended file or symbolic link from pathname.
<pre>long festat(     int <i>fildev</i>,     struct estat *<i>buffer</i> );</pre>	Get status of open extended file from file descriptor or unit.

## Performing Extended Arithmetic

Table A-27. C Calls for Performing Extended Arithmetic

Synopsis	Description
<code>esize_t eadd(   esize_t e1,   esize_t e2 );</code>	Add two extended integers.
<code>long ecmp(   esize_t e1,   esize_t e2 );</code>	Compare two extended integers.
<code>long ediv(   esize_t e,   long n );</code>	Divide extended integer by integer.
<code>long emod(   esize_t e,   long n );</code>	Give extended integer modulo an integer (remainder when <i>e</i> is divided by <i>n</i> ).
<code>esize_t emul(   esize_t e,   long n );</code>	Multiply extended integer by integer.
<code>esize_t esub(   esize_t e1,   esize_t e2 );</code>	Subtract two extended integers.
<code>void etos(   esize_t e,   char *s );</code>	Convert extended integer to string.
<code>esize_t stoe(   char *s );</code>	Convert string to extended integer.

## Getting Information About PFS File Systems

**Table A-28. C Calls for Getting Information About PFS™ File Systems**

Synopsis	Description
<pre>long getpfsinfo(     struct pfsmntinfo **attrbufp );</pre>	Get PFS-specific information about all mounted PFS file systems.
<pre>int statpfs(     char *path,     struct estatfs *fs_buffer,     struct statpfs *pfs_buffer,     unsigned int pfs_bufsize );</pre>	Get PFS-specific and non-PFS-specific information for the file system containing <i>path</i> .
<pre>long fstatpfs(     int fildes,     struct estatfs *fs_buffer,     struct statpfs *pfs_buffer,     unsigned int pfs_bufsize );</pre>	Get PFS-specific and non-PFS-specific information for the file system containing the file open on <i>fildes</i> .

## Managing Pthread Execution

Table A-29. C Calls for Managing Pthread Execution

Synopsis	Description
int <b>pthread_create</b> ( pthread_t * <i>thread</i> , pthread_attr_t <i>attr</i> , void *(* <i>routine</i> )(void * <i>arg</i> ), void * <i>arg</i> );	Creates a pthread.
pthread_t <b>pthread_self</b> (void);	Returns the ID of the calling pthread.
int <b>pthread_equal</b> ( pthread_t <i>thread1</i> , pthread_t <i>thread2</i> );	Compares two pthread identifiers.
void <b>pthread_yield</b> (void);	Allows the scheduler to run another pthread instead of the current one.
void <b>pthread_exit</b> ( void * <i>status</i> );	Terminates the calling pthread.
int <b>pthread_join</b> ( pthread_t <i>thread</i> , void ** <i>status</i> );	Waits for a pthread to terminate.
int <b>pthread_detach</b> ( pthread_t * <i>thread</i> );	Detaches a pthread.

## Managing Pthread Attributes

Table A-30. C Calls for Managing Pthread Attributes

Synopsis	Description
int <b>pthread_attr_create</b> ( pthread_attr_t * <i>attr</i> );	Creates a pthread attributes object.
int <b>pthread_attr_setstacksize</b> ( pthread_attr_t * <i>attr</i> , long <i>stacksize</i> );	Sets the value of the stack size attribute of a pthread attributes object.
int <b>pthread_attr_delete</b> ( pthread_attr_t * <i>attr</i> );	Deletes a pthread attributes object.
int <b>pthread_attr_getstacksize</b> ( pthread_attr_t <i>attr</i> );	Returns the value of the stack size attribute of a pthread attributes object.



## Managing Mutexes

Table A-31. C Calls for Managing Mutexes

Synopsis	Description
int <code>pthread_mutex_init</code> ( pthread_mutex_t * <i>mutex</i> , pthread_mutexattr_t <i>attr</i> );	Creates a mutex.
int <code>pthread_mutex_lock</code> ( pthread_mutex_t * <i>mutex</i> );	Locks a mutex.
int <code>pthread_mutex_trylock</code> ( pthread_mutex_t * <i>mutex</i> );	Tries once to lock a mutex.
int <code>pthread_mutex_unlock</code> ( pthread_mutex_t * <i>mutex</i> );	Unlocks a mutex.
int <code>pthread_mutex_destroy</code> ( pthread_mutex_t * <i>mutex</i> );	Deletes a mutex.
int <code>pthread_mutexattr_create</code> ( pthread_mutexattr_t * <i>attr</i> );	Creates a mutex attributes object.
int <code>pthread_mutexattr_delete</code> ( pthread_mutexattr_t * <i>attr</i> );	Deletes a mutex attributes object.

## Using Condition Variables to Synchronize Pthreads

Table A-32. C Calls for Using Condition Variables to Synchronize Pthreads

Synopsis	Description
<code>int pthread_cond_init( pthread_cond_t *cond, pthread_condattr_t attr );</code>	Creates a condition variable.
<code>int pthread_cond_wait( pthread_cond_t *cond, pthread_mutex_t *mutex );</code>	Waits on a condition variable.
<code>int pthread_cond_timedwait( pthread_cond_t *cond, pthread_mutex_t *mutex, struct timespec *abstime );</code>	Waits on a condition variable for a specified period of time.
<code>int pthread_cond_signal( pthread_cond_t *cond );</code>	Wakes up a pthread that is waiting on a condition variable.
<code>int pthread_cond_broadcast( pthread_cond_t *cond );</code>	Wakes up all pthreads that are waiting on a condition variable.
<code>int pthread_cond_destroy( pthread_cond_t *cond );</code>	Destroys a condition variable.
<code>int pthread_condattr_create( pthread_condattr_t *attr );</code>	Creates a condition variable attributes object.
<code>int pthread_condattr_delete( pthread_condattr_t *attr );</code>	Deletes a condition variable attributes object.

## Canceling Pthreads

Table A-33. C Calls for Canceling Pthreads

Synopsis	Description
<code>int pthread_cancel( pthread_t thread );</code>	Requests cancellation of a pthread.
<code>int pthread_setcancel( int state );</code>	Enables or disables the general cancelability of the calling pthread.
<code>int pthread_setsynccancel( int state );</code>	Enables or disables the asynchronous cancelability of the calling pthread.
<code>void pthread_testcancel(void);</code>	Creates a cancellation point in the calling pthread.

## Pthreads Cleanup Routines

Table A-34. C Calls for Pthreads Cleanup Routines

Synopsis	Description
void <b>pthread_cleanup_pop</b> ( int <i>execute</i> );	Removes a routine from the top of the cleanup stack of the calling pthread and optionally executes it.
void <b>pthread_cleanup_push</b> ( void (* <i>routine</i> )(void * <i>arg</i> ), void * <i>arg</i> );	Pushes a routine onto the cleanup stack of the calling pthread.

## Managing Pthread Keys

Table A-35. C Calls for Managing Pthread Keys

Synopsis	Description
int <b>pthread_keycreate</b> ( pthread_key_t * <i>key</i> , void (* <i>destructor</i> )(void * <i>value</i> ) );	Creates a key to be used with pthread-specific data.
int <b>pthread_setspecific</b> ( pthread_key_t <i>key</i> , void * <i>value</i> );	Binds a pthread-specific value to a key.
int <b>pthread_getspecific</b> ( pthread_key_t <i>key</i> , void ** <i>value</i> );	Returns the value bound to a key.

## Miscellaneous Pthread Calls

Table A-36. Miscellaneous Pthread Calls

Synopsis	Description
int <b>pthread_once</b> ( pthread_once_t * <i>once_block</i> , void(* <i>routine</i> )() );	Calls an initialization routine.
int <b>sigwait</b> ( sigset_t * <i>set</i> );	Suspends the calling pthread until one of a specified set of signals is received.

## Fortran System Call Summary

This section summarizes the Fortran versions of the system calls discussed in Chapter 3, Chapter 4, and Chapter 5. See the *Paragon™ Fortran System Calls Reference Manual* for more information on these calls.

### Process Characteristics

Table A-37. Fortran Calls for Process Characteristics

Synopsis	Description
INTEGER FUNCTION <b>MYNODE</b> ()	Obtain the calling process's node number.
INTEGER FUNCTION <b>NUMNODES</b> ()	Obtain the number of nodes allocated to the current application.
SUBROUTINE <b>SETPTYPE</b> ( <i>ptype</i> ) INTEGER <i>ptype</i>	Set the calling process's process type (only permitted if the process type is currently <b>INVALID_PTYPE</b> ).
INTEGER FUNCTION <b>MYPTYPE</b> ()	Obtain the calling process's process type.
INTEGER FUNCTION <b>MYHOST</b> ()	Obtain the controlling process's node number.

## Synchronous Send and Receive

Table A-38. Fortran Calls for Synchronous Send and Receive

Synopsis	Description
SUBROUTINE CSEND( <i>type, buf, count, node, ptype</i> )  INTEGER <i>type</i> INTEGER <i>buf</i> (*) INTEGER <i>count</i> INTEGER <i>node</i> INTEGER <i>ptype</i>	Send a message, waiting for completion.
SUBROUTINE CRECV( <i>typesel, buf, count</i> )  INTEGER <i>typesel</i> INTEGER <i>buf</i> (*) INTEGER <i>count</i>	Receive a message, waiting for completion.
INTEGER FUNCTION CSENDRECV( <i>type, sbuf, scount, node, ptype, typesel, rbuf, rcount</i> )  INTEGER <i>type</i> INTEGER <i>sbuf</i> (*) INTEGER <i>scount</i> INTEGER <i>node</i> INTEGER <i>ptype</i> INTEGER <i>typesel</i> INTEGER <i>rbuf</i> (*) INTEGER <i>rcount</i>	Send a message and post a receive for the reply. Wait for completion.
SUBROUTINE GSENDX( <i>type, buf, count, nodes, nodecount</i> )  INTEGER <i>type</i> INTEGER <i>buf</i> (*) INTEGER <i>count</i> INTEGER <i>nodes</i> (*) INTEGER <i>nodecount</i>	Send a message to a list of nodes, waiting for completion.

## Asynchronous Send and Receive

Table A-39. Fortran Calls for Asynchronous Send and Receive (1 of 2)

Synopsis	Description
<p>INTEGER FUNCTION <b>ISEND</b>(<i>type, buf, count, node, ptype</i>)</p> <p>INTEGER <i>type</i>            INTEGER <i>buf</i>(*)            INTEGER <i>count</i>            INTEGER <i>node</i>            INTEGER <i>ptype</i></p>	Send a message without waiting for completion.
<p>INTEGER FUNCTION <b>Irecv</b>(<i>typesel, buf, count</i>)</p> <p>INTEGER <i>typesel</i>            INTEGER <i>buf</i>(*)            INTEGER <i>count</i></p>	Receive a message without waiting for completion.
<p>INTEGER FUNCTION <b>ISENDRECV</b>(<i>type, sbuf, scount, node, ptype, typesel, rbuf, rcount</i>)</p> <p>INTEGER <i>type</i>            INTEGER <i>sbuf</i>(*)            INTEGER <i>scount</i>            INTEGER <i>node</i>            INTEGER <i>ptype</i>            INTEGER <i>typesel</i>            INTEGER <i>rbuf</i>(*)            INTEGER <i>rcount</i></p>	Send a message and post a receive for the reply without waiting for completion.
<p>INTEGER FUNCTION <b>MSGDONE</b>(<i>mid</i>)</p> <p>INTEGER <i>mid</i></p>	Determine whether a send or receive operation has completed.
<p>SUBROUTINE <b>MSGWAIT</b>(<i>mid</i>)</p> <p>INTEGER <i>mid</i></p>	Wait for completion of a send or receive operation.

**Table A-39. Fortran Calls for Asynchronous Send and Receive (2 of 2)**

Synopsis	Description
SUBROUTINE MSGIGNORE( <i>mid</i> )  INTEGER <i>mid</i>	Release a message ID as soon as a send or receive operation completes.
INTEGER FUNCTION MSGMERGE( <i>mid1</i> , <i>mid2</i> )  INTEGER <i>mid1</i> INTEGER <i>mid2</i>	Merge two message IDs into a single ID that can be used to wait for completion of both operations.

## Probing for Pending Messages

**Table A-40. Fortran Calls for Probing for Pending Messages**

Synopsis	Description
SUBROUTINE CPROBE( <i>typesel</i> )  INTEGER <i>typesel</i>	Wait for a message of a selected type to arrive.
INTEGER FUNCTION IPROBE( <i>typesel</i> )  INTEGER <i>typesel</i>	Determine whether a message of a selected type is pending.

## Getting Information About Pending or Received Messages

**Table A-41. Fortran Calls for Getting Information About Pending or Received Messages**

Synopsis	Description
INTEGER FUNCTION INFOCOUNT()	Return size in bytes of a pending or received message.
INTEGER FUNCTION INFONODE()	Return node number of the node that sent a pending or received message.
INTEGER FUNCTION INFOPTYPE()	Return process type of the process that sent a pending or received message.
INTEGER FUNCTION INFOTYPE()	Return message type of a pending or received message.

## Treating a Message as an Interrupt

Table A-42. Fortran Calls for Treating a Message as an Interrupt (1 of 2)

Synopsis	Description
<p>SUBROUTINE <b>HSEND</b>(<i>type, buf, count, node, ptype, handler</i>)</p> <p>INTEGER <i>type</i>            INTEGER <i>buf</i>(*)            INTEGER <i>count</i>            INTEGER <i>node</i>            INTEGER <i>ptype</i>            EXTERNAL <i>handler</i></p>	<p>Send a message and set up a handler procedure to be called when the send completes.</p>
<p>SUBROUTINE <b>HRCV</b> (<i>typesel, buf, count, handler</i>)</p> <p>INTEGER <i>typesel</i>            INTEGER <i>buf</i>(*)            INTEGER <i>count</i>            EXTERNAL <i>handler</i></p>	<p>Receive a message and set up a handler procedure to be called when the receive completes.</p>
<p>SUBROUTINE <b>HSENDRCV</b>(<i>type, sbuf, scount, node, ptype, typesel, rbuf, rcount, handler</i>)</p> <p>INTEGER <i>type</i>            INTEGER <i>sbuf</i>(*)            INTEGER <i>scount</i>            INTEGER <i>node</i>            INTEGER <i>ptype</i>            INTEGER <i>typesel</i>            INTEGER <i>rbuf</i>(*)            INTEGER <i>rcount</i>            EXTERNAL <i>handler</i></p>	<p>Send a message and post a receive for the reply. Set up a handler procedure to be called when the reply arrives.</p>



**Table A-42. Fortran Calls for Treating a Message as an Interrupt (2 of 2)**

Synopsis	Description
<b>INTEGER FUNCTION MASKTRAP</b> ( <i>state</i> ) <b>INTEGER</b> <i>state</i>	Enable or disable interrupts for message handlers. Required to prevent corruption of global variables.
<b>SUBROUTINE HSENDX</b> ( <i>type, buf, count, node, ptype, xhandler, hparam</i> ) <b>INTEGER</b> <i>type</i> <b>INTEGER</b> <i>buf</i> (*) <b>INTEGER</b> <i>count</i> <b>INTEGER</b> <i>node</i> <b>INTEGER</b> <i>ptype</i> <b>EXTERNAL</b> <i>xhandler</i> <b>INTEGER</b> <i>hparam</i>	Send a message and set up an extended handler procedure to be called with the value <i>hparam</i> when the send completes. Allows handler sharing.

## Extended Receive and Probe

**Table A-43. Fortran Calls for Extended Receive and Probe (1 of 2)**

Synopsis	Description
<b>SUBROUTINE CRECVX</b> ( <i>typesel, buf, count, nodesel, ptypesel, info</i> ) <b>INTEGER</b> <i>typesel</i> <b>INTEGER</b> <i>buf</i> (*) <b>INTEGER</b> <i>count</i> <b>INTEGER</b> <i>nodesel</i> <b>INTEGER</b> <i>ptypesel</i> <b>INTEGER</b> <i>info</i> (8)	Receive a message of a specified type from a specified sending node and process type, together with information about the message. Wait for completion.
<b>INTEGER FUNCTION IRECVX</b> ( <i>typesel, buf, count, nodesel, ptypesel, info</i> ) <b>INTEGER</b> <i>typesel</i> <b>INTEGER</b> <i>buf</i> (*) <b>INTEGER</b> <i>count</i> <b>INTEGER</b> <i>nodesel</i> <b>INTEGER</b> <i>ptypesel</i> <b>INTEGER</b> <i>info</i> (8)	Receive a message of a specified type from a specified sending node and process type, together with information about the message. Do not wait for completion.

Table A-43. Fortran Calls for Extended Receive and Probe (2 of 2)

Synopsis	Description
<p data-bbox="399 464 889 520"><b>SUBROUTINE HRECVX</b>(<i>typesel, buf, count, nodesel, ptypesel, xhandler, hparam</i>)</p> <p data-bbox="399 558 594 590">INTEGER <i>typesel</i></p> <p data-bbox="399 590 586 621">INTEGER <i>buf</i>(*)</p> <p data-bbox="399 621 578 653">INTEGER <i>count</i></p> <p data-bbox="399 653 602 684">INTEGER <i>nodesel</i></p> <p data-bbox="399 684 610 716">INTEGER <i>ptypesel</i></p> <p data-bbox="399 716 643 747">EXTERNAL <i>xhandler</i></p> <p data-bbox="399 747 602 779">INTEGER <i>hparam</i></p>	<p data-bbox="922 464 1406 621">Receive a message of a specified type from a specified sending node and process type. Set up an extended handler procedure to be called with information about the message and the value <i>hparam</i> when the receive completes.</p>
<p data-bbox="399 800 889 856"><b>SUBROUTINE CPROBEX</b>(<i>typesel, nodesel, ptypesel, info</i>)</p> <p data-bbox="399 894 594 926">INTEGER <i>typesel</i></p> <p data-bbox="399 926 602 957">INTEGER <i>nodesel</i></p> <p data-bbox="399 957 610 989">INTEGER <i>ptypesel</i></p> <p data-bbox="399 989 594 1020">INTEGER <i>info</i>(8)</p>	<p data-bbox="922 800 1406 894">Wait for a message of a specified type from a specified sending node and process type. Return information about the message.</p>
<p data-bbox="399 1045 870 1102"><b>INTEGER FUNCTION IPROBEX</b>(<i>typesel, nodesel, ptypesel, info</i>)</p> <p data-bbox="399 1140 594 1171">INTEGER <i>typesel</i></p> <p data-bbox="399 1171 602 1203">INTEGER <i>nodesel</i></p> <p data-bbox="399 1203 610 1234">INTEGER <i>ptypesel</i></p> <p data-bbox="399 1234 594 1266">INTEGER <i>info</i>(8)</p>	<p data-bbox="922 1045 1406 1161">Determine whether a message of a specified type from a specified sending node and process type is pending. If it is, return information about the message.</p>

## Global Operations

**Table A-44. Fortran Calls for Global Operations (1 of 3)**

Synopsis	Description
SUBROUTINE <b>GCOL</b> ( <i>x, xlen, y, ylen, ncnt</i> )  INTEGER <i>x</i> (*) INTEGER <i>xlen</i> INTEGER <i>y</i> (*) INTEGER <i>ylen</i> INTEGER <i>ncnt</i>	Concatenation.
SUBROUTINE <b>GCOLX</b> ( <i>x, xlens, y</i> )  INTEGER <i>x</i> (*) INTEGER <i>xlens</i> (*) INTEGER <i>y</i> (*)	Concatenation for contributions of known length.
SUBROUTINE <b>GDHIGH</b> ( <i>x, n, work</i> )  DOUBLE PRECISION <i>x</i> (*) INTEGER <i>n</i> DOUBLE PRECISION <i>work</i> (*)	Vector double precision MAX.
SUBROUTINE <b>GDLOW</b> ( <i>x, n, work</i> )  DOUBLE PRECISION <i>x</i> (*) INTEGER <i>n</i> DOUBLE PRECISION <i>work</i> (*)	Vector double precision MIN.
SUBROUTINE <b>GDPROD</b> ( <i>x, n, work</i> )  DOUBLE PRECISION <i>x</i> (*) INTEGER <i>n</i> DOUBLE PRECISION <i>work</i> (*)	Vector double precision MULTIPLY.
SUBROUTINE <b>GDSUM</b> ( <i>x, n, work</i> )  DOUBLE PRECISION <i>x</i> (*) INTEGER <i>n</i> DOUBLE PRECISION <i>work</i> (*)	Vector double precision SUM.
SUBROUTINE <b>GIAND</b> ( <i>x, n, work</i> )  INTEGER <i>x</i> (*) INTEGER <i>n</i> INTEGER <i>work</i> (*)	Vector integer bitwise AND.

Table A-44. Fortran Calls for Global Operations (2 of 3)

Synopsis	Description
SUBROUTINE <b>GIHIGH</b> ( <i>x, n, work</i> )  INTEGER <i>x</i> (*) INTEGER <i>n</i> INTEGER <i>work</i> (*)	Vector integer MAX.
SUBROUTINE <b>GILOW</b> ( <i>x, n, work</i> )  INTEGER <i>x</i> (*) INTEGER <i>n</i> INTEGER <i>work</i> (*)	Vector integer MIN.
SUBROUTINE <b>GIOR</b> ( <i>x, n, work</i> )  INTEGER <i>x</i> (*) INTEGER <i>n</i> INTEGER <i>work</i> (*)	Vector integer bitwise OR.
SUBROUTINE <b>GIPROD</b> ( <i>x, n, work</i> )  INTEGER <i>x</i> (*) INTEGER <i>n</i> INTEGER <i>work</i> (*)	Vector integer MULTIPLY.
SUBROUTINE <b>GISUM</b> ( <i>x, n, work</i> )  INTEGER <i>x</i> (*) INTEGER <i>n</i> INTEGER <i>work</i> (*)	Vector integer SUM.
SUBROUTINE <b>GLAND</b> ( <i>x, n, work</i> )  LOGICAL <i>x</i> (*) INTEGER <i>n</i> LOGICAL <i>work</i> (*)	Vector logical AND.
SUBROUTINE <b>GLOR</b> ( <i>x, n, work</i> )  LOGICAL <i>x</i> (*) INTEGER <i>n</i> LOGICAL <i>work</i> (*)	Vector logical inclusive OR.

Table A-44. Fortran Calls for Global Operations (3 of 3)

Synopsis	Description
SUBROUTINE <b>GOPF</b> ( <i>x, xlen, work, function</i> )  INTEGER <i>x</i> (*) INTEGER <i>xlen</i> INTEGER <i>work</i> (*) EXTERNAL <i>function</i>	Arbitrary commutative function.
SUBROUTINE <b>GSHIGH</b> ( <i>x, n, work</i> )  REAL <i>x</i> (*) INTEGER <i>n</i> REAL <i>work</i> (*)	Vector real MAX.
SUBROUTINE <b>GSLOW</b> ( <i>x, n, work</i> )  REAL <i>x</i> (*) INTEGER <i>n</i> REAL <i>work</i> (*)	Vector real MIN.
SUBROUTINE <b>GSPROD</b> ( <i>x, n, work</i> )  REAL <i>x</i> (*) INTEGER <i>n</i> REAL <i>work</i> (*)	Vector real MULTIPLY.
SUBROUTINE <b>GSSUM</b> ( <i>x, n, work</i> )  REAL <i>x</i> (*) INTEGER <i>n</i> REAL <i>work</i> (*)	Vector real SUM.
SUBROUTINE <b>GSYNC</b> ()	Global synchronization.

## Controlling Application Execution

**Table A-45. Fortran Calls for Controlling Application Execution (1 of 2)**

Synopsis	Description
INTEGER FUNCTION <b>NX_INITVE</b> ( <i>partition, size, account,</i> <i>argc, argv</i> )  CHARACTER <i>partition</i> *(*) INTEGER <i>size</i> CHARACTER <i>account</i> *(*) INTEGER <i>argc</i> INTEGER <i>argv</i>	Create a new application.
INTEGER FUNCTION <b>NX_INITVE_RECT</b> ( <i>partition, anchor,</i> <i>rows, cols, account, argc, argv</i> )  CHARACTER <i>partition</i> *(*) INTEGER <i>anchor</i> INTEGER <i>rows</i> INTEGER <i>cols</i> CHARACTER <i>account</i> *(*) INTEGER <i>argc</i> INTEGER <i>argv</i>	Create a new application with a rectangular shape.
INTEGER FUNCTION <b>NX_PRI</b> ( <i>pgroup,</i> <i>priority</i> )  INTEGER <i>pgroup</i> INTEGER <i>priority</i>	Set the priority of an application.
INTEGER FUNCTION <b>NX_NFORK</b> ( <i>node_list, numnodes,</i> <i>pctype, pid_list</i> )  INTEGER <i>node_list</i> (*) INTEGER <i>numnodes</i> INTEGER <i>pctype</i> INTEGER <i>pid_list</i> (*)	Copy the current process onto some or all nodes of an application.

Table A-45. Fortran Calls for Controlling Application Execution (2 of 2)

Synopsis	Description
INTEGER FUNCTION <b>NX_LOAD</b> ( <i>node_list</i> , <i>numnodes</i> , <i>ptype</i> , <i>pid_list</i> , <i>pathname</i> )  INTEGER <i>node_list</i> (*) INTEGER <i>numnodes</i> INTEGER <i>ptype</i> INTEGER <i>pid_list</i> (*) CHARACTER <i>pathname</i> *(*)	Execute a stored program on some or all nodes of an application.
INTEGER FUNCTION <b>NX_LOADVE</b> ( <i>node_list</i> , <i>numnodes</i> , <i>ptype</i> , <i>pid_list</i> , <i>pathname</i> , <i>argv</i> , <i>envp</i> )  INTEGER <i>node_list</i> (*) INTEGER <i>numnodes</i> INTEGER <i>ptype</i> INTEGER <i>pid_list</i> (*) CHARACTER <i>pathname</i> *(*) INTEGER <i>argv</i> INTEGER <i>envp</i>	Execute a stored program on some or all nodes of an application, with specified argument list and environment.
SUBROUTINE <b>NX_WAITALL</b> ()	Wait for all application processes.

## Getting Information About Applications

Table A-46. Fortran Calls for Getting Information About Applications

Synopsis	Description
INTEGER FUNCTION <b>NX_APP_RECT</b> ( <i>rows</i> , <i>cols</i> )  INTEGER <i>rows</i> INTEGER <i>cols</i>	Obtain the height and width of the rectangle of nodes allocated to the current application.
INTEGER FUNCTION <b>NX_APP_NODES</b> ( <i>pgroup</i> , <i>ptr</i> , <i>list_size</i> )  INTEGER <i>pgroup</i> POINTER ( <i>ptr</i> , <i>node_list</i> (1)) INTEGER <i>list_size</i>	List the nodes allocated to an application.

## Partition Management

**Table A-47. Fortran Calls for Partition Management (1 of 2)**

Synopsis	Description
INTEGER FUNCTION <b>NX_MKPART</b> ( <i>partition, size, type</i> )  CHARACTER <i>partition</i> (*) INTEGER <i>size</i> INTEGER <i>type</i>	Create a partition with a particular number of nodes.
INTEGER FUNCTION <b>NX_MKPART_RECT</b> ( <i>partition, rows, cols, type</i> )  CHARACTER <i>partition</i> (*) INTEGER <i>rows</i> INTEGER <i>cols</i> INTEGER <i>type</i>	Create a partition with a particular height and width.
INTEGER FUNCTION <b>NX_MKPART_MAP</b> ( <i>partition, numnodes, node_list, type</i> )  CHARACTER <i>partition</i> (*) INTEGER <i>numnodes</i> INTEGER <i>node_list</i> (*) INTEGER <i>type</i>	Create a partition with a specific set of nodes.
INTEGER FUNCTION <b>NX_RMPART</b> ( <i>pathname, force, recursive</i> )  CHARACTER <i>partition</i> (*) INTEGER <i>force</i> INTEGER <i>recursive</i>	Remove a partition.
INTEGER FUNCTION <b>NX_PART_ATTR</b> ( <i>partition, attributes</i> )  CHARACTER <i>partition</i> (*) RECORD /nx_part_info_t/ <i>attributes</i>	Get a partition's attributes.



Table A-47. Fortran Calls for Partition Management (2 of 2)

Synopsis	Description
INTEGER FUNCTION <b>NX_PART_NODES</b> ( <i>partition, ptr, list_size</i> )  CHARACTER <i>partition</i> *(*) POINTER ( <i>ptr, node_list</i> (1)) INTEGER <i>list_size</i>	List the root node numbers for the nodes of a partition.
INTEGER FUNCTION <b>NX_CHPART_NAME</b> ( <i>partition, name</i> )  CHARACTER <i>partition</i> *(*) CHARACTER <i>name</i> *(*)	Change a partition's name.
INTEGER FUNCTION <b>NX_CHPART_MOD</b> ( <i>partition, mode</i> )  CHARACTER <i>partition</i> *(*) INTEGER <i>mode</i>	Change a partition's protection modes.
INTEGER FUNCTION <b>NX_CHPART_EPL</b> ( <i>partition, priority</i> )  CHARACTER <i>partition</i> *(*) INTEGER <i>priority</i>	Change a partition's effective priority limit.
INTEGER FUNCTION <b>NX_CHPART_RQ</b> ( <i>partition, rollin_quantum</i> )  CHARACTER <i>partition</i> *(*) INTEGER <i>rollin_quantum</i>	Change a partition's rollin quantum.
INTEGER FUNCTION <b>NX_CHPART_OWNER</b> ( <i>partition, owner, group</i> )  CHARACTER <i>partition</i> *(*) INTEGER <i>owner</i> INTEGER <i>group</i>	Change a partition's owner and group.
INTEGER FUNCTION <b>NX_CHPART_SCHED</b> ( <i>partition, sched_type</i> )  CHARACTER <i>partition</i> *(*) INTEGER <i>sched_type</i>	Change a partition's rollin quantum.

## Finding Unusable Nodes

**Table A-48. Fortran Calls for Finding Unusable Nodes**

Synopsis	Description
INTEGER FUNCTION <b>NX_EMPTY_NODES</b> ( <i>ptr, list_size</i> )  POINTER ( <i>ptr, node_list</i> (1)) INTEGER <i>list_size</i>	List the nodes that are empty slots.
INTEGER FUNCTION <b>NX_FAILED_NODES</b> ( <i>ptr, list_size</i> )  POINTER ( <i>ptr, node_list</i> (1)) INTEGER <i>list_size</i>	List the nodes that failed to boot.

## Handling Errors

**Table A-49. Fortran Calls for Handling Errors**

Synopsis	Description
SUBROUTINE <b>NX_PERROR</b> ( <i>string</i> )  CHARACTER <i>string</i> *(*)	Print an error message corresponding to the current value of <i>errno</i> .

## Floating-Point Control

**Table A-50. Fortran Calls for Floating-Point Control**

Synopsis	Description
INTEGER FUNCTION <b>FPSETMASK</b> ( <i>mask</i> )  INTEGER <i>mask</i>	Set the floating-point exception mask for the calling process.

## Miscellaneous Calls

Table A-51. Miscellaneous Fortran Calls

Synopsis	Description
SUBROUTINE <b>FLICK</b> ()	Temporarily relinquish the CPU to another process.
SUBROUTINE <b>LED</b> ( <i>state</i> ) INTEGER <i>state</i>	Turn the node's green LED on or off.
DOUBLE PRECISION FUNCTION <b>DCLOCK</b> ()	Return time in seconds since booting the system.

## iPSC® and Touchstone DELTA Compatibility

Table A-52. Fortran Calls for iPSC® and Touchstone DELTA Compatibility (1 of 2)

Synopsis	Description
SUBROUTINE <b>FLUSHMSG</b> ( <i>typesel</i> , <i>nodesel</i> , <i>ptypesel</i> ) INTEGER <i>typesel</i> INTEGER <i>nodesel</i> INTEGER <i>ptypesel</i>	Flush specified messages from the system.
INTEGER FUNCTION <b>GINV</b> ( <i>graycode</i> ) INTEGER <i>graycode</i>	Return the position of an element in the binary-reflected gray code sequence. Inverse of <b>gray</b> ().
INTEGER FUNCTION <b>GRAY</b> ( <i>position</i> ) INTEGER <i>position</i>	Return the binary-reflected gray code for an integer.
SUBROUTINE <b>HWCLOCK</b> ( <i>hwtime</i> ) INTEGER <i>hwtime</i> (2)	Place the current value of the hardware counter into a 64-bit unsigned integer variable.
INTEGER FUNCTION <b>INFOPID</b> ()	Return the process type of the process that sent a pending or received message.
SUBROUTINE <b>KILLCUBE</b> ( <i>node</i> , <i>pid</i> ) INTEGER <i>node</i> INTEGER <i>pid</i>	Terminate and clear node process(es).

Table A-52. Fortran Calls for iPSC® and Touchstone DELTA Compatibility (2 of 2)

Synopsis	Description
SUBROUTINE KILLPROC( <i>node, pid</i> )  INTEGER <i>node</i> INTEGER <i>pid</i>	Terminate a node process.
SUBROUTINE LOAD( <i>filename, node, pid</i> )  CHARACTER <i>filename</i> *(*) INTEGER <i>node</i> INTEGER <i>pid</i>	Load a node process.
INTEGER FUNCTION MCLOCK()	Return the time in milliseconds.
SUBROUTINE MSGCANCEL( <i>mid</i> )  INTEGER <i>mid</i>	Cancel an asynchronous send or receive operation.
INTEGER FUNCTION MYPART( <i>rows, cols</i> )  INTEGER <i>rows</i> INTEGER <i>cols</i>	Obtain the height and width of the rectangle of nodes allocated to the current application.
INTEGER FUNCTION MYPID()	Return the process type of the calling process.
INTEGER FUNCTION NODEDIM()	Return the dimension of the current application (the number of nodes allocated to the application is $2^{\text{dimension}}$ ).
INTEGER FUNCTION RESTRICTVOL( <i>unit, nvol, volist</i> )  INTEGER <i>unit</i> INTEGER <i>nvol</i> INTEGER <i>volist</i> (*)	Return 0 (does nothing; provided for compatibility only).

## I/O Modes

Table A-53. Fortran Calls for I/O Modes

Synopsis	Description
SUBROUTINE <b>GOPEN</b> ( <i>unit, path, iomode</i> )  INTEGER <i>unit</i> CHARACTER <i>path</i> (*) INTEGER <i>iomode</i>	Open a file on all nodes and set its I/O mode.
SUBROUTINE <b>SETIOMODE</b> ( <i>unit, iomode</i> )  INTEGER <i>unit</i> INTEGER <i>iomode</i>	Set the I/O mode for a file.
INTEGER FUNCTION <b>IOMODE</b> ( <i>unit</i> )  INTEGER <i>unit</i>	Return the current I/O mode for a file.

## Reading and Writing Files in Parallel

Table A-54. Fortran Calls for Reading and Writing Files in Parallel (1 of 2)

Synopsis	Description
SUBROUTINE <b>CREAD</b> ( <i>unit, buffer, nbytes</i> )  INTEGER <i>unit</i> INTEGER <i>buffer</i> (*) INTEGER <i>nbytes</i>	Read from a file, waiting for completion.
SUBROUTINE <b>CWRITE</b> ( <i>unit, buffer, nbytes</i> )  INTEGER <i>unit</i> INTEGER <i>buffer</i> (*) INTEGER <i>nbytes</i>	Write to a file, waiting for completion.
SUBROUTINE <b>CREADV</b> ( <i>unit, iov, iovcnt</i> )  INTEGER <i>unit</i> INTEGER <i>iov</i> (*) INTEGER <i>iovcnt</i>	Read from a file to irregularly-scattered buffers, waiting for completion.

Table A-54. Fortran Calls for Reading and Writing Files in Parallel (2 of 2)

Synopsis	Description
<b>SUBROUTINE CWRITEV</b> ( <i>unit, iov, iovcnt</i> )  INTEGER <i>unit</i> INTEGER <i>iov</i> (*) INTEGER <i>iovcnt</i>	Write to a file from irregularly-scattered buffers, waiting for completion.
<b>INTEGER FUNCTION IREAD</b> ( <i>unit, buffer, nbytes</i> )  INTEGER <i>unit</i> INTEGER <i>buffer</i> (*) INTEGER <i>nbytes</i>	Read from a file without waiting for completion.
<b>INTEGER FUNCTION IWRITE</b> ( <i>unit, buffer, nbytes</i> )  INTEGER <i>unit</i> INTEGER <i>buffer</i> (*) INTEGER <i>nbytes</i>	Write to a file, waiting for completion.
<b>INTEGER FUNCTION IREADV</b> ( <i>unit, iov, iovcnt</i> )  INTEGER <i>unit</i> INTEGER <i>iov</i> (*) INTEGER <i>iovcnt</i>	Read from a file to irregularly-scattered buffers, without waiting for completion.
<b>INTEGER FUNCTION IWRITEV</b> ( <i>unit, iov, iovcnt</i> )  INTEGER <i>unit</i> INTEGER <i>iov</i> (*) INTEGER <i>iovcnt</i>	Write to a file from irregularly-scattered buffers, without waiting for completion.
<b>INTEGER FUNCTION IODONE</b> ( <i>id</i> )  INTEGER <i>id</i>	Determine whether an asynchronous I/O operation is complete. If complete, release the I/O ID.
<b>SUBROUTINE IOWAIT</b> ( <i>id</i> )  INTEGER <i>id</i>	Wait for completion of an asynchronous I/O operation and release the I/O ID.

## Detecting End-of-File and Moving the File Pointer

**Table A-55. Fortran Calls for Detecting End-of-File and Moving the File Pointer**

Synopsis	Description
INTEGER FUNCTION ISEOF( <i>unit</i> )  INTEGER <i>unit</i>	Test for end-of-file.
INTEGER FUNCTION LSEEK( <i>unit, offset, whence</i> )  INTEGER <i>unit</i> INTEGER <i>offset</i> INTEGER <i>whence</i>	Move the read/write file pointer.

## Flushing Fortran Buffered I/O

**Table A-56. Fortran Calls for Flushing Buffered I/O**

Synopsis	Description
SUBROUTINE FORCEFLUSH()	Cause all buffered I/O to be flushed if an exception occurs.
SUBROUTINE FORFLUSH( <i>unit</i> )  INTEGER <i>unit</i>	Flush all buffered I/O on a particular unit.

## Increasing the Size of a File

**Table A-57. Fortran Calls for Increasing the Size of a File**

Synopsis	Description
INTEGER FUNCTION LSIZE( <i>unit, offset, whence</i> )  INTEGER <i>unit</i> INTEGER <i>offset</i> INTEGER <i>whence</i>	Increase size of a file.

## Extended File Manipulation

**Table A-58. Fortran Calls for Extended File Manipulation**

Synopsis	Description
<p data-bbox="383 537 873 600">SUBROUTINE <b>ESEEK</b>(<i>unit, offset, whence, newpos</i>)</p> <p data-bbox="383 632 613 758">INTEGER <i>unit</i> INTEGER <i>offset</i>(2) INTEGER <i>whence</i> INTEGER <i>newpos</i>(2)</p>	<p data-bbox="902 537 1268 569">Move file pointer in extended file.</p>
<p data-bbox="383 779 862 842">SUBROUTINE <b>ESIZE</b>(<i>unit, offset, whence, newsize</i>)</p> <p data-bbox="383 873 618 999">INTEGER <i>unit</i> INTEGER <i>offset</i>(2) INTEGER <i>whence</i> INTEGER <i>newsize</i>(2)</p>	<p data-bbox="902 779 1219 810">Increase size of extended file.</p>



## Performing Extended Arithmetic

Table A-59. Fortran Calls for Performing Extended Arithmetic

Synopsis	Description
SUBROUTINE <b>EADD</b> ( <i>e1, e2, eresult</i> )  INTEGER <i>e1</i> (2) INTEGER <i>e2</i> (2) INTEGER <i>eresult</i> (2)	Add two extended integers.
INTEGER FUNCTION <b>ECMP</b> ( <i>e1, e2</i> )  INTEGER <i>e1</i> (2) INTEGER <i>e2</i> (2)	Compare two extended integers.
SUBROUTINE <b>EDIV</b> ( <i>e, n, result</i> )  INTEGER <i>e</i> (2) INTEGER <i>n</i> INTEGER <i>result</i>	Divide extended integer by integer.
SUBROUTINE <b>EMOD</b> ( <i>e, n, result</i> )  INTEGER <i>e</i> (2) INTEGER <i>n</i> INTEGER <i>result</i>	Give extended integer modulo an integer (remainder when <i>e</i> is divided by <i>n</i> ).
SUBROUTINE <b>EMUL</b> ( <i>e, n, eresult</i> )  INTEGER <i>e</i> (2) INTEGER <i>n</i> INTEGER <i>eresult</i> (2)	Multiply extended integer by integer.
SUBROUTINE <b>ESUB</b> ( <i>e1, e2, eresult</i> )  INTEGER <i>e1</i> (2) INTEGER <i>e2</i> (2) INTEGER <i>eresult</i> (2)	Subtract two extended integers.
SUBROUTINE <b>ETOS</b> ( <i>e, s</i> )  INTEGER <i>e</i> (2) CHARACTER <i>s</i> (*)	Convert extended integer to string.
SUBROUTINE <b>STOE</b> ( <i>s, e</i> )  CHARACTER <i>s</i> (*) INTEGER <i>e</i> (2)	Convert string to extended integer.



# iPSC<sup>®</sup> System Compatibility

**B**

## Introduction

This appendix gives you information you can use to port programs to Paragon<sup>™</sup> OSF/1 from the iPSC<sup>®</sup> series of supercomputers from Intel Supercomputer Systems Division.

This appendix lists the differences between iPSC system commands and system calls and those of Paragon<sup>™</sup> OSF/1, and suggests alternatives that you can use for commands and calls that are not supported. Commands and calls that are not listed here should work the same in Paragon OSF/1 as they do in the iPSC system.

## General Compatibility Issues

In general, iPSC system programs can simply be recompiled and executed on the Paragon system. However, keep in mind the following basic differences between the two systems:

- There is no SRM. The Diagnostic Station is used only for system administration; all software development is done either on remote workstations or on the Paragon system itself. Parallel applications are run only on the Paragon system.
- Host programs are not directly supported. See "Host Calls" on page B-9 for more information.
- The node network is a 2-D mesh rather than a hypercube. You might want to change the data distribution in your application to take advantage of the different system topology.
- An application can run on *any* number of nodes from 1 to the size of the compute partition (up to several thousand nodes). If your application depends on the number of nodes being a power of two or no greater than 128, you should re-write it so that it works on any number of nodes. If this is not possible, you should have the application print an error message if `numnodes()` is not a power of two or is too large for the application to handle.

- If a message arrives at a node before the receive for the message has been posted, the message is stored in a system buffer. In the iPSC system, the space available for these system buffers is the entire free physical memory of the node. In the Paragon system, this space is more limited (1M bytes by default). This limitation results from the fact that the Paragon system supports multiple processes per node.

## NOTE

Because of this limitation, iPSC system applications that use large amounts of system message buffering may slow down or hang on the Paragon system, especially when run on large numbers of nodes.

If this occurs, you can increase the system message buffering space with the **-mbf** switch, as described under “System Message Buffers” on page 8-16. However, it would be better to re-write the application so that receives are always posted before the message arrives, as discussed under “Avoid Message Buffering” on page 8-11.

- The term *process ID*, or *PID*, is used differently. In the iPSC system, each process has a *UNIX PID* used by the OS and an *NX PID* used for message passing. In the Paragon system, the “UNIX PID” is just called the *PID*, and the “NX PID” is called the *process type* or *ptype*. Although the names have changed, the software works the same. For example, **mypid()** and **infpid()** are supported as equivalents to **myptype()** and **infoptype()**. Exception: on the iPSC/860, the NX PID is always 0; in the Paragon system, the process type can be any integer from 0 to 2,147,483,647 ( $2^{31} - 1$ ) inclusive (but is usually 0).
- Force types (special message types that use a limited flow control technique) are fully supported and work the same as they do in the iPSC system. However, in the Paragon system regular messages are just as fast as force type messages, so force types are not needed for performance.

## New Features

Paragon OSF/1 offers the following features that were not available on the iPSC system. You can use these features to improve the performance and readability of your programs.

- You can use the complete set of OSF/1 commands on the Paragon system, as discussed in Chapter 2.
- You can execute an application on multiple nodes just by typing its name on the command line, using command-line switches to control its execution, as discussed under “Running Applications” on page 2-11.
- You can control the values of some important message-passing configuration parameters, as discussed under “Message-Passing Configuration Switches” on page 8-18.

- You can allocate groups of nodes of any size and shape, and control the scheduling characteristics of applications that run in them, as discussed under “Managing Partitions” on page 2-25 and “Managing Partitions” on page 4-27.
- You can have more than one process per node, as discussed under “Process Characteristics” on page 3-3. When sending messages, you specify a process by its *process type* (equivalent to the “NX PID” in the iPSC system).
- You can tell the system to discard an asynchronous message ID as soon as the send or receive completes with `msgignore()`, as discussed under “Asynchronous Send and Receive” on page 3-10.
- You can merge together a number of asynchronous message-passing requests and wait for all of them to complete in a single call with `msgmerge()`, as discussed under “Merging Message IDs” on page 3-13.
- You can pass a parameter to a message interrupt handler with `hsendx()`, as discussed under “Treating a Message as an Interrupt” on page 3-18.
- You can receive or probe for a message based on its sender, and receive information about a message along with the message, with the `...x()` calls, as discussed under “Extended Receive and Probe” on page 3-24.
- You can use system calls to control the execution characteristics of parallel programs, as discussed under “Managing Applications” on page 4-2.
- You can open a file on all nodes at once very efficiently with `gopen()`, as discussed under “Opening Files in Parallel” on page 5-9.
- You can read the same data from a file into all nodes at the same time very efficiently with the I/O mode `M_GLOBAL`, as discussed under “Using I/O Modes” on page 5-13.
- You can read data into or write data from a series of scattered memory buffers with the `...readv()` and `...writev()` calls, as discussed under “Reading and Writing Files in Parallel” on page 5-24.
- You can find out the characteristics of PFS file systems (which are more configurable than CFS) with the `getpfsinfo()` and `statpfs()` calls, as discussed under “Getting Information About PFS File Systems” on page 5-39.
- You can use the HIPPI and FDDI network interfaces, as discussed in the *Paragon™ High-Performance Parallel Interface Manual* and *Paragon™ Fiber Distributed Data Interface Installation and Configuration Guide*.
- You can use the Paragon application development tools to help you port and optimize your code, as discussed in the *Paragon™ Application Tools User's Guide*.

## Compilers

The Paragon OSF/1 compilers work the same as the iPSC system compilers, with the following exceptions:

- The compilers, linker, and other tools are now available on the Paragon system as well as on workstations. They can be called by the standard names (**cc**, **f77**, **ld**, and so on) as well as the names used in cross-development (**icc**, **if77**, **ld860**, and so on).
- The environment variable that specifies the root of the compiler directory tree is called **PARAGON\_XDEV** rather than **IPSC\_XDEV**. The default for this variable is now **/usr/paragon/XDEV** rather than **/usr/ipsc/XDEV**.
- The compiler files are now found in the directory **\$PARAGON\_XDEV/paragon** rather than **\$IPSC\_XDEV/i860**. For example, your execution search path (**path** or **PATH** environment variable) should include the directory **\$PARAGON\_XDEV/paragon/bin.arch** (where **arch** identifies the architecture of the system, such as **paragon** or **sun4**) rather than **\$IPSC\_XDEV/i860/bin.arch** or **\$IPSC\_XDEV/i860/bin**.
- The **-p** switch is now ignored. See the *Paragon™ Application Tools User's Guide* for information on profiling.
- The default for quad-alignment has been changed from **-Mnoquad** to **-Mquad**. This change results in up to four times better performance for some code.
- The new switch **-nx** has been added. This switch generates a program that automatically starts itself on multiple nodes, as discussed under "Compiling and Linking Applications" on page 2-5. The switch **-node** is currently accepted as a synonym for **-nx**, but this support may be dropped in a future release.
- You can now have a file called **.icfrc** in your home directory that defines the default compiler switches for you.

See the *Paragon™ Fortran Compiler User's Guide* or *Paragon™ C Compiler User's Guide* for more information on the Paragon OSF/1 compilers.

### NOTE

You cannot use the Paragon OSF/1 cross-compilers to produce programs for the iPSC system, and you cannot use the iPSC system cross-compilers to produce programs for Paragon OSF/1.

If you develop programs for the iPSC system as well for Paragon OSF/1, you must be sure that your execution search path (**PATH** or **path** variable) is set appropriately for your current target system. To compile a program for Paragon OSF/1, the variable **PARAGON\_XDEV** must be set appropriately

and your execution search path must include `$PARAGON_XDEV/paragon/bin.arch`; to compile a program for the iPSC system, the variable `IPSC_XDEV` must be set appropriately and your execution search path must include `$IPSC_XDEV/i860/bin.arch` instead. Be sure that your execution search path does not include both these directories at the same time.

## Commands

In general, all of the standard commands of UNIX System V are supported by Paragon OSF/1, but none of the iPSC-system-specific commands are supported. However, many of these commands are not needed in Paragon OSF/1, or have equivalent standard commands in OSF/1.

## Cube Control Commands

The usage model of Paragon OSF/1 is different from that of the iPSC system. Instead of allocating a cube with a certain number of nodes, loading a program onto the cube, and then releasing the cube, you run a parallel application simply by typing its name on the Paragon OSF/1 command line. You can use command-line arguments to control its execution characteristics (such as the number of nodes on which it runs), and you can use standard OSF/1 process control commands such as **kill** to control the program. (See Chapter 2 for more information on running and controlling applications in Paragon OSF/1.)

For this reason, the following iPSC system commands, which create and control cubes, are not supported in Paragon OSF/1:

- |                   |   |
|-------------------|---|
| <b>archcube</b>   | This command is not needed in Paragon OSF/1 because all nodes currently have the same architecture.   |
| <b>attachcube</b> | This command is not needed in Paragon OSF/1 because you do not have to attach to a cube before you can use it.  |
| <b>cubeinfo</b>   | Use the <b>lspart</b> command to list the available partitions. See "Listing Subpartitions" on page 2-49 for more information.  |
| <b>getcube</b>    | Use the <b>-sz</b> switch on the application command line to specify the number of nodes allocated to the application. See "Specifying Application Size" on page 2-15 for more information. |

The **mkpart** command is similar to **getcube** in that it allocates a partition (a group of nodes). However, partitions are not the same as cubes: partitions can overlap, and a partition can be used by several applications at once. Depending on the policies of your site, you may or may not be allowed to allocate partitions. See "Making Partitions" on page 2-39 for more information.

<b>killcube</b>	Use the OSF/1 <b>kill</b> command to kill a running application, or press your interrupt key (<Ctrl-c> or <Del>). See "Managing Running Applications" on page 2-23 for more information.
<b>load</b>	Type an application's filename on the command line to run it on multiple nodes. See "Running Applications" on page 2-11 for more information.
<b>newserver</b>	This command is not needed in Paragon OSF/1 because you can use the usual OSF/1 I/O redirection characters to redirect an application's output. See "I/O Redirection" on page 2-12 for more information.
<b>relcube</b>	This command is not needed in Paragon OSF/1 because you do not have to release a cube once you have used it. The nodes allocated to an application are automatically released when all the processes in the application have terminated.  The <b>rmupart</b> command is similar to <b>relcube</b> in that it deallocates a partition (a group of nodes). However, partitions are not the same as cubes: partitions can overlap, and a partition can be used by several applications at once. Depending on the policies of your site, you may or may not be allowed to remove partitions. See "Removing Partitions" on page 2-45 for more information.
<b>startcube</b>	This command has no equivalent in Paragon OSF/1. There is no way to load an application into the nodes' memory without starting it.
<b>syslog</b>	This command is not needed in Paragon OSF/1 because you can use the usual OSF/1 I/O redirection characters to redirect an application's output. The standard I/O of a node process is connected to the same files or devices as the standard I/O of its controlling process. See "I/O Redirection" on page 2-12 for more information.
<b>waitcube</b>	This command is not needed in Paragon OSF/1 because, by default, your command prompt does not return until the application has completed. Also, you can redirect the output of any program with the usual OSF/1 I/O redirection characters (see "I/O Redirection" on page 2-12 for more information).



## CFS Commands

The following iPSC system commands, which control the Concurrent File System and the SRM tape drive, are not supported in Paragon OSF/1:

<b>cptape</b>	Use the <b>cpio</b> command instead. See <b>cpio</b> in the <i>OSF/1 Command Reference</i> for more information.
<b>showvol</b>	Use the <b>showfs</b> command instead. See “Displaying File System Attributes” on page 5-5 for more information.
<b>star</b>	Use the <b>tar</b> command instead. (Note that you must use the <b>-E</b> switch to archive a file larger than 2G–1 bytes.) See <b>tar</b> in the <i>Paragon™ Commands Reference Manual</i> and <i>OSF/1 Command Reference</i> for more information.
<b>stream</b>	This command is not needed in Paragon OSF/1 because there is no streaming tape drive.
<b>tapemode</b>	This command currently has no equivalent in Paragon OSF/1. There is no way to display or change the operating mode of the system's tape drives.

## System Administration Commands

The following iPSC system commands, which are used for system administration, are not supported in Paragon OSF/1:

<b>cbackup</b>	Use the <b>dump</b> command instead. See <b>dump</b> in the <i>OSF/1 System and Network Administrator's Reference</i> for more information.
<b>cfschk</b>	Use the <b>fsck</b> command instead. See <b>fsck</b> in the <i>OSF/1 System and Network Administrator's Reference</i> for more information.
<b>crestore</b>	Use the <b>rdump</b> command instead. See <b>rdump</b> in the <i>OSF/1 System and Network Administrator's Reference</i> for more information.
<b>makewhatis</b>	Use the <b>catman</b> command instead. See <b>catman</b> in the <i>OSF/1 Command Reference</i> for more information.
<b>mkcfs</b>	Use the <b>newfs</b> command instead. See <b>newfs</b> in the <i>OSF/1 System and Network Administrator's Reference</i> for more information.

**mkdev** Use the **mknod** command instead. See **mknod** in the *OSF/1 System and Network Administrator's Reference* for more information.

**plogon and plogoff**

These commands currently have no equivalent in Paragon OSF/1. There is currently no way to log creation and deletion of partitions or running of applications. However, you can use the **syslogd** daemon to log other system activity. See **syslogd** in the *OSF/1 System and Network Administrator's Reference* for more information.

## Remote Host Commands

The following iPSC system commands, which are used for program development on remote hosts, are not supported in Paragon OSF/1:

**rf77** Use the **if77** command instead. See the *Paragon™ Fortran Compiler User's Guide* for more information.

**rcc** Use the **icc** command instead. See the *Paragon™ C Compiler User's Guide* for more information.

**rld** Use the **ld860** command instead. See the *Paragon™ Fortran Compiler User's Guide* or *Paragon™ C Compiler User's Guide* for more information.

**ras** Use the **as860** command instead. See the *Paragon™ Fortran Compiler User's Guide* or *Paragon™ C Compiler User's Guide* for more information.

**rar** Use the **ar860** command instead. See the *Paragon™ Fortran Compiler User's Guide* or *Paragon™ C Compiler User's Guide* for more information.

## Miscellaneous Commands

The following iPSC system commands are not supported in Paragon OSF/1:

**less** Use the **more** command instead. See **more** in the *OSF/1 Command Reference* for more information.

**manpath** Use the **MANPATH** environment variable instead. See **man** in the *OSF/1 Command Reference* for more information.

<b>nsh</b>	Use the <b>rlogin</b> or <b>telnet</b> command to log into the Paragon system from your workstation. See <b>rlogin</b> or <b>telnet</b> in your workstation's documentation for more information.
<b>rebootcube</b>	This command has no equivalent in Paragon OSF/1. There is no way for ordinary users to reboot the system.

## System Calls

In general, all of the standard system calls of UNIX System V and most of the iPSC-system-specific system calls are supported by Paragon OSF/1. This section suggests alternatives for the unsupported calls.

### NOTE

Some iPSC calls are provided for backward compatibility only, and are not intended for use in new programs. These calls are not documented in the online manpages or in the *Paragon™ C System Calls Reference Manual* or *Paragon™ Fortran System Calls Reference Manual*. See "iPSC® and Touchstone DELTA Compatibility Calls" on page 4-52 for a list of these calls.

## Include Files

Paragon OSF/1 does not support the iPSC system include files `<cube.h>` or `<fcube.h>`. You should replace any reference to `<cube.h>` with `<nx.h>`, and any reference to `<fcube.h>` with `<fnx.h>`.

## Host Calls

Applications in Paragon OSF/1 do not usually have host programs. The usual programming model in Paragon OSF/1 is to write a single program (which corresponds to a "node program" in the iPSC system), link it with `-nx`, and execute the program on a group of nodes by typing its name (see "Running Applications" on page 2-11 for more information). You may be able to eliminate all references to the following unsupported calls by rewriting your program to use this programming model. If your application requires a separate host program, you can rewrite your host program into a *controlling process* (see "Managing Applications" on page 4-2 for more information).

For this reason, the **-host** switch to the **cc** and **f77** commands is not supported (there is no separate host library; host programs use the same library as node programs). Also, the following iPSC system calls, which are used in host programs, are not supported in Paragon OSF/1:

**attachcube()** This call currently has no equivalent in Paragon OSF/1. Unlike a host program, a controlling process cannot be associated with more than one application. Consider re-writing your host program as two or more separate programs, each of which creates one application and communicates with the other host program(s) using pipes, signals, or some other OSF/1 interprocess communication method. See “Managing Applications” on page 4-2 for information on creating and controlling applications using system calls.

**cubeinfo()** This call currently has no equivalent in Paragon OSF/1. However, because allocation of nodes in Paragon OSF/1 is not exclusive, it is not usually necessary for programs to know how other users have allocated nodes. To get information on your own application (equivalent to the “current cube”), you can use calls such as **numnodes()**.

**getcube()** Use **nx\_initve()** instead. See “Managing Applications” on page 4-2 for information on **nx\_initve()**.

**killcube()** This call is supported, but can only be used to kill and flush all processes on all nodes (**killcube(-1,-1)**).

You can use **kill()** to kill a single process, as discussed for **killproc()** below.

**killproc()** This call is supported, but can only be used to kill all processes on all nodes (**killproc(-1,-1)**).

You can use **kill()** to kill a single process, given its OSF/1 process ID. **kill()** is supported in both C and Fortran. To determine the OSF/1 process ID of a process created by **nx\_nfork()**, **nx\_load()**, or **nx\_loadve()**, use the values stored into the *pid\_array* argument. These calls store the OSF/1 PIDs of the processes created into the elements of this array, as discussed under “Using PIDs” on page 4-14.

For example, to kill the process on node number *node*:

```
#include <signal.h>

n = nx_nfork(NULL, -1, ptype, pid_array);
.
.
.
kill(pid_array[node], SIGKILL);
```

Note that process types (*ptype* in this example) in Paragon OSF/1 are equivalent to NX PIDs in the iPSC system. PIDs (*pid\_array* in this example) in Paragon OSF/1 are standard UNIX process IDs.

See the *OSF/1 Programmer's Reference* for information on `kill()`; see "Managing Applications" on page 4-2 for information on `nx_nfork()`, `nx_load()`, and `nx_loadve()`.

**killsyslog()** Use `freopen()` instead, to close the standard output and standard error output and reopen them to `/dev/tty`. See `freopen()` in the *OSF/1 Programmer's Reference* for more information.

`freopen()` is not currently supported for Fortran programs. However, it is supported for C programs. You can write a C "wrapper" function, as follows:

```
#include <stdio.h>

void killsyslog_() {
    freopen("/dev/tty", "w", stdout);
    freopen("/dev/tty", "w", stderr);
}
```

Note the underscore at the end of the function name. Once you have compiled this function and linked it into your Fortran program, you can call `killsyslog()` as described in the iPSC system documentation.

**newserver()** This call is not necessary in Paragon OSF/1. The standard I/O of a controlling process (host process) is connected to the same files or devices as the standard I/O of its node processes.

**relcube()** This call is not necessary in Paragon OSF/1. The nodes allocated to an application are automatically released when all the processes in the application have terminated.

**setpid()** Use `setptype()` instead. "Process Characteristics" on page 3-3 for information on `setptype()`, and "Message Passing Between Controlling Process and Application Processes" on page 4-25 for information on using `setptype()` in a controlling process.

**setsyslog()** This call is not necessary in Paragon OSF/1. The standard I/O of a controlling process (host process) is connected to the same files or devices as the standard I/O of its node processes.

**waitall()**

To wait for all processes on all nodes (**waitall(-1, -1)**), call **nx\_waitall()**. See “Waiting for Application Processes with **nx\_waitall()**” on page 4-14 for more information.

To wait for a single node process (**waitall(node, pid)**), use the OSF/1 system call **waitpid()** to wait for the process with a particular OSF/1 process ID. To determine the PID of a process created by **nx\_nfork()**, **nx\_load()**, or **nx\_loadve()**, use the values stored into the *pid\_array* argument. These calls store the OSF/1 PIDs of the processes created into the elements of this array, as discussed under “Using PIDs” on page 4-14.

For example, to wait for the process on node number *node*:

```
n = nx_nfork(NULL, -1, ptype, pid_array);
.
.
.
waitpid(pid_array[node], &status, 0);
```

Note that process types (*ptype* in this example) in Paragon OSF/1 are equivalent to NX PIDs in the iPSC system. PIDs (*pid\_array* in this example) in Paragon OSF/1 are standard UNIX process IDs.

See the *OSF/1 Programmer's Reference* for information on **wait()** and **waitpid()**; see “Managing Applications” on page 4-2 for information on **nx\_nfork()**, **nx\_load()**, and **nx\_loadve()**.

**wait()** is supported in both C and Fortran, but **waitpid()** is not currently supported in Fortran. You can make **waitpid()** callable from Fortran by writing a C “wrapper” function, as follows:

```
#include <sys/types.h>
#include <sys/wait.h>

int waitpid_(int *process_id,
             int *status_location,
             int *options) {
    return((int)waitpid((pid_t)*process_id,
                       status_location,
                       *options);
}
```

Note the underscore at the end of the function name. Once you have compiled this file and linked it into your Fortran program, you can call **waitpid()** as described in the *OSF/1 Programmer's Reference*. The wrapper function **waitpid()** takes three **integer\*4** parameters and returns an **integer\*4** value.

**waitone()**

To wait for the first node process in the entire application to complete (**waitone(-1, -1, cnode, cpid, ccode)**), use the OSF/1 system call **wait()**. For example:

```
n = nx_nfork(nodes, NUMNODES, ptype, pids);
  .
  .
  .
pid = wait(&status);
```

After this call, the status of the first process to complete is stored in *status* and its OSF/1 process ID is stored in *pid*. To determine the process's node number, look for the value of *pid* in the *pids* array returned by **nx\_nfork()**, **nx\_load()**, or **nx\_loadve()**.

To wait for a single node process (**waitone(node, pid, cnode, cpid, ccode)**), use the same technique described for **waitall(node, pid)**:

```
n = nx_nfork(NULL, -1, ptype, pid_array);
  .
  .
  .
pid = waitpid(pid_array[node], &status, 0);
```

In this case, the status of the process is stored in *status* and its OSF/1 process ID is stored in *pid*. To determine the process's node number, look for the value of *pid* in *pid\_array* as described above.

See the *OSF/1 Programmer's Reference* for information on **wait()** and **waitpid()**; see "Managing Applications" on page 4-2 for information on **nx\_nfork()**, **nx\_load()**, and **nx\_loadve()**. **wait()** is supported in both C and Fortran, but **waitpid()** is not; to call **waitpid()** from Fortran, use the technique discussed previously under **waitall()**.

## Byte-Swapping Calls

The calls listed in Table B-1, which swap bytes between the format used on the cube and the format used on some remote hosts, are not supported in the current release of Paragon OSF/1.

**Table B-1. Unsupported iPSC® System Byte-Swapping Calls**

<b>createstruc()</b>	<b>CTOHF()</b>	<b>HTOCC()</b>	<b>HTOCL()</b>
<b>CTOHC()</b>	<b>CTOHL()</b>	<b>HTOCD()</b>	<b>HTOCS()</b>
<b>CTOHD()</b>	<b>CTOHS()</b>	<b>HTOCF()</b>	<b>relstruc()</b>

You can use the standard OSF/1 system calls **htonl()**, **htons()**, **ntohl()**, and **ntohs()** to swap bytes between the standard format for your machine and the Internet network format. See **htonl()**, **htons()**, **ntohl()**, and **ntohs()** in the *OSF/1 Programmer's Reference* for more information.

**htonl()**, **htons()**, **ntohl()**, and **ntohs()** are not currently supported for Fortran programs. However, they are supported for C programs. You can make them callable from Fortran by writing C “wrapper” functions, as follows:

```
#include <netinet/in.h>

long htonl_(long *hostlong) {
    return((long)htonl((unsigned long)*hostlong));
}

short htons_(short *hostshort) {
    return((short)htons((unsigned short)*hostshort));
}

long ntohl_(long *netlong) {
    return((long)ntohl((unsigned long)*netlong));
}

short ntohs_(short *netshort) {
    return((short)ntohs((unsigned short)*netshort));
}
```

Note the underscore at the end of each function name. Once you have compiled this file and linked it into your Fortran program, you can call **htonl()**, **htons()**, **ntohl()**, and **ntohs()** as described in the *OSF/1 Programmer's Reference*. The wrapper functions **htonl()** and **ntohl()** take an **integer\*4** parameter and return an **integer\*4** value; the wrapper functions **htons()** and **ntohs()** take an **integer\*2** parameter and return an **integer\*2** value.



## Floating-Point Control Calls

The Paragon OSF/1 C system calls `fpgetsticky()` and `fpsetsticky()`, which get and set the i860 microprocessor's *floating-point exception sticky flags*, and `fpgetmask()` and `fpsetmask()`, which get and set the *floating-point exception mask*, do not support the exception value `FP_X_DNML`, which represents a denormalization exception in the iPSC system.

The Paragon OSF/1 Fortran system call `fpsetmask()` also does not support the denormalization exception, and uses different numeric values to represent the various exceptions than the corresponding iPSC system call. See "Controlling Floating-Point Behavior" on page 4-46 for the correct values for Paragon OSF/1.

## CFS Calls

In Paragon OSF/1, the iPSC system's Concurrent File System (CFS) has been replaced by the Parallel File System (PFS). PFS calls are compatible with CFS calls; however, PFS offers additional functionality (see Chapter 5 for more information). This section lists the differences that may affect some programs that use CFS calls.

### `iread()` and `iwrite()`

These calls work in Paragon OSF/1 just as they do in the iPSC system. In both systems, the number of I/O IDs is limited; however, the limit in Paragon OSF/1 is much smaller than in the iPSC system. (In the iPSC system the limit is 5000; in Paragon OSF/1 it is at least 256, but may vary from release to release.) For this reason, it is very important that you use `iodone()` or `iowait()` to release each ID as soon as possible after you use it. If you program in C, you can use `_iread()` or `_iwrite()` to detect the "too many requests" error (`EQNOMID`).

### `open()`

Many iPSC system programs use code like the following to open a file on all nodes:

```
if(mynode() == 0) {
    fd = open("myfile", O_CREAT | O_RDWR, 0644);
    gsync();
} else {
    gsync();
    fd = open("myfile", O_RDWR, 0644);
}
setiomode(fd, iomode);
```

The `open()` call works the same in Paragon OSF/1 as it does in the iPSC system. However, if this code is executed on many nodes, the large number of `open()` requests arriving simultaneously at the I/O node can cause the I/O node to slow down, hang, or even crash. This can even cause the system to crash.

You should always use the **gopen()** call instead of this type of code. For example, you should replace the lines shown above with the following:

```
fd = gopen("myfile", O_CREAT | O_RDWR,
          iomode, 0644);
```

**gopen()** opens a file simultaneously on all nodes and sets its I/O mode in a single operation. It is much more efficient than having each node call **open()**, and avoids this type of system crash completely.

Note that **gopen()** opens the same file on each node. If each node is opening its own file, you must still use **open()**. However, you should try to avoid using **open()** together with **gsync()**, to prevent all the **open()** requests from arriving at the I/O node at the same time.

## Miscellaneous Calls

The following iPSC system calls work differently or are not supported in Paragon OSF/1:

- dclock()** This call works in Paragon OSF/1 just as it does in the iPSC system: it returns the time since the system was booted, in seconds. However, in a gang-scheduled partition your application may be *rolled out* and then *rolled in* again. While it is rolled out, the application is stopped but the **dclock()** clock keeps going (reflecting "wall-clock" time), which means that **dclock()** cannot be used to determine the amount of time your application has actually been running.
- Using **dclock()** in a gang-scheduled partition may result in incorrect MFLOPS estimates. You can use the **time** command, **getrusage()** system call (C only), or the **etime()** or **dtime()** routine (Fortran only) instead to determine your application's CPU usage. See the *OSF/1 Command Reference* for information on **time**, the *OSF/1 Programmer's Reference* for information on **getrusage()**, and the *Paragon™ Fortran Compiler User's Guide* for information on **etime()** and **dtime()**.
- flushmsg()** This call currently has no equivalent in Paragon OSF/1. It may be supported in a future release.
- getiphosts()** This call currently has no equivalent in Paragon OSF/1. However, because the OSF/1 operating system automatically routes network traffic using all available Ethernet ports, it is not usually necessary to know the network names of the available ports.
- gixor()** This call is not supported in Paragon OSF/1. The exclusive OR operator is not associative, and gives unpredictable results when used on more than two nodes.

- glxor()** This call is not supported in Paragon OSF/1. The exclusive OR operator is not associative, and gives unpredictable results when used on more than two nodes.
- handler()** Use the **signal()** system call instead (**signal()** is supported for both C and Fortran). See **signal()** in the *OSF/1 Programmer's Reference* for information on signal handling; see **signal()** in the *Paragon™ Fortran Compiler User's Guide* for information on the Fortran interface to **signal()**.
- plogon()** and **plogoff()** These calls currently have no equivalent in Paragon OSF/1. There is currently no way to automatically log creation and deletion of partitions or running of applications. However, you can use the **syslog()** call to log activities under program control. See **syslog()** in the *OSF/1 Programmer's Reference* for more information.
- setiphost()** This call is not necessary in Paragon OSF/1. The OSF/1 operating system automatically routes network traffic using all available Ethernet ports; it is not necessary to select one port to perform network operations.
- setpgrp()** There are two different versions of this call in Paragon OSF/1. The standard version of **setpgrp()**, found in *libbsd.a*, is equivalent to **setpgid()** and is not compatible with the iPSC/860 version. The System V version of **setpgrp()**, found in *libsys5.a*, is equivalent to **setsid()** and is compatible with the iPSC/860 version. To get the iPSC/860-compatible version, be sure to use the switch **-lsys5** when linking.

## Summary

Table B-2 summarizes the Paragon OSF/1 equivalents for the unsupported iPSC system commands.

Table B-2. Summary of Unsupported iPSC® System Commands (1 of 2)

iPSC® System Command	Paragon™ OSF/1 Equivalent
<b>archcube</b>	(none)
<b>attachcube</b>	(none)
<b>cbackup</b>	<b>dump</b>
<b>cfschk</b>	<b>fsck</b>
<b>cptape</b>	<b>cpio</b>
<b>crestore</b>	<b>rdump</b>
<b>cubeinfo</b>	<b>lspart</b>

**Table B-2. Summary of Unsupported iPSC® System Commands (2 of 2)**

<b>iPSC® System Command</b>	<b>Paragon™ OSF/1 Equivalent</b>
<b>getcube</b>	-sz switch on application command line
<b>killcube</b>	<b>kill</b>
<b>less</b>	<b>more</b>
<b>load</b>	Application's filename
<b>makewhatis</b>	<b>catman</b>
<b>manpath</b>	<i>MANPATH</i> environment variable
<b>mkcfs</b>	<b>newfs</b>
<b>mkdev</b>	<b>mknod</b>
<b>newserver</b>	I/O redirection characters
<b>nsh</b>	<b>rlogin or telnet</b>
<b>plogoff</b>	(none)
<b>plogon</b>	(none)
<b>rar</b>	<b>ar860</b>
<b>ras</b>	<b>as860</b>
<b>rcc</b>	<b>icc</b>
<b>rebootcube</b>	(none)
<b>relcube</b>	(none)
<b>rf77</b>	<b>if77</b>
<b>rld</b>	<b>ld860</b>
<b>showvol</b>	<b>showfs</b>
<b>star</b>	<b>tar</b>
<b>startcube</b>	(none)
<b>stream</b>	(none)
<b>syslog</b>	I/O redirection characters
<b>tapemode</b>	(none)
<b>waitcube</b>	(none)

Table B-3 summarizes the Paragon OSF/1 equivalents for the unsupported iPSC system calls.

**Table B-3. Summary of Unsupported iPSC® System Calls (1 of 2)**

<b>iPSC® System Call</b>	<b>Paragon™ OSF/1 Equivalent</b>
<b>attachcube()</b>	(none)
<b>cubeinfo()</b>	(none)
<b>dclock()</b>	Supported, but use <b>getrusage()</b> (C) or <b>etime()/dtime()</b> (Fortran) to determine CPU usage in gang-scheduled partitions.
<b>fpgetsticky(), fpsetsticky(), fpgetmask(), fpsetmask()</b>	Supported, except for <b>FP_X_DNML</b> , but Fortran mask values are different.
<b>flushmsg()</b>	(none)
<b>getcube()</b>	<b>nx_initve()</b>
<b>getiphosts()</b>	(none)
<b>gixor()</b>	(none)
<b>glxor()</b>	(none)
<b>handler()</b>	<b>signal()</b>
<b>iread(), iwrite()</b>	Supported, but number of I/O IDs is much smaller.
<b>killcube()</b>	Use <b>killcube(-1,-1)</b> to kill and flush all processes; use <b>kill()</b> to kill one process
<b>killsyslog()</b>	(none)
<b>killproc()</b>	Use <b>killproc(-1,-1)</b> to kill all processes; use <b>kill()</b> to kill one process
<b>newserver()</b>	<b>freopen()</b>
<b>open()</b>	Supported, but use <b>gopen()</b> instead if possible.
<b>plogoff()</b>	(none)
<b>plogon()</b>	(none)
<b>relcube()</b>	(none)
<b>setiphost()</b>	(none)
<b>setpgrp()</b>	Supported, but be sure to link with <b>-lsys5</b> to get the correct version.

**Table B-3. Summary of Unsupported iPSC® System Calls (2 of 2)**

<b>iPSC® System Call</b>	<b>Paragon™ OSF/1 Equivalent</b>
<b>setpid()</b>	<b>setpctype()</b>
<b>setsyslog()</b>	(none)
<b>waitall()</b>	Use <b>nx_waitall()</b> to wait for all processes; use <b>wait()</b> or <b>waitpid()</b> to wait for one process
<b>waitone()</b>	<b>wait()</b> or <b>waitpid()</b>
Byte-swapping calls	<b>htonl(), htons(), ntohl(), and ntohs()</b>

# Index

## Symbols

- ### in filenames 5-31
- . (dot) in partition pathnames 2-28
- . (root) partition 2-26
- .compute partition 2-28
- .F extension 2-9
- .service partition 2-28
- /cfs directory 5-13
- /dev/io0/rmt6 device 5-46
- /pfs file system 5-5
- /usr/ccs/lib directory 2-8
- /usr/include directory 2-8
- /usr/lib directory 2-8
- /usr/paragon/XDEV directory 2-9
- /usr/tmp directory 5-13
- \;file (second program in an application) 2-21
- \_\_NODE preprocessor symbol 2-5
- \_cread() system call 5-26
- \_creadv() system call 5-26
- \_crecv() system call 4-42
- \_cwrite() system call 5-26
- \_cwritev() system call 5-26
- \_r calls 6-10
- \_REENTRANT preprocessor symbol 6-5, 6-41
- 1
  - as error return 4-42
  - as message type 3-6
  - as node number 3-3, 4-25
  - as process type 3-4
  - as sending node number 3-24
  - as sending process type 3-24
- 64-bit integers 5-37

## A

- absolute partition pathname 2-28
- access methods 5-13
  - synchronization of 5-48
- access() system call 5-31
- accessing contiguous memory locations 8-5
- active and inactive applications 2-35
- address space 1-5
- algorithms 8-1
- aligning application buffers 8-12
- aligning I/O buffers 8-25
- ALLOCATE statement 8-8, 8-15
- allocating memory 8-8

- allocating nodes to a partition 2-30, 2-40
- allocating nodes to an application 2-15, 4-4
- allocating space to a file 5-7, 5-32
- allocator configuration parameters 2-36
- alternate node topologies 7-6
- anonymous files 5-13
- applicable documents vii
- application buffers
  - aligning 8-12
- "application" command 2-13, A-2
- applications 1-1, 2-1
  - active and inactive 2-35
  - allocating nodes to an application 4-4
  - Bourne shell 2-24
  - compiling and linking 2-5
  - compiling, linking, and executing 2-3
  - contiguous nodes 2-16, 4-5, 4-28
  - control decomposition 7-5
  - controlling execution characteristics 2-13
  - controlling process 2-24, 2-28, 4-4, 4-21
  - controlling with system calls 4-2
  - creating and controlling 4-4
  - debugging 2-34
  - decomposition 7-3
  - default partition 2-14
  - designing 7-1
  - designing a communication strategy 7-6
  - domain decomposition 7-3
  - error handling 4-42
  - error messages 2-12
  - executing 2-11
  - gth switch 8-18
  - I/O redirection 2-12
  - independent of number of nodes 7-5
  - interactive 2-34
  - killing application processes 4-23
  - listing the applications in a partition 2-51
  - lnx compiler switch 2-11
  - load balancing 7-3
  - managing running applications 2-23
  - matrix\*vector example 7-11
  - mbf switch 8-16
  - mea switch 8-17
  - message buffers 3-14
  - message passing with controlling process 4-25
  - mex switch 8-19
  - noc switch 8-17
  - node numbers 3-3
  - nqueens example 7-13
  - nx compiler switch 2-11
  - on switch 2-19
  - order of switches 2-13
  - overlapping 2-36
  - partition of 2-14
  - perfectly-parallel 7-2



- performance improvement techniques 8-1
  - pi example 7-7
  - pkt switch 8-16
  - plk switch 8-3, 8-15
  - pn switch 2-22
  - pri switch 2-17
  - priority of 2-17, 2-35, 4-9
  - process type of 2-18, 3-4
  - pt switch 2-18, 3-4
  - rectangular dimensions of 4-16
  - rectangular size 2-16
  - removing partitions containing 2-45
  - running in a particular partition 2-22
  - running multiple programs 2-21
  - running on a subset of the nodes 2-18
  - sct switch 8-18
  - separating the user interface from the computation 7-3
  - shell scripts 2-24
  - size of 2-15
  - sth switch 8-18
  - sz switch 2-16, 3-3
  - waiting for application processes 4-14
  - arbitration between processes 2-33
  - arbitration mechanisms 6-1
  - archcube command B-5
  - architecture of your workstation 2-6
  - argc and argv parameters 4-5, 4-13
  - arithmetic, extended 5-37
  - arrays
    - accessing contiguous memory locations 8-5
    - dynamic allocation of 8-8
    - large 8-8
    - memory layout in C and Fortran 8-5
  - arrow LEDs 1-4
  - assembly language programming 8-1, 8-7
  - asynchronous and synchronous calls 8-10
  - asynchronous and synchronous I/O calls 8-24
  - asynchronous cancelability 6-29
  - asynchronous file I/O calls 5-27
  - asynchronous message-passing calls 3-6, 3-7, 3-10
    - with interrupt handler 3-7
  - attachcube command B-5
  - attachcube() system call B-10
  - attributes of file systems 5-5
  - attributes of pthreads 6-15
  - available nodes 2-48
  - avoiding virtual memory paging 8-3
- B**
- backward compatibility calls B-9
  - backward library references 2-10
  - bad nodes 2-31, 2-48, 4-40
  - balancing the load among the nodes 7-3
  - Basic Linear Algebra Subroutines (BLAS) 7-13
  - Basic Math Library (libkmath.a) 8-6
  - bg command 2-23
  - binary files in PFS file systems 5-4
  - BLAS library 8-6
  - blocking 3-8, 5-48
    - on child processes 4-14
    - with pthreads 6-37
  - blocks, file system 8-25
  - bold text vi
  - Bourne shell (sh) 2-24
    - with pthreads 6-39
  - brackets, in syntax descriptions vii
  - broadcast 3-9

- buffering
  - of Fortran I/O 5-30
  - of messages 3-14
  - of standard I/O 2-12
- buffers for I/O 8-25
- buffers for messages 8-11
  - aligning 8-12
- buffers, message 8-16
- bureaucracy in node programs 7-5
- bytes read or written 5-26
- byte-swapping calls B-14
  
- C**
- C library, reentrant 6-2
- C preprocessor on a Fortran program 2-9
- C programs
  - error handling 4-42, 8-7
  - file descriptors 5-8
  - including nx.h 2-8
  - memory access considerations 8-5
  - pointers to message buffers 8-13
  - structure padding 8-13
- cache lines 8-5, 8-12
- caches
  - data 8-5
  - instruction 8-6
- caching 8-5
- canceling pthreads 6-28
- cat command 5-5, 5-35
- cbackup command B-7
  
- cc command 2-4, 2-5, A-1
  - host switch B-10
  - l switch 2-9, 2-10
  - L switch 2-9, 2-10
  - Mquad switch B-4
  - node switch B-4
  - nx switch 2-5, B-4
  - order of switches 2-10
  - p switch B-4
- CFS 5-1, B-7
- /cfs directory 5-13
- CFS\_MOUNT environment variable 5-13
- cfschk command B-7
- changing partition characteristics 2-54, 4-36
- changing process type 3-4
- characteristics of a partition 2-29
  - default 2-29
- characteristics of processes 3-3
- chdir() system call 5-31, 8-9
  - with pthreads 6-9
- chess example 7-13
- chgrp command 5-35
- child partitions 2-29, 2-34
  - creating 2-39
  - listing 2-49
  - removing 2-45
  - removing partitions containing 2-45
- child processes 4-10, 4-11, 4-14
- chmod command 5-35
- chmod() system call 5-31, 8-9
- chown command 5-35
- chown() system call 5-31
- chpart command 2-29, 2-54, A-2
- CLASSPACK 8-6
- cleanup routines for pthreads 6-32

- clock, global 4-50
- close() system call 5-28
  - synchronization 5-48
- closing parallel files 5-15, 5-28
- code segment 8-15
- commands 2-1
  - compiling and linking applications 2-5
  - cross-development 2-5
  - executing applications 2-11
  - iPSC system compatibility B-5
  - managing partitions 2-25
  - managing running applications 2-23
  - native 2-5
  - on the Intel supercomputer 2-2
  - on workstations 2-2
  - PFS 5-5
  - summary A-1
- commons in message passing 3-17
- communication
  - overlapping with computation 8-10
- compatibility with the iPSC system B-1
- compiler switches 2-5, 2-8, 2-10, B-4
- compilers, iPSC system 2-7
  - compatibility with B-4
- compiling and linking
  - optimization 8-3
- compiling and linking applications 1-5, 2-5
  - host switch B-10
  - lnx switch 2-6
    - order of 2-10
  - Mquad switch B-4
  - node switch B-4
  - nx switch 2-5, B-4
  - p switch B-4
  - quick example 2-3
  - specifying include file pathnames 2-8
  - specifying library pathnames 2-8
  - tips 2-8
  - with pthreads 6-5
- complete (synchronous) system calls
  - file I/O 5-24
  - message passing 3-7
- compress command 5-35
- computation
  - overlapping with communication 8-10
- computational kernel of an application 7-3
- compute nodes 1-2
- compute partition 2-2, 2-27, 2-28
- Concurrent File System 5-1, B-7
- condition variables 6-21
- configuration parameters of the allocator 2-36
- configuring message passing 8-18
- configuring your environment
  - for cross-development 2-6
  - for online manual pages 2-6
- contiguous memory locations 8-5
- contiguous nodes 2-16, 2-40, 4-5, 4-28
- contiguous partitions 2-31
- control decomposition 7-5
  - example 7-14
- controlling application execution with system calls
  - 4-2
- controlling process 2-24, 2-28, 2-51, 4-4, 4-21
  - global operations 4-25
  - message passing 4-25
  - node number of 4-25
  - process type of 4-25
- controlling tape devices 5-44
- controlling the application's execution
  - characteristics 2-13
- coprocessor 8-10
- copying processes onto nodes 4-10
- core dumps 4-44

- core files in PFS file systems 5-4
  - correspondents parameter 8-17
  - count parameter 3-5
  - Courier font vi
  - cp command 5-5, 5-35
  - cpio command 5-35
  - cprobe() system call 3-14, A-7, A-31
  - cprobex() system call 3-16, 3-24, A-9, A-34
  - cptape command B-7
  - cread() system call 5-15, 5-25, A-19, A-45
    - synchronization 5-48
  - creadv() system call 5-25, 5-48, A-19, A-45
  - creat() system call 5-31
  - createstruc() system call B-14
  - creating an application 4-4
  - creating partitions 2-39, 4-28
  - crecv() system call 3-8, A-5, A-29
    - message buffering 8-11
  - crecvx() system call 3-16, 3-24, A-9, A-33
  - crestore command B-7
  - critical code 3-22
  - cross-compilers 2-2
  - cross-development facility 1-6
    - commands 2-5
    - configuring your environment 2-6
  - csend() system call 3-8, A-5, A-29
    - message buffering 8-11
  - csendrecv() system call 3-8, A-5, A-29
  - cthreads 6-2
  - CTOH...() system calls B-14
  - <Ctrl-c> key 2-23
  - <Ctrl-z> key 2-23
  - cube control commands B-5
  - cube.h file B-9
  - cubeinfo command B-5
  - cubeinfo() system call B-10
  - current partition 2-28
  - cwrite() system call 5-15, 5-24, 5-25, 5-31, A-19, A-45
    - synchronization 5-48
  - cwritev() system call 5-25, 5-48, A-19, A-46
- ## D
- data cache 8-5
  - data locality 8-10
  - data segment 8-15
  - dclock() system call 4-50, A-16, A-43, B-16
  - dead nodes 2-31, 2-48, 4-40
  - deadlock 8-17
  - dealing out data to the nodes 7-3
  - debugging applications 2-34
  - declaring large arrays 8-8
  - decomposition 7-3
    - control decomposition 7-5
    - domain decomposition 7-3
  - default application size 2-17
  - default characteristics of a partition 2-29
  - default partition 2-2, 2-14
    - determining 2-15
    - listing applications in 2-51
    - listing subpartitions of 2-49
    - setting 2-14
    - showing characteristics of 2-46
  - <De> key 2-23
  - DELTA System 4-52

- designing a communication strategy 7-6
  - designing a parallel application 7-1
  - destroying partitions 2-45, 4-30
  - detecting end-of-file 5-29
  - determining your default partition 2-15
  - /dev/io0/rmt6 device 5-46
  - devices, disk 5-2
  - df command 5-7
  - diff command 5-35
  - differences between iPSC and Paragon B-1
  - disk space allocated to a file 5-7, 5-32
  - disks and file systems 5-2
  - displaying file system attributes 5-5
  - distributed memory 7-2
  - distributing computation among the nodes 7-3
  - distributing data among the nodes 7-3
  - documents, related vii
  - domain decomposition 7-3
    - example 7-8
  - dot (.) in partition pathnames 2-28
  - dot (.) partition (root partition) 2-26
  - du command 5-35
  - dumping core 4-44
  - dynamic algorithm selection 3-27
  - dynamic memory allocation 8-8
- E**
- eadd() system call 5-37, A-22, A-49
  - ecmp() system call 5-37, 5-38, A-22, A-49
  - ed command 5-35
  - ediv() system call 5-37, 5-38, A-22, A-49
  - effective priority limit 2-35, 2-44, 2-56
  - efficiency of PFS files 8-23
  - ellipses (...), in syntax descriptions vii
  - emod() system call 5-37, 5-38, A-22, A-49
  - emul() system call 5-37, A-22, A-49
  - end-of-file 5-29
  - environment variables
    - CFS\_MOUNT 5-13
    - for compiling and linking 2-6
    - for online manual pages 2-6
    - IPSC\_XDEV B-4
    - MANPATH 2-6
    - NX\_DFLT\_PART 2-3, 2-12, 2-14
    - NX\_DFLT\_SIZE 2-12, 2-17, 2-22
    - of child processes 4-13
    - PARAGON\_XDEV 2-6, 2-9
    - PATH 2-6
  - envp parameter 4-13
  - errno variable 4-42
    - with pthreads 6-41
  - "error 216 occurred, unknown" error 2-12
  - error handling 4-42, 8-7
    - in parallel file I/O calls 5-26
    - with pthreads 6-41
  - error messages 2-12
  - eseek() system call 5-36, A-21, A-48
    - synchronization 5-48
  - esize() system call 5-36, A-21, A-48
  - esize\_t structure 5-37
  - estat structure 5-36
  - estat() system call 5-36, A-21
  - estat.h file 5-37
  - estatfs structure 5-41
  - esub() system call 5-37, A-22, A-49

- Ethernet interface 1-2
  - etos() system call 5-37, A-22, A-49
  - ex command 5-35
  - example of compiling and linking 2-3
  - examples
    - iomodes 5-17
    - matrix\*vector 7-11
    - nqueens 7-13
    - pi 7-7
    - pthread 6-18
    - triangle 7-18
  - "exceeded allocator configuration parameters" error 2-36
  - "exceeds partition resources" error 2-12, 2-16
  - exception mask 4-48
  - exceptions 5-31
  - exec() system call 3-5
    - with pthread 6-9
  - executable files in PFS file systems 5-4
  - execute (x) permission on a partition 2-33
  - executing applications 2-3, 2-11
    - after cross-compilation 1-6
    - controlling 2-13
  - execution search path 2-6
  - execution timing 4-50
  - exit() system call, with pthread 6-9, 6-42
  - extended arithmetic 5-37
  - extended files 5-33
  - extended receive and probe 3-24
- F**
- .F extension 2-9
  - f77 command 2-4, 2-5, A-1
    - host switch B-10
    - l switch 2-9, 2-10
    - L switch 2-9, 2-10
    - Mquad switch B-4
    - node switch B-4
    - nx switch 2-5, B-4
    - order of switches 2-10
    - p switch B-4
  - failed nodes 2-31, 2-48, 4-40
  - fault LEDs 1-3
  - fcntl() system call 5-4, 5-34
  - fcube.h file B-9
  - festat() system call 5-36, A-21
  - FFT library 8-6
  - fg command 2-23
  - fgetpos() system call 5-34
  - FIFO size 8-12
  - file I/O, parallel (see also "parallel file I/O") 5-1
  - file pointers 5-14
  - file system blocks 8-25
  - file systems 5-2
    - attributes of 5-5
    - getting information about PFS file systems 5-39
  - fileID parameter 5-8
  - filenames, length of 5-3
  - files
    - extended 5-33
    - file descriptors 5-8
    - file pointers 5-14
      - moving 5-29
    - maximum open at once 5-4
    - size of 5-7, 5-32
  - find command 5-35
  - fixed-size records 5-16

- flick() system call 4-50, A-16, A-43
  - floating-point control calls 4-46
  - flock() system call 5-4
  - flockfile() system call 6-10
  - flow control of messages 8-13
  - flushing Fortran buffered I/O 5-30
  - flushmsg() system call 4-52, A-17, A-43, B-16
  - fnx.h file 2-8
  - force types 3-6, B-2
  - forceflush() system call 5-30, A-47
  - forflush() system call 5-30, A-47
  - fork() system call 2-6, 3-5, 4-10, 5-13
    - with pthreads 6-9
  - forking processes onto nodes 4-10
  - form="formatted" parameter 5-11
  - form="unformatted" parameter 5-11
  - formatted files 5-11
  - fort.nnn files 5-13
  - Fortran programs
    - error handling 4-42
    - file I/O on parallel files 5-24
    - flushing buffered I/O 5-30
    - including fnx.h 2-8
    - memory access considerations 8-5
    - message passing with Fortran commons 3-17
    - opening new files 5-12
    - opening parallel files 5-11
    - parallel file I/O calls 5-25
    - preprocessing 2-9
    - sequential files 5-26
    - unformatted files 5-26
    - units 5-8
  - fpgetmask() system call 4-46, A-16
  - fpgetround() system call 4-46, A-16
  - fpgetsticky() system call 4-46, A-16
  - fpsetmask() system call 4-46, 5-31, A-16, A-42
  - fpsetround() system call 4-46, A-16
  - fpsetsticky() system call 4-46, A-16
  - fread() system call 8-24
  - free nodes 2-48
  - front panel LEDs 1-3
  - fseek() system call 5-34
  - fsetpos() system call 5-34
  - fsplit command A-3
  - fstat() system call 5-34
  - fstatfs() system call 5-41, 8-25
  - fstatvfs() system call 5-39, 8-26, A-23
  - ftell() system call 5-34
  - FTNxxxxxxx.nn files 5-13
  - ftp command 1-6, 2-7, 5-35
  - ftruncate() system call 5-34
  - full stripe size 8-25
  - funlockfile() system call 6-10
- G**
- gang scheduling 2-35, 2-43, 2-55
    - with pthreads 6-4
  - Gauss-Seidel method 8-11
  - gcol() system call 3-27, A-10, A-35
  - gcolx() system call 3-27, 7-6, 7-13, A-10, A-35
  - gdhigh() system call 3-27, A-10, A-35
  - gdlow() system call 3-27, A-10, A-35
  - gdprod() system call 3-27, A-10, A-35
  - gdsum() system call 3-27, 3-28, 7-6, A-10, A-35
  - general cancelability 6-28

getcube command B-5  
 getcube() system call B-10  
 getiphosts() system call B-16  
 getmntinfo() system call 5-39  
 getpfsinfo() system call 5-39, A-23  
 getrusage() system call 8-9  
 getting information about PFS file systems 5-39  
 GFLOPS 8-2  
 giand() system call 3-27, A-10, A-35  
 gigabyte files 5-33  
 gihigh() system call 3-27, A-10, A-36  
 gilow() system call 3-27, A-11, A-36  
 ginv() system call 4-52, A-17, A-43  
 gior() system call 3-27, A-11, A-36  
 giprod() system call 3-27, A-11, A-36  
 gisum() system call 3-27, A-11, A-36  
 give\_threshold parameter 8-18  
 gixor() system call B-16  
 gland() system call 3-27, A-11, A-36  
 global clock 4-50  
 global operations 3-4, 3-27, 5-9, 5-13, 5-14, 5-48, 7-6  
     and controlling process 4-25  
     effect on-noc switch 8-22  
     with -on switch 2-20  
     with pthreads 6-4, 6-37  
 global predicate variables 6-21  
 glor() system call 3-27, A-11, A-36  
 glxor() system call B-17  
 gopen() system call 5-9, 8-23, A-18, A-45  
     synchronization 5-48  
     with pthreads 6-4

gopf() system call 3-27, A-11, A-37  
 gprof command 8-2  
 gray() system call 4-52, A-17, A-43  
 green LEDs 1-3  
 group of a partition 2-32, 2-56  
 groups of processes 4-22  
 gsendx() system call 3-9, A-5, A-29  
 gshigh() system call 3-27, A-11, A-37  
 gslow() system call 3-27, A-11, A-37  
 gsprod() system call 3-27, A-12, A-37  
 gssum() system call 3-27, A-12, A-37  
 gsync() system call 3-27, A-12, A-37  
     with open() B-16  
     with pthreads 6-4  
 -gth switch 8-18

## H

handled message-passing calls 3-7  
 handled types 3-19  
 handler() system call B-17  
 handling errors 4-42, 8-7  
     with pthreads 6-41  
 hard disk 5-2  
 hardware 1-2  
 hardware failures 2-31  
 heap 8-15  
 "hello, world" program 2-4  
 hierarchical partition structure 2-28  
 host calls B-9  
 -host switch B-10  
 hparam parameter 3-21



**hrecv()** system call 3-18, A-8, A-32  
     with pthreads 6-38  
**hrecvx()** system call 3-24, A-9, A-34  
**hsend()** system call 3-18, 3-20, A-8, A-32  
     with pthreads 6-38  
**hsendrecv()** system call 3-18, 3-20, A-8, A-32  
**hsendx()** system call 3-20, A-8, A-33  
**HTOC...()** system calls B-14  
**hwclock()** system call 4-52, A-17, A-43

## I

**I/O buffers** 8-25  
**I/O calls, efficiency of** 8-24  
**I/O IDs** 5-27  
**I/O interfaces** 1-2  
**I/O modes** 5-9, 5-13  
     efficiency of 8-24  
     example 5-17  
     inheritance across fork() 5-13  
     M\_GLOBAL 5-17  
     M\_LOG 5-15  
     M\_RECORD 5-16, 8-26  
     M\_SYNC 5-15  
     M\_UNIX 5-14  
     standard I/O 2-12  
     synchronization of 5-48  
**I/O nodes** 5-2  
**I/O partition** 2-26  
**I/O redirection** 2-12  
**I/O request size** 8-23, 8-25  
**I/O to parallel files** 5-24  
**I/O, parallel (see also "parallel file I/O")** 5-1

**i860 microprocessor** 1-2, 8-5  
     cache line 8-12  
     data cache 8-5  
     FIFO size 8-12  
     floating-point control registers 4-47  
     instruction cache 8-6  
     physical memory page 8-12

**icc command** 2-5, A-1  
     environment variables 2-6  
     -l switch 2-9, 2-10  
     -Knoieee switch 8-4  
     -L switch 2-9, 2-10  
     -Mnostride0 switch 8-4  
     -Mquad switch B-4  
     -Mvect switch 8-4  
     -node switch B-4  
     -nx switch 2-5, B-4  
     -O switch 8-4  
     order of switches 2-10  
     -p switch B-4

**ID of a message** 3-6

**IEEE math library** 8-4

**IEEE NaN** 4-47

**if/else blocks, efficiency of** 8-6

**if77 command** 2-5, A-1  
     environment variables 2-6  
     -l switch 2-9, 2-10  
     -Knoieee switch 8-4  
     -L switch 2-9, 2-10  
     -lkmath switch 7-13  
     -Mnostride0 switch 8-4  
     -Mquad switch B-4  
     -Mvect switch 8-4  
     -node switch B-4  
     -nx switch 2-5, B-4  
     -O switch 8-4  
     order of switches 2-10  
     -p switch B-4

**image enhancement** 7-3

**improving performance** 8-1

- inactive applications 2-35
- include directories 2-8
- include files
  - cube.h B-9
  - estat.h 5-37
  - fcube.h B-9
  - fnx.h 2-8
  - mtio.h 5-44
  - nx.h 2-8
- incomplete (asynchronous) system calls
  - file I/O 5-24
  - message passing 3-7
- increasing problem size 8-5
- increasing the size of a file 5-7, 5-32
- info parameter 3-25
- info...() system calls, with pthreads 6-38
- infocount() system call 3-15, A-7, A-31
- infnode() system call 3-15, A-7, A-31
- infopid() system call 4-52, A-17, A-43
- infoftype() system call 3-15, A-7, A-31
- information about messages 3-15
- infotype() system call 3-15, A-7, A-31
- innermost loops 8-4
- instruction cache 8-6
- instruction pointer 6-1
- Intel supercomputer
  - hardware 1-2
  - software 1-4
  - using commands on 2-2
- interactive applications 2-34
- interconnect network 1-2
- interfaces 1-2
- interrupt key 2-23
- interrupts
  - preventing 3-22
  - treating messages as interrupts 3-18
- INVALID\_PTYPE constant 3-5
- ioctl() system call 5-44
- iDONE() system call 5-24, 5-27, 5-28, A-19, A-46
- iomode() system call 5-13, A-18, A-45
- iowait() system call 5-24, 5-27, 5-28, A-19, A-46
- iprobe() system call 3-14, A-7, A-31
- iprobex() system call 3-16, 3-24, A-9, A-34
- IPSC system
  - CFS compatibility 5-1
  - commands B-5
  - compatibility calls 4-52
  - compatibility with B-1
  - compilers 2-7, B-4
  - IPSC\_XDEV environment variable B-4
  - system calls B-9
- iread() system call 5-27, A-19, A-46, B-15
  - synchronization 5-48
- ireadv() system call 5-27, 5-48, A-19, A-46
- irecv() system call 3-10, A-6, A-30
- irecvx() system call 3-24, A-9, A-33
- isend() system call 3-10, A-6, A-30
- isendrecv() system call 3-10, A-6, A-30
- iseof() system call 5-15, 5-24, 5-26, 5-29, A-20, A-47
  - synchronization 5-48
- isnan() system call 4-46, A-16
- isnand() system call 4-46, A-16
- isnanf() system call 4-46, A-16
- italic text vi
- iwrite() system call 5-27, 5-31, A-19, A-46, B-15
  - synchronization 5-48

`iwritev()` system call 5-27, 5-48, A-19, A-46

## K

kernel of an application 7-3

kernel threads 6-3, 6-16

kill command 1-4, 2-24, 2-51

`kill()` system call 4-15, 4-23

killcube command B-6

`killcube()` system call 4-52, A-17, A-43, B-10

killing application processes 4-23

`killproc()` system call 4-52, A-17, A-44, B-10

`killsyslog()` system call B-11

-Knoieee switch 8-4

## L

`led()` system call 4-52, A-16, A-43

LEDs 1-3

length of a filename or pathname 5-3

length of a message 3-5, 3-15

less command B-8

`lestat()` system call 5-36, A-21

`libc_r.a` 6-2, 6-6  
error handling 6-42

lib-coff directory 2-8

`libkmath.a` 8-6

`libmach.a` 6-5

`libnx.a` 2-5  
with pthreads 6-3

`libpthreads.a` 6-2, 6-11  
error handling 6-42

## libraries

Basic Math Library 8-6

BLAS 7-13

command-line switches 2-10

IEEE math library 8-4

`libkmath.a` 7-13

`libnx.a` 2-5

pthreads package 6-5

search path for 2-10

Signal Processing Library 8-6

specifying 2-8

`libsignal.a` 8-6

life of a process type 3-5

limitations of PFS 5-4

limitations of pthreads 6-3

link switches 2-10

`link()` system call 5-31

linking an application 2-3

single-pass linker 2-10

specifying library pathnames 2-8

with pthreads 6-5

listing partitions 2-49

listing the applications in a partition 2-51

-lkmath switch 7-13

-lnx switch 2-6

effect on execution 2-11

order of 2-10

with pthreads 6-5

load balancing 2-33, 7-3

load command B-6

`load()` system call 4-52, A-17, A-44

loading processes onto nodes 4-11

locality of data 8-10

locking a process in memory 8-15

locking and unlocking pthreads 6-16

- locking data into memory 8-15
- logical node numbers 2-30
- loops, innermost 8-4
- loops, size of 8-6
- ls command 5-5, 5-35
- lseek() system call 5-15, 5-17, 5-26, 5-29, 5-34, A-20, A-47
  - synchronization 5-48
- lsize command 5-7, A-3
- lsize() system call 5-32, A-20, A-47
- lspart command 2-31, 2-49, A-2
  - r switch 2-50
- lstat() system call 5-34
  
- M**
- M\_GLOBAL I/O mode 5-17, 8-24
- M\_LOG I/O mode 5-15
- M\_RECORD I/O mode 5-16, 8-24, 8-26
- M\_SYNC I/O mode 5-15, 8-24
- M\_UNIX I/O mode 5-14, 8-24
- Mach kernel interface 6-5
- Mach threads 6-2
- madvise() system call 5-34
- magnetic tapes, controlling 5-44
- main thread 6-12
- maintaining data locality 8-10
- makewhatis command B-7
- making partitions 2-39, 4-28
- making the program independent of the number of nodes 7-5
- malloc() system call 8-8, 8-13, 8-15
  
- manager-worker decomposition 7-5, 7-14
- managing partitions 2-25
  - with system calls 4-27
- managing running applications 2-23
- manpath command B-8
- MANPATH environment variable 2-6
- manual pages, configuring your environment for 2-6
- masktrap() system call 3-22, A-8, A-33
- math library, IEEE 8-4
- matrix\*vector example 7-11
- maximum capacity of a PFS file system 5-3
- maximum length of a filename or pathname 5-3
- maximum number of open files 5-4
- maximum size of a PFS file 5-3
- mbf switch 8-16
- mclock() system call 4-52, A-17, A-44
- mea switch 8-17
- memory
  - accessing contiguous memory locations 8-5
  - allocated to message buffers 8-16
  - distributed 7-2
  - dynamic allocation 8-8
  - locking 8-15
  - locking data into memory 8-15
  - of nodes 1-2
  - physical 1-5
  - physical pages 8-5, 8-12
  - static allocation 8-8
  - virtual 1-5, 8-3, 8-15
- memory\_each parameter 8-17
- memory\_export ipc\_option 8-19
- merging message IDs 3-13
- message buffers 8-16
- message coprocessor 8-10

- message handlers 3-20
- message\_buffer parameter 8-16
- message-passing configuration switches 8-18
- message-passing flow control 8-13
- message-passing system calls 3-1
  - with pthreads 6-4, 6-37
- messages 1-1, 7-2
  - as interrupts 3-18
  - asynchronous calls 3-10
  - buffers 3-14, 8-11
    - aligning 8-12
  - configuration options 8-18
  - designing a communication strategy 7-6
  - exchanging with controlling process 4-25
  - force types 3-6
  - getting information about 3-15
  - handled types 3-19
  - memory allocated to message buffers 8-16
  - merging message IDs 3-13
  - message characteristics 3-5
  - message ID 3-6
  - message IDs 3-10
  - message length 3-5, 3-15
  - message order 3-7
  - message passing with Fortran commons 3-17
  - message type 3-6, 3-15
  - names of message-passing calls 3-7
  - pending messages 3-14, 8-11
  - performance improvement techniques 8-7, 8-21
  - pthreads 6-37
  - releasing message IDs 3-12
  - synchronous calls 3-8
  - typesel masks 3-6
  - zero-length messages 3-6
- mex switch 8-19
- MFLOPS 8-2
- miscellaneous system calls 4-50
- mkcfs command B-7
- mkdev command B-8
- mkdir() system call 5-31
- mknod() system call 5-31
- mkpart command 2-29, 2-39, A-2
  - epl switch 2-44
  - mod switch 2-42
  - nd switch 2-40
  - rq switch 2-43
  - sps switch 2-43
  - ss switch 2-43
  - sz switch 2-40
- mmap() system call 5-4, 5-34
- Mnostride0 switch 8-4
- modes for I/O 5-13
  - synchronization of 5-48
- modes of a partition 2-32, 2-42, 2-56
- monospace text vi
- more command 5-35
- mount points 5-2
- moving the file pointer 5-29
- mp\_switches 8-18
- mprotect() system call 5-34
- Mquad switch B-4
- msgcancel() system call 4-52, A-17, A-44
- msgdone() system call 3-10, 3-16, A-6, A-30
- msgignore() system call 3-10, A-6, A-31
- msginfo array 3-16, 3-25
- msgmerge() system call 3-13, A-6, A-31
- msgwait() system call 3-10, 3-16, A-6, A-30
- msync() system call 5-34
- MT operations 5-44
- mtio.h file 5-44
- multi-node performance 8-7, 8-21
- multiple nodes 3-9

multiple programs in an application 2-21  
 munmap() system call 5-34  
 mutexes 6-1, 6-16  
 mv command 5-5, 5-35  
 -Mvect switch 8-4  
 myapp (any application) command 2-11  
 myapp.c 2-4  
 myapp.f 2-4  
 myhost() system call 4-25, A-4, A-28  
 mynode() system call 2-17, 3-3, 8-2, A-4, A-28  
 mypart() system call 2-17, 4-17, 4-52, A-17, A-44  
 mypid() system call 4-52, A-17, A-44  
 myptype() system call 2-18, 3-4, 8-2, A-4, A-28

## N

name of a partition 2-30, 2-56  
 named commons in message passing 3-17  
 names of message-passing calls 3-7  
 NaN (Not-a-Number) 4-47  
 native commands 2-5  
 new features in Paragon OSF/1 B-2  
 new files 5-12  
 newfs command 5-35  
 newserver command B-6  
 newserver() system call B-11  
 NEXTPATH() macro 5-40  
 NFS (Network File System) 1-6, 2-7, 5-2  
   accessing PFS files 5-4  
   parallel I/O to 5-24  
 -noc switch 8-17  
 node interconnect network 1-2

node numbers 3-3  
   in filenames 5-31  
   in overlapping partitions 2-32  
   logical 2-30  
   of a received or pending message 3-15  
   of controlling process 4-25  
   physical 2-30  
   within a partition 2-30  
 node parameter 3-3  
 \_\_NODE preprocessor symbol 2-5  
 node programs  
   error handling 8-7  
 -node switch B-4  
 nodedim() system call 4-52, A-18, A-44  
 nodes 1-1, 1-2, 7-2  
   allocated to a partition 2-30, 2-40  
   allocated to an application 2-15, 4-4  
   compute nodes 1-2  
   contiguous 2-16, 2-40, 4-5, 4-28  
   copying processes onto nodes 4-10  
   designing a communication strategy 7-6  
   free 2-48  
   I/O nodes 5-2  
   load balancing 7-3  
   loading processes onto nodes 4-11  
   making programs independent of number of  
     nodes 7-5  
   node numbers 3-3  
   node topologies 7-6  
   operating system 1-4  
   partitions 2-25  
   running application processes on a subset 2-18  
   service nodes 1-2  
   unusable nodes 2-31, 2-48, 4-40  
 nodesel parameter 3-3  
 nodespecs 2-19, 2-40  
 noieee switch 8-4  
 noncontiguous nodes 2-16, 2-40, 4-5, 4-28  
 noncontiguous partitions 2-31

- non-parallel programs 2-11, 2-27
  - nostride0 switch 8-4
  - Not-a-Number (NaN) 4-47
  - notational conventions used in the manual vi
  - nqueens example 7-13
  - nsh command B-9
  - number of bytes read or written 5-26
  - numbers, extended 5-37
  - numerical methods 8-1
  - numnodes() system call 2-17, 3-3, A-4, A-28
  - nx switch 2-5, B-4
    - actions performed by 4-3
    - and nx\_initve() 4-4
    - command-line switches 2-13, 8-18
    - effect on execution 2-11
    - with pthreads 6-5
  - nx.h file 2-8
  - nx\_...() system calls, error handling of 4-42
  - nx\_app\_nodes() system call 4-16, A-13, A-39
  - nx\_app\_rect() system call 2-17, 4-16, 8-10, A-13, A-39
  - nx\_chpart\_epl() system call 4-36, A-14, A-41
  - nx\_chpart\_mod() system call 4-36, A-14, A-41
  - nx\_chpart\_name() system call 4-36, A-14, A-41
  - nx\_chpart\_owner() system call 4-36, A-15, A-41
  - nx\_chpart\_rq() system call 4-36, A-14, A-41
  - nx\_chpart\_sched() system call 4-36, A-15, A-41
  - NX\_DFLT\_PART environment variable 2-3, 2-12, 2-14
  - NX\_DFLT\_SIZE environment variable 2-12, 2-17, 2-22
  - nx\_empty\_nodes() system call 4-40, A-15, A-42
  - nx\_failed\_nodes() system call 4-40, A-15, A-42
  - nx\_initve() system call 4-3, 4-4, 4-21, A-12, A-38
    - linking 2-6
    - with pthreads 6-39
  - nx\_initve\_rect() system call 4-3, 4-7, A-12, A-38
  - nx\_load() system call 2-20, 3-5, 4-3, 4-11, 4-14, A-13, A-39
  - nx\_loadve() system call 2-20, 3-5, 4-3, 4-13, 4-14, A-13, A-39
  - nx\_mkpart() system call 4-28, A-14, A-40
  - nx\_mkpart\_map() system call 4-28, A-14, A-40
  - nx\_mkpart\_rect() system call 4-28, A-14, A-40
  - nx\_nfork() system call 2-20, 3-5, 4-3, 4-10, 4-14, A-12, A-38
    - with pthreads 6-39
  - nx\_part\_attr() system call 4-31, A-14, A-40
  - nx\_part\_nodes() system call 4-31, A-14, A-41
  - nx\_perror() system call 4-42, A-15, A-42
    - with pthreads 6-42
  - nx\_pri() system call 4-3, 4-9, A-12, A-38
  - nx\_pspart() system call 4-16, A-13
  - nx\_rmpart() system call 4-30, A-14, A-40
  - nx\_waitall() system call 4-3, 4-14, 4-45, A-13, A-39
- O**
- on switch 2-19
  - open() system call 5-10, 5-11, 5-31, 8-23, B-15
  - opening parallel files 5-9
    - ### in filenames 5-31
    - special considerations for Fortran 5-11
    - with standard operations 5-11
  - operating system 1-4
  - "Operation not supported by this file system" error 5-4

- optimization
  - compiler 8-3
- optimizations 8-3
- order of application switches 2-13
- order of compiler switches 2-10
- order of messages 3-7
- organization of the manual v
- OSF/1 operating system 1-4
  - commands 2-2
- OSF/1 PIDs 4-15
- other system calls 4-1
- overlapping computation and communication 8-10
- overlapping partitions or applications 2-36
- owner of a partition 2-32, 2-56
  
- P**
- p switch B-4
- packet\_size parameter 8-16
- packetization 8-16
- padding in common blocks 3-17
- pages of physical memory 8-5
- paging 8-3
  - preventing 8-15
- Paragon OSF/1 operating system 1-4
  - commands 2-1
  - message-passing system calls 3-1
  - new features B-2
  - other system calls 4-1
  - parallel file I/O 5-1
  - programming model 7-2
- Paragon system
  - hardware 1-2
  - software 1-4
- PARAGON\_XDEV environment variable 2-6, 2-9
  
- \$PARAGON\_XDEV/paragon directory 2-8
- ParaGraph performance visualization tool 8-10
- parallel applications 1-1, 2-1
- parallel file I/O 5-1
  - ### in filenames 5-31
  - asynchronous I/O calls 5-27
    - efficiency of 8-24
  - closing files 5-28
  - detecting end-of-file 5-29
  - efficiency of 8-24
  - error handling 5-26
  - file pointers 5-14
  - flushing Fortran buffered I/O 5-30
  - formatted vs. unformatted I/O 5-11
  - I/O modes 5-9, 5-13
    - efficiency of 8-24
  - I/O performance 8-23
  - in Fortran programs 5-25
  - increasing the size of a file 5-7, 5-32
  - manipulating extended files 5-36
  - moving the file pointer 5-29
  - new files 5-12
  - opening files 5-9
    - with standard operations 5-11
  - reading and writing files 5-24
  - scattered read and write 5-25
  - special considerations for Fortran 5-11
  - synchronizing calls 5-9, 5-13, 5-14
  - synchronizing operations 5-48
  - synchronous I/O calls 5-25
  - system calls 5-8
  - tapes, controlling 5-44
  - to NFS files 5-24
  - to the user's terminal 5-31
  - unnamed files 5-13
  - with pthreads 6-4, 6-38
- parallel file system (see also "PFS") 5-1
- parallel programming techniques 7-2
- parent partition 2-29
- "partition permission denied" error 2-12



- partitions 1-6, 2-2, 2-25
  - allocating nodes to applications 2-15
  - changing partition characteristics 2-54, 4-36
  - characteristics 2-29
  - child partitions 2-29, 2-34
  - compute partition 2-2, 2-27, 2-28
  - contiguous and noncontiguous 2-31
  - contiguous nodes 2-40
  - current 2-28
  - default characteristics 2-29
  - default partition 2-2, 2-14
    - determining 2-15
    - setting 2-14
  - dot (.) partition (root partition) 2-26
  - effective priority limit 2-35, 2-44, 2-56
  - error messages 2-14
  - execute (x) permission 2-33
  - free nodes 2-48
  - gang-scheduled 2-35, 2-43, 2-55
  - hierarchical structure 2-28
  - I/O partition 2-26
  - listing 2-49
  - listing the applications in a partition 2-51
  - making partitions 2-39, 4-28
  - managing 2-25
  - managing with system calls 4-27
  - name of a partition 2-30, 2-56
  - nodes 2-30
  - nodes allocated to a partition 2-40
  - overlapping 2-36
  - owner and group 2-32, 2-56
  - parent partition 2-29
  - pathnames 2-28
  - permission bits 2-32, 2-42, 2-56
  - priority 2-35
  - protection modes 2-32, 2-42, 2-56
  - read (r) permission 2-33
  - removing partitions 2-45, 4-30
  - rollin quantum 2-34, 2-35, 2-43, 2-55
  - root partition 2-26, 2-29
    - shape of 2-26
  - running applications in 2-22
  - scheduling characteristics 2-33, 2-43
  - service partition 2-2, 2-24, 2-27
  - showing partition characteristics 2-46
  - space-shared 2-37, 2-43, 2-56
  - special 2-26
  - standard-scheduled 2-34, 2-43
  - subpartitions 2-29, 2-34
  - unusable nodes 2-31, 2-48, 4-40
  - write (w) permission 2-33
- passing information to the handler 3-20
- PATH environment variable 2-6
- pathnames of partitions 2-28
- pathnames, length of 5-3
- pending messages 3-14, 8-11
  - getting information about 3-15
- perfectly-parallel applications 1-6, 7-2
- performance improvement techniques 8-1
- performance of PFS files 8-23
- performance visualization 8-10
- performing extended arithmetic 5-37
- permissions of a partition 2-32, 2-42, 2-56
- per-node vector size 8-5
- perorr() system call, with pthreads 6-9, 6-42

**PFS 5-3**

- accessing via NFS 5-4
- commands 5-5
- core files in 5-4
- executable files in 5-4
- file systems 5-3
- filename and pathname length 5-3
- files 5-3
- getting information about PFS file systems 5-39
- limitations of 5-4
- manipulating extended files 5-36
- maximum capacity 5-3
- maximum file size 5-3
- maximum number of open files 5-4
- mount points 5-2
- opening PFS files 5-9
  - with standard operations 5-11
- performance 8-23
- special files in 5-3
- stripe directories 5-3
- stripe units 5-3

**/pfs file system 5-5****PFS file systems**

- block size 8-25
- striping 8-25

**pfsmntinfo structure 5-39****physical memory 1-5****physical memory page 8-12****physical memory pages 8-5****physical nodes 2-30****physical topology 7-6****pi example 7-7****PIDs (process IDs) 4-14**

- contrasted with process types 4-15

**-pkt switch 8-16****-plk switch 8-3, 8-15****plock() system call 8-15****plogon and plogoff commands B-8****plogon() and plogoff() system calls B-17****pmake command A-3****-pn switch 2-22****pointers to message buffers 8-13****porting iPSC programs B-1****porting serial codes 7-5****POSIX threads 6-1****preallocating disk space 5-33****preprocessing a Fortran program 2-9****preprocessor symbol \_\_NODE 2-5****preventing interrupts 3-22****-pri switch 2-17****priority**

- effective priority limit of a partition 2-35, 2-44, 2-56
- of a partition 2-35
- of a process 2-35
- of an application 2-17, 2-35, 4-9

**probing for pending messages 3-14**

- extended 3-24

**problem decomposition 7-3****problem size 8-5****process group IDs 2-51****process group leaders 4-4, 4-22****process IDs (PIDs) 4-14****process locking 8-15****process types 3-4**

- changing 3-4
- contrasted with OSF/1 PIDs 4-15
- INVALID\_PTYPE 3-5
- life of 3-5
- of a received or pending message 3-15
- of an application 2-18
- of controlling process 4-25
- with pthreads 3-5, 6-37

- processes
  - arbitration between 2-33
  - characteristics 3-3
  - child processes 4-10, 4-11, 4-14
  - controlling process 2-24, 2-28, 2-51, 4-4, 4-21
  - copying processes onto nodes 4-10
  - loading processes onto nodes 4-11
  - PIDs (process IDs) 4-14
  - priority of 2-35
  - process groups 4-22
  - process types 3-4
  - threads 1-1, 6-1
  - waiting for application processes 4-14
- processor time 2-17
- processors 1-1
- prof command 8-2
- profiling B-4
- profiling tools 8-2
- program development tools 1-6
- programming model 1-6, 7-2
- programming techniques 7-2
- programs, non-parallel 2-1, 2-27
- protection modes of a partition 2-32, 2-42, 2-56
- ps command 2-24
- pspart command 2-51, A-2
  - r switch 2-54
- pt switch 2-18, 3-4
- pthread\_attr\_create() system call 6-15, A-24
- pthread\_attr\_delete() system call 6-15, A-24
- pthread\_attr\_getstacksize() system call 6-15, A-24
- pthread\_attr\_setstacksize() system call 6-15, A-24
- pthread\_cancel() system call 6-28, A-26
- pthread\_cleanup\_pop() system call 6-32, A-27
- pthread\_cleanup\_push() system call 6-32, A-27
- pthread\_cond\_broadcast() system call 6-21, A-26
- pthread\_cond\_destroy() system call 6-21, A-26
- pthread\_cond\_init() system call 6-21, A-26
- pthread\_cond\_signal() system call 6-21, A-26
- pthread\_cond\_timedwait() system call 6-21, A-26
- pthread\_cond\_wait() system call 6-21, A-26
- pthread\_condattr\_create() system call 6-23, A-26
- pthread\_condattr\_delete() system call 6-23, A-26
- pthread\_create() system call 6-13, A-24
- pthread\_detach() system call 6-13, A-24
- pthread\_equal() system call 6-13, A-24
- pthread\_exit() system call 6-13, A-24
- pthread\_getspecific() system call 6-33, A-27
- pthread\_join() system call 6-13, A-24
- pthread\_keycreate() system call 6-33, A-27
- pthread\_mutex\_destroy() system call 6-16, A-25
- pthread\_mutex\_init() system call 6-16, A-25
- pthread\_mutex\_lock() system call 6-16, A-25
- pthread\_mutex\_trylock() system call 6-16, A-25
- pthread\_mutex\_unlock() system call 6-16, A-25
- pthread\_mutexattr\_create() system call 6-17, A-25
- pthread\_mutexattr\_delete() system call 6-17, A-25
- pthread\_once() system call 6-34, A-27
- pthread\_self() system call 6-13, A-24
- pthread\_setasynccancel() system call 6-28, A-26
- pthread\_setcancel() system call 6-28, A-26
- pthread\_setspecific() system call 6-33, A-27
- pthread\_testcancel() system call 6-28, A-26
- pthread\_yield() system call 6-13, A-24

**pthread**

- attributes of 6-15
- Bourne shell 6-39
- canceling 6-28
- cleanup routines 6-32
- compiling and linking 6-5
- condition variables 6-21
- data types and symbols 6-11
- error handling 6-41
- example program 6-18
- execution of 6-13
- file I/O 6-38
- global predicate variables 6-21
- kernel threads 6-3
- keys 6-33
- libraries 6-2
- limitations of 6-3
- locking and unlocking 6-16
- main thread 6-12
- message-passing calls 6-37
- mutexes 6-16
- non-thread-safe code 6-37
- priority of 6-15
- process types 3-5
- pthread library calls 6-11
- pthread package 6-1
- pthread-specific data objects 6-33
- recommended safe operating environment 6-4
- reentrant C library 6-6
- signals 6-34, 6-39
- stack size of 6-15
- synchronization of 6-21

ptype (see also "process type") 2-18

ptype parameter 3-4

ptypesel parameter 3-4

**Q**

quad-alignment B-4

queens example 7-13

quick example 2-3

quotaon command 5-4

**R**

\_r calls 6-10

RAID (Redundant Array of Inexpensive Disks) 5-2

rar command B-8

ras command B-8

rcc command B-8

rcmd command 1-6

rcp command 1-6, 2-7, 5-35

read (r) permission on a partition 2-33

read statement 5-24

read() system call 5-25  
with pthreads 6-4

reading files in parallel 5-24

readlink() system call 5-31

readv() system call 5-25

rebootcube command B-9

receiving messages 3-8  
extended 3-24

record size 5-16

rectangular applications 2-16

rectangular dimensions of an application 4-16

recursively listing applications in subpartitions 2-54

recursively listing subpartitions 2-50

recursively removing subpartitions 2-45

red LEDs 1-3

redirecting I/O 2-12

reentrant C library 6-2, 6-6

\_REENTRANT preprocessor symbol 6-5, 6-41

reentrant software 6-1

related documents vii

relative partition pathname 2-28

relcube command B-6

relcube() system call B-11

releasing control of the processor 4-50

releasing I/O IDs 5-28

releasing message IDs 3-12

relstruc() system call B-14

remote host commands B-8

removing partitions 2-45, 4-30

rename() system call 5-31

repeated use of system calls 8-2

"request overlaps with nodes in use" error 2-12,  
2-16, 2-34

request size 8-23, 8-25

resources 6-1

restrictvol() system call 4-52, A-18, A-44

rewinding a tape 5-46

rf77 command B-8

ring topology 7-6

rld command B-8

rlogin command 1-5

rm command 5-35

rmdir() system call 5-31

rmpart command 2-45, A-2

- f switch 2-45
- r switch 2-45

rollin and rollout 2-2, 4-50

rollin quantum 2-34, 2-35

- of a partition 2-43, 2-55

root account 2-26, 2-33, 2-56

root partition 2-26, 2-29

- shape of 2-26

rounding mode 4-47

RPM global clock 4-50

rsh command 1-6

running a program on a subset of the nodes 2-18

running applications 2-11

- consisting of multiple programs 2-21
- in a particular partition 2-22
- removing partitions containing 2-45

## S

sat command A-3

scattered read and write 5-25

scheduling characteristics of a partition 2-33, 2-43

scheduling mechanisms 1-1

"scheduling parameters conflict with allocator  
configuration" error 2-36

scratch files 5-13

SCSI interface 1-2, 5-2

- sct switch 8-18

sdot() BLAS function 7-13

search path 2-6

search path for libraries 2-10

seeking on a file 5-29

send\_avail value 8-17

send\_count parameter 8-17

send\_threshold parameter 8-17

sending messages 3-8

sending to multiple nodes 3-9

separating the user interface from the computation  
7-3

- sequential files 5-26
- serial codes, porting 7-5
- serializing calls 8-9
- service nodes 1-2
- service partition 2-2, 2-11, 2-24, 2-27
- setiomode() system call 5-13, A-18, A-45
  - synchronization 5-48
  - with pthreads 6-4
- setiphost() system call B-17
- setpart alias 2-15
- setpgrp() system call B-17
- setpid() system call B-11
- setptype() system call 3-4, A-4, A-28
  - in controlling process 4-25
- setsyslog() system call B-11
- setting your default partition 2-14
- sh command 2-24
  - with pthreads 6-39
- shadow buffers 8-11
- shape of an application 4-16
- shape of the root partition 2-26
- shell 2-27
- shell scripts 2-24
- shepherd process 2-24
- shows command 5-5, 8-26, A-3
- showing partition characteristics 2-46
- showpart command 2-29, 2-31, 2-46, 2-50, A-2
  - f switch 2-48
- showvol command B-7
- Signal Processing Library (libsignal.a) 8-6
- signal() system call 4-44
- signals, with pthreads 6-5, 6-34, 6-39
- sigwait() system call 6-5, 6-34, A-27
- single program multiple data (SPMD) programming model 1-6, 7-2
- single system image 1-4, 7-2
- single-node performance 8-2
- 64-bit integers 5-37
- size
  - of a file 5-7, 5-32
  - of a message 3-5
  - of a packet 8-16
  - of a partition 2-40
  - of an application 2-15
    - rectangular 2-16
- size of a problem 8-5
- sizeof operator 3-6
- sleep system call, with pthreads 6-9
- software failures 2-31
- space sharing 2-37, 2-43, 2-56
- special partitions 2-26
- specifying application priority 2-17
- specifying application size 2-15
- specifying nodes allocated to a partition 2-40
- specifying process type 2-18
- speed of calculation 8-2
- speedup of a parallel program 7-2
- SRM B-1
- stack 6-1, 8-15
- stack size of pthreads 6-15
- standard include directory 2-9
- standard input and output, redirecting 2-12
- standard scheduling 2-34, 2-43
- star command B-7

- startcube command B-6
- start-up routine 2-5
- stat() system call 5-31, 5-34, 8-9
- statfs() system call 5-31, 5-41, 8-25
- static memory allocation 8-8
- statvfs structure 5-39
- statvfs() system call 5-39, 8-26, A-23
- status of a tape device 5-46
- status="new" parameter 5-12
- status="scratch" parameter 5-13
- sth switch 8-18
- sticky flags 4-48
- stoe() system call 5-37, 5-38, A-22, A-49
- stream command B-7
- stripe directories 5-3, 8-25
- stripe example 8-26
- stripe factor 8-25
- stripe unit 5-3, 8-25
- strip-mining loops 8-5
- structures, padding of 8-13
- subpartitions 2-29, 2-34
  - creating 2-39
  - listing 2-49
  - listing the applications in a subpartition 2-54
  - removing 2-45
  - removing partitions containing 2-45
- succ() function 7-7
- summaries of commands and system calls A-1
- supercomputer
  - hardware 1-2
  - software 1-4
  - using commands on 2-2
- suspend key 2-23
- SVR3.2 B-9
- switches
  - compiler 2-8
  - compiler optimization 8-3
  - gth switch 8-18
  - host switch B-10
  - in nx\_initve() 4-6
  - mbf switch 8-16
  - mea switch 8-17
  - mex switch 8-19
  - mp\_switches 8-18
  - Mquad switch B-4
  - noc switch 8-17
  - node switch B-4
  - nx switch 2-5, 2-13, 8-18, B-4
  - on switch 2-19
  - order of switches 2-13
  - p switch B-4
  - pkt switch 8-16
  - plk switch 8-3, 8-15
  - pn switch 2-22
  - pri switch 2-17
  - pt switch 2-18, 3-4
  - sct switch 8-18
  - sth switch 8-18
  - sz switch 2-16, 3-3
- symlink() system call 5-31
- synchronization of pthreads 6-21
- synchronizing calls 5-9, 5-13, 5-14
- synchronizing operations 5-29, 5-30
  - summary 5-48
- synchronous and asynchronous calls 8-10
- synchronous and asynchronous I/O calls 8-24
- synchronous file I/O calls 5-25
- synchronous message-passing calls 3-7, 3-8
  - with pthreads 6-37
- sys/estat.h file 5-37
- sysacct command 5-4
- syslog command B-6

system administrator 2-26, 2-33, 2-56

system buffers 3-14, 8-11

#### system calls

- asynchronous file I/O calls 5-27
- asynchronous message-passing calls 3-10
- backward compatibility 4-52, B-9
- closing files in parallel 5-28
- controlling application execution 4-2
- controlling tape devices 5-44
- detecting end-of-file 5-29
- error handling 4-42, 8-7
  - in parallel file I/O calls 5-26
- extended arithmetic 5-37
- floating-point control 4-46
- flushing Fortran buffered I/O 5-30
- global operations 3-27, 7-6
- I/O modes 5-13
- I/O to parallel files 5-24
- increasing the size of a file 5-32
- information about messages 3-15
- iPSC system compatibility 4-52, B-9
- manipulating extended files 5-36
- message buffers 3-14, 8-11
- message passing with Fortran commons 3-17
- message-passing 3-1
- miscellaneous 4-50
- moving the file pointer 5-29
- names of message-passing calls 3-7
- opening files in parallel 5-9
- other system calls 4-1
- parallel file I/O 5-8
- parallel file I/O synchronization 5-48
- partition management 4-27
- reading and writing files in parallel 5-24
- repeated 8-2
- summary of C system calls A-4
- summary of Fortran system calls A-28
- synchronization 5-48
- synchronous file I/O calls 5-25
- synchronous message-passing calls 3-8
- timing 4-50
- treating messages as interrupts 3-18
- underscore versions 4-42, 8-7
  - with pthreads 6-43

system hardware 1-2

system message buffers 8-16

system software 1-4

System V UNIX B-9

-sz switch 2-16, 3-3

## T

tail command 5-35

tape devices, controlling 5-44

tapemode command B-7

tar command 5-35

task decomposition 7-5

techniques for improving performance 8-1

techniques for parallel programming 7-2

telnet command 1-5

temporarily releasing control of the processor 4-50

terminal I/O 5-31

terminology 2-1

threads 1-1, 6-1  
(see also "pthreads")

thread-safe software 6-1

tiling 2-37

timing execution 4-50

tips for compiling and linking 2-8

tools for program development 1-6

topics in this manual v

topologies 7-6

Touchstone DELTA System 4-52

treating a message as an interrupt 3-18

tree search 7-5



triangle example 7-18  
 truncate() system call 5-31, 5-34  
 type of a message 3-6, 3-15  
 type parameter 3-6  
 typesel masks 3-6  
 typesel parameter 3-6

## U

UFS file systems 5-2  
 underscore versions of system calls 4-42, 8-7  
     with pthreads 6-43  
 understanding message-passing flow control 8-13  
 unformatted files 5-11, 5-26  
 unit stride 8-4  
 units (Fortran I/O) 5-8  
 UNIX System V B-9  
 unlink() system call 5-31  
 unlocked\_...() system calls 6-10  
 unlocked\_fseek() system call 5-34  
 unlocking pthreads 6-16  
 unnamed files 5-13  
 unusable nodes 2-31, 2-48, 4-40  
 uppercase .F extension 2-9  
 user interface of an application 7-3  
 user model 1-5  
     differences from iPSC system B-5  
 using Paragon OSF/1 commands  
     on the Intel supercomputer 2-2  
     on workstations 2-2  
 using PIDs 4-14  
 using the default partition 2-14

/usr/ccs/lib directory 2-8  
 /usr/include directory 2-8  
 /usr/lib directory 2-8  
 /usr/paragon/XDEV directory 2-9  
 /usr/tmp directory 5-13  
 utimes() system call 5-31

## V

variables  
     CFS\_MOUNT 5-13  
     errno 4-42  
         with pthreads 6-41  
     IPSC\_XDEV B-4  
     MANPATH 2-6  
     NX\_DFLT\_PART 2-12, 2-14  
     NX\_DFLT\_SIZE 2-12, 2-17, 2-22  
     PARAGON\_XDEV 2-6, 2-9  
     PATH 2-6  
 vector library 8-6  
 vector multiplication 7-11  
 vector operations 3-27  
 vector size 8-5  
 vi command 5-35  
 virtual memory 1-5, 8-3, 8-15  
 virtual topology 7-6  
 visualization of performance 8-10  
 vm\_stat command 8-3

## W

wait() system call, with pthreads 6-9  
 waitall() system call B-12  
 waitcube command B-6  
 waiting for application processes 4-14

**waitone() system call B-13**

**wildcards in partition pathnames 2-28**

**wiring memory 8-15**

**workstations**

architecture of 2-6

using commands on 2-2

**workstations, working at 1-6**

**write (w) permission on a partition 2-33**

**write statement 5-24**

**write() system call 5-25**

with pthreads 6-4

**writev() system call 5-25**

**writing files in parallel 5-24**

## **X**

**x (execute) permission on a partition 2-33**

**xprof and xgprof commands 8-2**

## **Y**

**yellow LEDs 1-4**

## **Z**

**zero-length messages 3-6**