

System Interfaces Guide

2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.



Amazon Ex. 1012
IPR Petition - USP 7,784,058

© 1995 Sun Microsystems, Inc. 2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

All rights reserved. This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX® system, licensed from UNIX System Laboratories, Inc., a wholly owned subsidiary of Novell, Inc., and from the Berkeley 4.3 BSD system, licensed from the University of California. Third-party software, including font technology in this product, is protected by copyright and licensed from Sun's suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19. The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS

Sun, Sun Microsystems, the Sun logo, SunSoft, the SunSoft logo, Solaris, SunOS, OpenWindows, DeskSet, ONC, ONC+, and NFS are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd. OPEN LOOK is a registered trademark of Novell, Inc. PostScript and Display PostScript are trademarks of Adobe Systems, Inc.

All SPARC trademarks are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. SPARCcenter, SPARCcluster, SPARCcompiler, SPARCdesign, SPARC811, SPARCengine, SPARCprinter, SPARCserver, SPARCstation, SPARCstorage, SPARCworks, microSPARC, microSPARC-II, and UltraSPARC are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

X Window System is a trademark of X Consortium, Inc.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN. THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAMS(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.



Contents

1. Introduction to the API	1
The Programming Interface	1
Interface Functions	2
Libraries	2
Static libraries.....	2
Dynamic libraries.....	3
Interface Taxonomy	3
Standard Classification	4
Public Classification	4
Classification	4
Obsolete Classification	5
2. Processes.....	7
Overview.....	7
Functions.....	8
Spawning new processes	9

Runtime Linking	11
Process Scheduling	13
Error Handling.....	14
Signals	14
Overview	14
3. Process Scheduler.....	23
Overview of the Process Scheduler.....	25
Time-Sharing Class	26
System Class.....	27
Real-time Class.....	27
Commands and Functions	27
The <code>priocntl</code> Command	30
The <code>priocntl</code> Function.....	34
The <code>priocntlset</code> Function	44
Interaction with Other Functions.....	46
Kernel Processes	46
<code>fork</code> and <code>exec</code>	47
<code>nice</code>	47
<code>init</code>	47
Performance	47
Process State Transition.....	48
Software Latencies.....	50
4. Input/Output Interfaces.....	51
Files and I/O	51

Basic File I/O	52
Advanced File I/O	53
File System Control	54
File and Record Locking	55
Supported File Systems	55
Choosing A Lock Type	55
Terminology	56
Setting a File Lock	56
Opening a File for Record Locking	58
Setting and Removing Record Locks.	58
Getting Lock Information	59
Forking Locks.	60
Deadlock Handling	61
Selecting Advisory or Mandatory Locking	61
Cautions about Mandatory Locking	62
Terminal I/O	63
5. System V IPC.	65
Permissions.	66
IPC Functions, Key Arguments, and Creation Flags.	67
Messages	68
Structure of a Message Queue	69
Initializing a Message Queue with <code>msgget()</code>	71
Controlling Message Queues with <code>msgctl()</code>	72
Sending and Receiving Messages	73

Semaphores	75
Structure of a Semaphore Set	77
Initializing a Semaphore Set with <code>semget()</code>	78
Controlling Semaphores with <code>semctl()</code>	80
Performing Semaphore Operations with <code>semop()</code>	82
System V Shared Memory	83
Structure of a System V Shared Memory Segment	84
Using <code>shmget()</code> to Access a Shared Memory Segment....	85
Controlling a Shared Memory Segment with <code>shmctl()</code> ...	87
Attaching and Detaching a Shared Memory Segment with <code>shmat()</code> and <code>shmdt()</code>	88
6. Memory Management	91
Overview of the Virtual Memory System.....	91
Virtual Memory, Address Spaces, and Mapping	91
Networking, Heterogeneity, and Coherence	93
Memory Management Interfaces	94
Creating and Using Mappings.....	94
Removing Mappings.....	100
Cache Control.....	100
Other Mapping Functions	103
Address Space Layout	103
7. Realtime Programming and Administration	107
Basic Rules of Realtime Applications	107
Degrading Response Time	108

Runaway Realtime Processes	111
I/O Behavior	111
Scheduling	112
Dispatch Latency	112
System Calls That Control Scheduling	121
Utilities that Control Scheduling	122
Configuring Scheduling	124
Memory Locking	127
Overview	127
High Performance I/O	129
Asynchronous I/O	129
Synchronized I/O	131
Interprocess Communication	133
Overview	133
Signals	134
Pipes	134
IPC Message Queues	135
IPC Semaphores	136
Shared Memory	136
Choice of IPC Mechanism	138
Asynchronous Networking	138
Modes of Networking	138
Networking Programming Models	139
Asynchronous Connectionless-Mode Service	140

Asynchronous Connection-Mode Service	141
Asynchronous Open	143
Timers	145
Timestamp Functions	145
Interval Timer Functions.....	146
A. Full Code Examples	149
Index.....	181

Tables

Table 2-1	Process Functions	8
Table 2-2	Signal Functions.....	15
Table 3-1	Valid <code>priocntl</code> ID Types.....	30
Table 3-2	Valid <i>idtype</i> Values	31
Table 3-3	Class-Specific Options for <code>priocntl</code>	32
Table 3-4	Valid <code>priocntl.h</code> <i>idtypes</i>	35
Table 3-5	Valid <code>cmd</code> Values.....	35
Table 3-6	What <code>PC_GETPARMS</code> Returns	41
Table 3-7	Special Values for <code>rt_tqnsecs</code>	44
Table 4-1	Basic File I/O Functions.....	52
Table 4-2	Advanced File I/O Functions	53
Table 4-3	File System Control Functions.....	54
Table 4-4	Terminal I/O Functions	63
Table 5-1	IPC Reference Manual Pages.....	66
Table 5-2	Octal Permission Values.....	67
Table 7-1	Realtime System Dispatch Latency with SunOS 5.x.....	115

Table 7-2	Class Options for the dispadmin(1M) Utility	123
-----------	---	-----

Figures

Figure 3-1	SunOS 5.x Process Scheduler	25
Figure 3-2	Process Priorities (Programmer's View)	28
Figure 3-3	Process State Transition Diagram	48
Figure 5-1	Structure of a Message Queue	69
Figure 5-2	Structure of a Semaphore	77
Figure 5-3	Structure of a Shared Memory Segment	84
Figure 5-4	Synopsis of <code>shmctl()</code>	87
Figure 6-1	Traditional UNIX System Address-Space Layout	104
Figure 6-2	Address-Space Layout	105
Figure 7-1	Unbounded Priority Inversion	110
Figure 7-2	Application Response Time	113
Figure 7-3	Internal Dispatch Latency	114
Figure 7-4	Dispatch Priorities for Scheduling Classes	116
Figure 7-5	The Kernel Dispatch Queue	118

Code Samples

Code Example 4-1	Basic file I/O	52
Code Example 4-2	Seek code	53
Code Example 5-1	IPC Permissions Data Structure	66
Code Example 5-2	Message Queue Control Structure	70
Code Example 5-3	Message Header Structure.....	71
Code Example 5-4	Example of msgget() call	71
Code Example 5-5	Example msgctl() calls.....	73
Code Example 5-6	Example msgsnd() and msgrcv() calls.....	74
Code Example 5-7	Synopsis of semget().....	79
Code Example 5-8	Example semget() call.....	79
Code Example 5-9	Example semctl() call.....	81
Code Example 5-10	Example semop() call.....	83
Code Example 5-11	Sample Program to Illustrate shmget().....	86
Code Example 5-12	Sample shmctl() call.....	87
Code Example 5-13	Sample shmat() and shmdt() calls.....	89
Code Example 7-1	Asynchronous Network Transfers	140

Code Example 7-2	Asynchronous Connection Request	142
Code Example 7-3	Asynchronous Listening	143
Code Example 7-4	File Descriptor Transfer	144
Code Example 7-5	Controlling Timer Interrupts.	147
Code Example A-1	Sample Program to Illustrate <code>msgget()</code>	149
Code Example A-2	Sample Program to Illustrate <code>msgctl()</code>	150
Code Example A-3	Sample Program to Illustrate <code>msgsnd()</code> and <code>msgrcv()</code> .	153
Code Example A-4	Sample Program to Illustrate <code>semget()</code>	157
Code Example A-5	Sample Program to Illustrate <code>semctl()</code>	159
Code Example A-6	Sample Program to Illustrate <code>semop()</code>	164
Code Example A-7	Sample Program to Illustrate <code>shmget()</code>	168
Code Example A-8	Sample Program to Illustrate <code>shmctl()</code>	170
Code Example A-9	Sample Program to Illustrate <code>shmat()</code> and <code>shmdt()</code> .	173
Code Example A-10	Example of Record Locking With Lock Promotion . .	178
Code Example A-11	Record Write Locks With <code>lockf()</code>	179

Preface

Purpose

Read this guide for information about system interfaces provided by SunOS libraries. Rather than teaching you to write programs, this guide supplements programming texts by concentrating on other elements that are part of getting programs into operation.

Audience and Prerequisite Knowledge

This guide addresses programmers. Expert programmers, such as those developing system software, might find that this guide lacks the depth of information they need. Expert programmers should see the *Solaris 2.5 Reference Manual AnswerBook*.

Knowledge of terminal use, of a UNIX system editor, and of the UNIX system directory and file structure is assumed. Read the *Solaris User's Guide* to review these basic tools and concepts.

The C Connection

The SunOS system supports many programming languages. Nevertheless, the relationship between this operating system and C has always been and remains very close.

Most of the code in the operating system is written in the C language. So, while this guide is intended to be useful to you no matter what language you are using, most of the examples assume you are programming in C.

Hardware And Software Dependency

Except for hardware-specific information such as addresses, most of the text in this book applies to any computer running the Solaris 2.x operating system.

If commands work differently in your system environment, your system might be running a different software release. If some commands do not seem to exist, they might be in packages that are not installed on your system—talk to your system administrators to find out what commands you have available.

Typeface Conventions

The following conventions are used in this guide:

- Prompts and error messages from the system are printed in **listing type like this**.
- Information you type as a command or in response to prompts is shown in **boldface listing type like this**. Type everything shown in boldface exactly as it appears in the text.
- Parts of a command shown in *italic text like this* refer to a variable that you have to substitute from a selection. It is up to you to make the correct substitution.
- Output to the screen by the system or an application are in **currier**, inputs from the keyboard are in **currier bold**:

```
$pwd  
/home/traveler/scotty
```

- You are expected to press the RETURN key after entering a command or menu choice, so the RETURN key is not explicitly shown in these cases. If, however, you are expected to press RETURN without typing any text, the notation is shown.
- Control characters are shown by the string “CTRL-” followed by the appropriate character, such as D (this is known as CTRL-D). To enter a control character, hold down the key marked CTRL (or CONTROL) and press the D key.

-
- The default prompt signs for an ordinary user and `root` are the dollar sign or percent sign (\$ or %) and the number sign (#). When the # prompt is used in an example, the command illustrated can be executed only by `root`.

Command References

When a command is mentioned in a section of the text for the first time, a reference to the manual section where the command is formally described is included in parentheses: `command(section)`. Numbered sections are in the *Solaris 2.5 Reference Manual AnswerBook*.

For example, “See `priocntl(2)`” tells you to look at the `priocntl` page in section 2 of the *Solaris 2.5 Reference Manual AnswerBook*.

Information in the Examples

While every effort has been made to present displays of information just as they appear on your terminal, it is possible that your system might produce slightly different output. Some displays depend on a particular machine configuration that might differ from yours.

Introduction to the API



A SunSoft goal is to define the Architectural Interfaces of Solaris. There are two reasons:

- The system interface is an effective “contract” with our customers. We tell our customers exactly what we offer them to use, and help ensure that only the official (intended) interface is used.
- We use the definition to ensure that we honor this contract. Interfaces we offer in a particular version of the product are preserved in future versions of the product. Thus, we maintain upward compatibility in subsequent releases of Solaris.

The Programming Interface

Solaris offers many kinds of “interface”, such as: the programming interface, elements of the user interface, protocols, and rules about naming and the locations of objects in the file system. One of the most important interfaces to the system is the programming interface - the one offered to developers. The programming interface has two major parts: one seen by developers of applications, which we call the API, and one seen by developers of system components such as device drivers and platform support modules, which we call the SPI (system programming interface).

Each programming interface to Solaris is also “visible” to the developer at two levels, source level and binary. When we use the acronyms API and SPI, we indicate the source level programming interface to the system. We use the terms Application Binary Interface (ABI) and System Binary Interface (SBI) to

indicate the binary interfaces corresponding to the respective source level programming interfaces. (Because the phrase “the ABI” can be confused with other binary interfaces, we refer to the “Solaris ABI” only by name.)

Interface Functions

The Solaris 2.x functions discussed in this manual are the interfaces between the services provided by the kernel and application programs. The functions described in Sections 2 and 3 of the *Solaris 2.5 Reference Manual AnswerBook* are an application’s interface to the Solaris 2.x operating system. These functions are how an application uses facilities such as the file system, interprocess communication primitives, and multitasking mechanisms. This manual is one of a set that describe major elements of the API. Other manuals in the set are *STREAMS Programming Guide*, *Multithreaded Programming Guide*, *Transport Interfaces Programming Guide*, etc.

When you use the library routines described in sections 2 and 3 of the *Solaris 2.5 Reference Manual AnswerBook*, the details of their implementation are transparent to the program. For example, the function `read` underlies the `fread` implementation in the standard C library.

A C program is automatically linked to the invoked functions when you compile the program. The procedure might be different for programs written in other languages. See the *Linker and Libraries Guide* for more information.

Libraries

Solaris provides both static and dynamic implementations of libraries. Static libraries do not provide an interface, they provide only an implementation. The application programming interface of Solaris is made available to developers through the shared libraries (also called shared objects). In the runtime environment, a dynamic executable and shared objects are processed by the runtime linker, to produce a runnable process. The official API to the system is the interface between an application and the dynamic shared libraries.

Static libraries

The traditional, static, implementation of libraries (.a files or archives), do not separate the application programming interface from its implementation (the contents of the library). When an application is linked to a static library, the

object code that implements that library is bound into the executable object resulting from the build. The source-level programming interface to the library may be preserved, but the application must be relinked to produce an executable that runs on a later version of an operating system. Future binary compatibility is only assured when shared libraries are used.

The presence of static libraries is a historical artifact and there is no mechanism to define their interfaces in a way that is separate from their implementation. For this reason, use of static libraries should be avoided by new applications.

Dynamic libraries

Unlike the static libraries, shared libraries do separate the application programming interface from the implementation. The interface is bound to an implementation of the library only at runtime. This allows SunSoft to evolve the library's implementation - such as changing internal interfaces, while maintaining the API and preserving binary compatibility with applications built against it.

Interface Taxonomy

The Interface Taxonomy classifies commitment level of an interface. The commitment level identifies who may, or how to, use the interface. Definitions:.

Open specification	An interface specification that we publish, which customers can use freely (build products that use our implementation of the interface). Others are free to provide alternative implementations without licensing or legal restrictions.
Closed specification	An interface specification that we do not publish. On which we do not want customers to build products and of which we do not want others to build alternative implementations.
Compatible change	A change to an interface or its implementation has no effect on previously valid programs.
Incompatible change	A change to an interface or its implementation that makes previously valid programs invalid. This may include bug fixes or performance degradation. This does not include programs which depend on unspecified "artifacts of the implementation".

Standard Classification

Specification:	Open
Incompatible Change:	major release (X.0)
Examples:	POSIX, ANSI-C, Solaris ABI, SCD, SVID, XPG, X11, DK1, VMEbus, Ethernet

Standard interfaces are those whose specification is controlled by a group outside of Sun. This includes standards such as POSIX and ANSI C, as well as industry specifications from groups such as X/Open, the MIT X-Consortium, and the OMG.

Public Classification

Specification:	Open
Incompatible Change:	major release (X.0)
Examples:	Sun DDI, XView, ToolTalk, NFS protocol, Sbus, OBP

These are interfaces whose specification is completely under Sun's control. We publish the specification of these interfaces and commit to remain compatible with them.

Classification

Specification:	Open
Incompatible Change:	minor release (X.Y)
Examples:	VFS interface, vm kernel interfaces, libkvm

Uncommitted interfaces are available for use by customers, but they lack the commitment that comes with a Public interface. We often publish specifications of the current versions of these interfaces in a "this is how it works in this release, but we may change it next release" form. Use these interfaces at your own risk in experiments.

Some of these interfaces are ones we would like to elevate to “Public” status, but at the moment we don’t feel confident enough in them to commit to the compatibility constraints of a Public interface. The VFS and vm interfaces are examples of this.

Obsolete Classification

Specification:	None
Incompatible Change:	Minor release (.X.0)
Examples:	RFS

An interface no longer in general use. An existing interface can be downgraded from some other status (such as Public or Standard) to Obsolete through a standard proactive program to communicate the change in commitment to customers.

A change in commitment requires one year’s notice to the customer base and the Sun product development community of the intended obsoleting of the interface. A full year must elapse before delivering a product that contains a change incompatible with the present status of the interface.

Acceptable means of customer notice includes letters to customers on support contracts, release notes or product documentation, or announcements to customer forums appropriate for the interface in question.

The notice of obsolescence is considered to be “public” information in that it is freely available to the customers. It is not intended that this require specific actions to “publish” the information, such as press releases or similar forms of publicity.

Realtime Programming and Administration



This chapter describes writing and porting realtime applications to run under SunOS 5.x. This chapter is written for programmers experienced in writing realtime applications and administrators familiar with realtime processing and the SunOS system.

Basic Rules of Realtime Applications

Realtime response is guaranteed when certain conditions are met. This section identifies these conditions and some of the more significant design errors that can cause problems or disable a system.

Most of the potential problems described here can degrade the response time of the system. One of the potential problems can freeze a workstation. Other, more subtle mistakes are priority inversion and system overload (too much to do).

A SunOS realtime process:

- runs in the RT scheduling class, as described in “Scheduling” on page 112
- locks down all the memory in its process address space, as described in “Memory Locking” on page 127
- is from a statically-linked program or from a program in which all dynamic binding is completed early, as described in “Shared Libraries” on page 109

Realtime operations are described in this chapter in terms of single-threaded processes, but the description can also apply to multithreaded processes (for detailed information about multithreaded processes, see the *Multithreaded Programming Guide*). To guarantee realtime scheduling of a thread, it must be created as a bound thread, and the thread's LWP must be run in the RT scheduling class. The locking of memory and early dynamic binding is effective for all threads in a process.

When a process is the highest priority realtime process, it:

- acquires the processor within the guaranteed dispatch latency period of becoming runnable (see “Dispatch Latency” on page 112)
- continues to run for as long as it remains the highest priority runnable process

A realtime process can lose control of the processor or can be unable to gain control of the processor because of other events on the system. These events include external events (such as interrupts), resource starvation, waiting on external events (synchronous I/O), and preemption by a higher priority process.

Realtime scheduling generally does not apply to system initialization and termination services such as `open(2)` and `close(2)`.

Degrading Response Time

The problems described in this section all increase the response time of the system to varying extents. The degradation can be serious enough to cause an application to miss a critical deadline.

Realtime processing can also significantly impact the operation of aspects of other applications active on a system running a realtime application. Since realtime processes have higher priority, time-sharing processes can be prevented from running for significant amounts of time. This can cause interactive activities, such as displays and keyboard response time, to be noticeably slowed.

System Response Time

System response under SunOS 5.x provides no bounds to the timing of I/O events. This means that synchronous I/O calls should never be included in any program segment whose execution is time critical. Even program segments that permit very large time bounds must not perform synchronous I/O. Mass storage I/O is such a case, where causing a read or write operation hangs the system while the operation takes place.

A common application mistake is to perform I/O to get error message text from disk. This should be done from an independent non-realtime process or thread.

Interrupt Servicing

Interrupt priorities are independent of process priorities. Prioritizing processes does not carry through to prioritizing the services of hardware interrupts that result from the actions of the processes. This means that interrupt processing for a device controlled by a realtime process is not necessarily done before interrupt processing for another device controlled by a timeshare process.

Shared Libraries

Time-sharing processes can save significant amounts of memory by using dynamically linked, shared libraries. This type of linking is implemented through a form of file mapping. Dynamically linked library routines cause implicit reads.

Realtime programs can use shared libraries, yet avoid dynamic binding, by setting the environment variable `LD_BIND_NOW` to a non-NULL value when the program is invoked. This forces all dynamic linking to be bound before the program begins execution. See the *Linker and Libraries Guide* for more information.

Priority Inversion

A time-sharing process can block a realtime process by acquiring a resource that is required by a realtime process. Priority inversion is a condition that occurs when a higher priority process is blocked by a lower priority process.

The term *blocking* describes a situation in which a process must wait for one or more processes to relinquish control of resources. If this blocking is prolonged, even for lower level resources, deadlines might be missed.

By way of illustration, consider the case in Figure 7-1 where a high priority process wanting to use a shared resource gets blocked when a lower priority process holds the resource, and the lower priority process is preempted by an intermediate priority process. This condition can persist for a long time, arbitrarily long, in fact, since the amount of time the high priority process must wait for the resource depends not only on the duration of the critical section being executed by the lower priority process, but on the duration until the intermediate process blocks. Any number of intermediate processes can be involved.




Figure 7-1 Unbounded Priority Inversion

Sticky Locks

A page is permanently locked into memory when its lock count reaches 65535 (0xFFFF). The value 0xFFFF is implementation-defined and might change in future releases. Pages locked this way cannot be unlocked.

Runaway Realtime Processes

Runaway realtime processes can cause the system to halt or can slow the system response so much that the system appears to halt.

Note – If you have a runaway process on a SPARC system, try typing (Stop-A). You might have to repeat this procedure many times. If this doesn't work, disconnect the keyboard.

When a high priority realtime process will not relinquish control of the CPU, there is no simple way to regain control of the system until the infinite loop is forced to terminate. Such a runaway process will not respond to the control-C kill sequence.

Caution – Attempts to use a shell set at a higher priority than a runaway process will not succeed. The STREAMS processes that govern `tty` management are running at system priority, and so will not get scheduled. Therefore, keyboard input is not received by the shell, even when the shell is running at a higher priority.

I/O Behavior

Asynchronous I/O

There is no guarantee that asynchronous I/O operations will be done in the sequence in which they are queued to the kernel. Nor is there any guarantee that asynchronous operations will be returned to the caller in the sequence in which they were done.

If a single buffer is specified for a rapid sequence of calls to `aioread(3)`, there is no guarantee about the state of the buffer between the time that the first call is made and the time that the last result is signaled to the caller.

An individual `aio_result_t` structure can be used only for one asynchronous read or write at a time.

Realtime Files

SunOS 5.x provides no facilities to assure that files will be allocated as physically contiguous.

For regular files, the `read()` and `write()` operations are always buffered. An application can use `mmap()` and `msync()` to effect direct I/O transfers between secondary storage and process memory.

Scheduling

Realtime scheduling constraints are necessary to manage data acquisition or process control hardware. The realtime environment requires that a process be able to react to external events in a bounded amount of time. Such constraints can exceed the capabilities of a kernel designed to provide a “fair” distribution of the processing resources to a set of time-sharing processes.

This section describes the SunOS 5.x realtime scheduler, its priority queue, and how to use system calls and utilities that control scheduling. For more information about the functions described in this section, see the *man Pages(3): Library Routines*.

Dispatch Latency

The most significant element in scheduling behavior for realtime applications is the provision of a real-time scheduling class. The standard time-sharing scheduling class is not suitable for realtime applications because this scheduling class treats every process equally and has a limited notion of priority. Realtime applications require a scheduling class in which process priorities are taken as absolute and are changed only by explicit application operations.

The term dispatch latency describes the amount of time it takes for a system to respond to a request for a process to begin operation. With a scheduler written specifically to honor application priorities, realtime applications can be developed with a bounded dispatch latency.

Figure 7-2 illustrates the amount of time it takes an application to respond to a request from an external event.




Figure 7-2 Application Response Time

The overall application response time is composed of the interrupt response time, the dispatch latency, and the time it takes the application itself to determine its response.

The interrupt response time for an application includes both the interrupt latency of the system and the device driver's own interrupt processing time. The interrupt latency is determined by the longest interval that the system must run with interrupts disabled; this is minimized in SunOS 5.x using synchronization primitives that do not commonly require a raised processor interrupt level.

During interrupt processing, the driver's interrupt routine wakes up the high priority process and returns when finished. The system detects that a process with higher priority than the interrupted process is now dispatchable and arranges to dispatch that process. The time to switch context from a lower priority process to a higher priority process is included in the dispatch latency time.

Figure 7-3 illustrates the internal dispatch latency / application response time of a system, defined in terms of the amount of time it takes for a system to respond to an internal event. The dispatch latency of an internal event represents the amount of time required for one process to wake up another higher priority process, and for the system to dispatch the higher priority process.

The application response time is the amount of time it takes for a driver to wake up a higher priority process, have a low priority process release resources, reschedule the higher priority task, calculate the response, and dispatch the task.

Note – Interrupts can arrive and be processed during the dispatch latency interval. This processing increases the application response time, but is not attributed to the dispatch latency measurement, and so is not bounded by the dispatch latency guarantee.




Figure 7-3 Internal Dispatch Latency

With the new scheduling techniques provided with realtime SunOS 5.x, the system dispatch latency time is within specified bounds.

As you can see in Table 7-1, dispatch latency improves with a bounded number of processes.

Table 7-1 Realtime System Dispatch Latency with SunOS 5.x

Workstation	Dispatch Latency	
	Bounded Number of Processes	Arbitrary Number of Processes
SPARCstation 1	< 2.0 milliseconds in a system with fewer than 8 active processes	4.5 milliseconds
SPARCstation 1+	< 2.0 milliseconds in a system with fewer than 8 active processes	4.0 milliseconds
SPARCstation IPX	< 1.0 milliseconds in a system with fewer than 8 active processes	2.2 milliseconds
SPARCstation 2	< 1.0 milliseconds in a system with fewer than 16 active processes	2.0 milliseconds

Tests for dispatch latency and experience with such critical environments as manufacturing and data acquisition have proven that the Sun workstation is an effective platform for the development of realtime applications. (We apologize that the examples are not of current products.)

Scheduling Classes

The SunOS 5.x kernel dispatches processes by priority. The scheduler (or dispatcher) supports the concept of scheduling classes. Classes are defined as Realtime (RT), System (SYS), and Time-Sharing (TS). Each class has a unique scheduling policy for dispatching processes within its class.

The kernel dispatches highest priority processes first. By default, realtime processes have precedence over sys and TS processes, but administrators can configure systems so that TS and RT processes have overlapping priorities.

Figure 7-4 illustrates the concept of classes as viewed by the SunOS 5.x kernel.




Figure 7-4 Dispatch Priorities for Scheduling Classes

At highest priority are the hardware interrupts; these cannot be controlled by software. The interrupt processing routines are dispatched directly and immediately from interrupts, without regard to the priority of the current process.

Realtime processes have the highest default software priority. Processes in the RT class have a priority and *time quantum* value. RT processes are scheduled strictly on the basis of these parameters. As long as an RT process is ready to run, no SYS or TS process can run. Fixed priority scheduling allows critical processes to run in a predetermined order until completion. These priorities never change unless an application changes them.

An RT class process inherits the parent's time quantum, whether finite or infinite. A process with a finite time quantum runs until the time quantum expires or the process terminates, blocks (while waiting for an I/O event), or is

preempted by a higher priority runnable realtime process. A process with an infinite time quantum ceases execution only when it terminates, blocks, or is preempted.

The `sys` class exists to schedule the execution of special system processes, such as paging, `STREAMS`, and the swapper. It is not possible to change the class of a process to the `sys` class. The `sys` class of processes has fixed priorities established by the kernel when the processes are started.

At lowest priority are the time-sharing (TS) processes. TS class processes are scheduled dynamically, with a few hundred milliseconds for each time slice. The TS scheduler switches context in round-robin fashion often enough to give every process an equal opportunity to run, depending upon its time slice value, its process history (when the process was last put to sleep), and considerations for CPU utilization. Default time-sharing policy gives larger time slices to processes with lower priority.

A child process inherits the scheduling class and attributes of the parent process through `fork(2)`. A process' scheduling class and attributes are unchanged by `exec(2)`.

Different algorithms dispatch each scheduling class. Class dependent routines are called by the kernel to make decisions about CPU process scheduling. The kernel is class-independent, and takes the highest priority process off its queue. Each class is responsible for calculating a process' priority value for its class. This value is placed into the dispatch priority variable of that process.

As Figure 7-5 illustrates, each class algorithm has its own method of nominating the highest priority process to place on the global run queue.




Figure 7-5 The Kernel Dispatch Queue

Each class has a set of priority levels that apply to processes in that class. A class-specific mapping maps these priorities into a set of global priorities. It is not required that a set of global scheduling priority maps start with zero, nor that they be contiguous.

By default, the global priority values for time-sharing (TS) processes range from -20 to +20, mapped into the kernel from 0-40, with temporary assignments as high as 99. The default priorities for realtime (RT) processes range from 0-59, and are mapped into the kernel from 100 to 159. The kernel's class-independent code runs the process with the highest global priority on the queue.

Dispatch Queue

The dispatch queue is a linear linked list of processes with the same global priority. Each process is invoked with class specific information attached to it. A process is dispatched from the kernel dispatch table based upon its global priority.

Dispatching Processes

When a process is dispatched, the process' context is mapped into memory along with its memory management information, its registers, and its stack. Then execution begins. Memory management information is in the form of hardware registers containing data needed to perform virtual memory translations for the currently running process.

Preemption

When a higher priority process becomes dispatchable, the kernel interrupts its computation and forces the context switch, preempting the currently running process. A process can be preempted at any time if the kernel finds that a higher priority process is now dispatchable.

For example, suppose that process A performs a read from a peripheral device. Process A is put into the *sleep* state by the kernel. The kernel then finds that a lower priority process B is runnable, so process B is dispatched and begins execution. Eventually, the peripheral device interrupts, and the driver of the device is entered. The device driver makes process A runnable and returns. Rather than returning to the interrupted process B, the kernel now preempts B from processing and resumes execution of the awakened process A.

Another interesting situation occurs when several processes contend for kernel resources. When a lower priority process releases a resource for which a higher priority realtime process is waiting, the kernel immediately preempts the lower priority process and resumes execution of the higher priority process.

Kernel Priority Inversion

Priority inversion occurs when a higher priority process is blocked by one or more lower priority processes for a long time. The use of synchronization primitives such as mutual-exclusion locks in the SunOS 5.x kernel can lead to priority inversion.

The term *blocking* describes the situation in which a process must wait for one or more processes to relinquish resources. If this blocking continues, it can lead to deadlines being missed, even for low levels of utilization.

The problem of priority inversion has been addressed for mutual-exclusion locks for the SunOS 5.x kernel by implementing a basic priority inheritance policy. The policy states that a lower priority process inherits the priority of a higher priority process when the lower priority process blocks the execution of the higher priority process. This places an upper bound on the amount of time a process can remain blocked. The policy is a property of the kernel's behavior, not a solution that a programmer institutes through system calls or function execution. User-level processes can still exhibit priority inversion, however.

User Priority Inversion

There is no mechanism by which processes synchronizing with other processes will automatically inherit the priority of waiting processes. An application can bound its priority inversion by using priority ceiling emulation.

Under this model, the application associates a priority with each synchronization object, which is typically the highest priority of any process that can block on that object.

Each process then uses the following sequence when manipulating the shared resources

```
/*
 * raise process priority to maximum of current level
 * and synchronization object level
 */
...

/*
 * acquire synchronization object
 */
...
```

```
/*
 * execute the critical section
 */
...

/*
 * release synchronized object
 */
...

/*
 * return to previous process priority level
 */
...
```

System Calls That Control Scheduling

System calls implemented for realtime scheduling include the library calls and functions listed in this section. For more detail about using these, see the *man Pages(3): Library Routines*.

Using `priocntl(2)`

Control over scheduling of active classes is handled with `priocntl(2)`. Class attributes are inherited over `fork(2)` and `exec(2)`, along with scheduling parameters and permissions required for priority control. These characteristics are true for both the RT and the TS classes.

The `priocntl(2)` function provides an interface for specifying a realtime process, a set of processes, or a class to which the system call will apply. The `priocntlset(2)` system call also provides the more general interface for specifying an entire set of processes to which the system call is to apply.

The *idtype* and *id* arguments are used together to specify the set of processes on the queue. Depending upon the value of *idtype*, *id* can have values for a single process ID, a parent process ID, a process group ID, a session ID, a class ID, a user ID, a group ID, or a lightweight process ID.

The command arguments of `priocntl` can be one of: `PC_GETCID`, `PC_GETCLINFO`, `PC_GETPARMS`, or `PC_SETPARMS`. The real or effective ID of the calling process must match that of the affected process or processes, or must have super-user privilege.

PC_GETCID

This command takes the name field of a structure that contains a recognizable class name (RT for realtime and TS for time-sharing). The class ID and an array of class attribute data are returned.

PC_GETCLINFO

This command takes the ID field of a structure that contains a recognizable class identifier. The class name and an array of class attribute data are returned.

PC_GETPARMS

This command returns the scheduling class identifier and/or the class specific scheduling parameters of one of the specified processes. Even though `idtype` & `id` might specify a big set, `PC_GETPARMS` returns the parameter of only one process. It is up to the class to select which one.

PC_SETPARMS

This command sets the scheduling class and/or the class specific scheduling parameters of the specified process or processes.

Utilities that Control Scheduling

The administrative utilities that control process scheduling are `dispadmin(1M)` and `priocntl(1)`. Both these utilities support the `priocntl(2)` system call with compatible options and loadable modules. Using these utilities provides system administration functions that control realtime process scheduling during runtime. For more details about using these utilities, see the *man Pages(1): User Commands* and the *System Administration Guide, Volume II* guide.

Using `priocntl(1)`

The `priocntl(1)` command sets and retrieves scheduler parameters for processes.

Using `dispadmin(1M)`

The `dispadmin(1M)` utility displays all current process scheduling classes by including the `-l` command line option during runtime. Process scheduling can also be changed for the class specified after the `-c` option, using `RT` as the argument for the realtime class.

The following options are also available:

Table 7-2 Class Options for the `dispadmin(1M)` Utility

<i>option</i>	<i>meaning</i>
<code>-l</code>	lists scheduler classes currently configured
<code>-c</code>	specifies the class whose parameters are to be displayed or changed
<code>-g</code>	gets the dispatch parameters for the specified class
<code>-r</code>	when using <code>-g</code> , specifies time quantum resolution
<code>-s</code>	specifies a file where values can be located

A class specific file containing the dispatch parameters can also be loaded during runtime. Use this file to establish a new set of priorities replacing the default values established during boot time. This class specific file must assert the arguments in the format used by the `-g` option. Parameters for the `RT` class are found in the `rt_dptbl(4)`, and are listed in the example at the end of this section.

To add an `RT` class file to the system, the following modules must be present:

- An `rt_init()` routine in the class module which loads the `rt_dptbl`.
- A `rt_dptbl` module that provides the dispatch parameters and a routine to return pointers to `config_rt_dptbl`.

- The `dispadmin` executable.

Then load the class specific module with the following command, where `<module_name>` is the class specific module.

```
modload /kernel/sched/<module_name>
```

Then invoke the `dispadmin` command

```
# dispadmin -c RT -s <file_name>
```

The file must describe a table with the same number of entries as the table that is being overwritten.

Configuring Scheduling

Associated with each scheduling class is a parameter table, `config_rt_dptbl` (RT), and `config_ts_dptbl` (TS). These tables are configurable by using a loadable module at boot time, or with `dispadmin(1M)` during runtime.

The Dispatcher Parameter Table

The in-core table for realtime establishes the properties for RT scheduling. The `config_rt_dptbl` structure consists of an array of parameters, `struct rt_dpent`, one for each of the `n` priority levels. The properties of a given priority level `i` are specified by the `i`th parameter structure in the array, `config_rt_dptbl[i]`.

A parameter structure consists of the following members (also described in the `/usr/include/sys/rt.h` header file):

`rt_globpri`

The global scheduling priority associated with this priority level. The `rt_globpri` values cannot be changed with `dispadmin(1M)`.

`rt_quantum`

The length of the time quantum allocated to processes at this level in ticks (see “Timestamp Functions” on page 145). The time quantum value is only a default or starting value for processes at a particular level. The time quantum of a realtime process can be changed by using the `priocntl(1)` command or the `priocntl(2)` system call.

Reconfiguring config_rt_dptbl

A realtime administrator can change the behavior of the realtime portion of the scheduler by reconfiguring the `config_rt_dptbl` at any time. Two methods are described here.

The first method is to reconfigure the `config_rt_dptbl` parameter table with a loadable module which contains a new dispatch table loaded at boot time. The module containing the dispatch table is a separate module. This is the only method that can be used to change the number of realtime priority levels or the set of global scheduling priorities used by the realtime class. Note that changing the `config_rt_dptbl` affects the realtime processes that you set after the table gets updated.

A second method for examining or modifying the realtime parameter table on a running system is through using the `dispadmin(1M)` command. Invoking `dispadmin` for the realtime class allows retrieval of the current `rt_quantum` values in the current `config_rt_dptbl` configuration from the kernel's in-core table. When overwriting the current in-core table, the configuration file used for input to `dispadmin` must conform to the specific format described in the manual page for `config_rt_dptbl` found in the *man Pages(1M): System Administration Commands*.

Following is an example of prioritized processes `rtdpent_t` with their associated time quantum `config_rt_dptbl[]` value as they might appear in `config_rt_dptbl[]`:

<code>rtdpent_t</code>	<code>rt_dptbl[] = {</code>	
	<code>/* prilevel Time quantum */</code>	
	100, 100,	134, 40,
	101, 100,	135, 40,
	102, 100,	136, 40,
	103, 100,	137, 40,
	104, 100,	138, 40,
	105, 100,	139, 40,
	106, 100,	140, 20,
	107, 100,	141, 20,
	108, 100,	142, 20,
	109, 100,	143, 20,
	110, 80,	144, 20,
	111, 80,	145, 20,
	112, 80,	146, 20,
	113, 80,	147, 20,
	114, 80,	148, 20,
	115, 80,	149, 20,
	116, 80,	150, 10,
	117, 80,	151, 10,
	118, 80,	152, 10,
	119, 80,	153, 10,
	120, 60,	154, 10,
	121, 60,	155, 10,
	122, 60,	156, 10,
	123, 60,	157, 10,
	124, 60,	158, 10,
	125, 60,	159, 10,
	126, 60,	
	127, 60,	
	128, 60,	
	129, 60,	
	130, 40,	
	131, 40,	
	132, 40,	
	133, 40,	
	}	

Memory Locking

Locking memory is one of the most important issues for realtime applications. In a realtime environment, a process must be able to guarantee continuous memory residence to reduce latency and to prevent paging and swapping.

This section describes the memory locking mechanisms available to realtime applications in SunOS 5.x. For more details about using memory management functions and calls, see the *man Pages(3): Library Routines* for pertinent manual pages.

Overview

Under SunOS 5.x, the memory residency of a process is determined by its current state, the total available physical memory, the number of active processes, and the processes' demand for memory. This is appropriate in a time-share environment, but it is often unacceptable for a realtime process. In a realtime environment, a process must be able to guarantee memory residence for all or part of itself to reduce its memory access and dispatch latency.

For realtime in SunOS 5.x, memory locking is provided by a set of library routines that allow a process running with superuser privileges to lock specified portions of its virtual address space into physical memory. Pages locked in this manner are exempt from paging until they are unlocked or the process exits.

There is a system-wide limit on the number of pages that can be locked at any time. This is a tunable parameter whose default value is calculated at boot time. It is based on the number of page frames less another percentage (currently set at ten percent).

Locking a Page

A call to `mlock(3)` requests that one segment of memory be locked into the system's physical memory. The pages that make up the specified segment are faulted in and the lock count of each is incremented. Any page with a lock count greater than 0 is exempt from paging activity.

A particular page can be locked multiple times by multiple processes through different mappings. If two different processes lock the same page, the page remains locked until both processes remove their locks. However, within a given mapping, page locks do not nest. Multiple calls of locking functions on the same address by the same process are removed by a single unlock request.

If the mapping through which a lock has been performed is removed, the memory segment is implicitly unlocked. When a page is deleted through closing or truncating the file, it is also unlocked implicitly.

Locks are not inherited by a child process after a `fork(2)` call is made. So, if a process with memory locked forks a child, the child must perform a memory locking operation in its own behalf to lock its own pages. Otherwise, the child process incurs copy-on-write page faults, which are the usual penalties associated with forking a process.

Unlocking a Page

To unlock a page of memory, a process requests that a segment of locked virtual pages be released by a call to `munlock(3)`. The lock counts of the specified physical pages are decremented. Once the lock count of a page has been decremented to 0, the page is swapped normally.

Locking All Pages

A superuser process can request that all mappings within its address space be locked by a call to `mlockall(3)`. If the flag `MCL_CURRENT` is set, all the existing memory mappings are locked. If the flag `MCL_FUTURE` is set, every mapping that is added to or that replaces an existing mapping is locked into memory.

Sticky Locks

A page is permanently locked into memory when its lock count reaches 65535 (0xFFFF). The value 0xFFFF is implementation defined and might change in future releases. Pages locked in this manner cannot be unlocked. Reboot the system to recover.

High Performance I/O

This section describes I/O with realtime processes. With SunOS 5.x, several functions and calls are available within the libraries supplied to perform fast, asynchronous I/O operations. For robustness, SunOS provides file synchronization operations and modes to prevent information loss and data inconsistency. See the *man Pages(3): Library Routines* for more detailed information.

Asynchronous I/O

Standard UNIX I/O is generally synchronous to the application programmer. An application that calls `read(2)` or `write(2)` is not usually allowed to proceed until that system call has finished, successfully or otherwise.

Realtime applications need *asynchronous* bounded I/O behavior. A process that issues an asynchronous I/O call does not wait until the I/O operation has been completed before it is allowed to proceed. Instead, the caller is notified that the I/O operation has finished at a later time while the process is doing something else.

Asynchronous I/O applies to any SunOS file. Files are opened in the synchronous way and no special flagging is required. An asynchronous I/O transfer is composed of three elements: call, request, and operation. The application calls an asynchronous I/O function, the request for the I/O is placed on a queue, and the call returns immediately. At some point, the system dequeues the request and initiates the I/O operation itself.

Asynchronous and standard I/O requests can be intermingled on any file descriptor. Note, however, that the system does not necessarily maintain any particular sequence of read and write requests. The system can and does arbitrarily resequence any and all pending read and write requests. If a specific sequence is required for the application, the application must assure the completion of prior operations before issuing the dependent requests.

Notification (SIGIO)

When an asynchronous I/O call returns successfully, the I/O operation has only been placed on the queue, waiting to be done. The actual operation also has a return value and a potential error identifier. These are the values that would have been returned to the caller as the result of a synchronous call.

When the I/O is finished, the return value and error value are stored at a location given by the user at the time of the request as a pointer to an `aio_result_t`. The structure of the `aio_result_t` is defined in `<sys/asynch.h>`

```
typedef struct aio_result_t {
    int aio_return; /* return value of read or write */
    int aio_errno; /* errno generated by the IO */
} aio_result_t;
```

When `aio_result_t` has been updated, a `SIGIO` signal is delivered to the process that made the I/O request.

Note that a person with two or more asynchronous I/O operations pending has no certain way to determine which request or even whether either request is the cause of the `SIGIO` signal. A process receiving a `SIGIO` should check all its conditions which could be generating the `SIGIO` signal.

Using aioread(3)

The `aioread(3)` function is the asynchronous version of `read(2)`. In addition to the normal read arguments, `aioread` takes the arguments specifying a file position and the address of an `aio_result_t` structure at which the system is to store the result information about the operation.

The file position specifies a seek to be performed within the file before the operation. If the `aioread` call succeeds, the file pointer is updated to the position that would have resulted in a successful seek and read. The file pointer is also updated when a read fails to allow for subsequent read requests.

Using aiowrite(3)

The `aiowrite(3)` function is the asynchronous version of `write(2)`. In addition to the normal write arguments, `aiowrite` takes arguments specifying a file position and the address of an `aio_result_t` structure at which the system is to store the result information about the operation.

The file position specifies a seek to be performed within the file before the operation. If the `aiowrite` call succeeds, the file pointer is updated to the position that would have resulted in a successful seek and write. The file pointer is also updated when a write fails to allow for subsequent write requests.

Using aiocancel(3)

The `aiocancel(3)` function attempts to cancel the asynchronous request whose `aio_result_t` structure is given as an argument. An `aiocancel` call succeeds only if the request is still queued. If the operation is in progress, `aiocancel` fails.

Using aiowait(3)

A call to the `aiowait(3)` function blocks the calling process until at least one outstanding asynchronous I/O operation is completed. The timeout parameter points to a maximum interval to wait for I/O completion. A timeout value of zero specifies that no wait is wanted. The `aiowait` function returns a pointer to the `aio_result_t` structure for the completed operation.

Using poll(2)

When you prefer to poll devices rather than to depend on a `SIGIO` interrupt, use the `poll(2)` system call. You can also poll to determine the origin of an `SIGIO` interrupt.

Using close(2)

Files are closed by a call to `close(2)`. The call to `close` cancels any outstanding asynchronous I/O request that can be cancelled. The `close` function waits on an operation that cannot be cancelled. When a call to `close` returns, there is no asynchronous I/O pending for the file descriptor.

Only asynchronous I/O requests that are queued to the specified file descriptor are cancelled when a file is closed. Any I/O requests that are pending for other file descriptors are not cancelled.

Synchronized I/O

Applications may need to guarantee that information has been written to stable storage, or that file updates are performed in a particular order. Synchronized I/O provides for these needs.

Modes of Synchronization

Under SunOS 5.x, data is successfully transferred for a write operation to a regular file when the system ensures that all data written is readable on any subsequent open of the file (even one that follows a system or power failure) in the absence of a failure of the physical storage medium. Data is successfully transferred for a read operation when an image of the data on the physical storage medium is available to the requesting process. An I/O operation is complete when either the associated data been successfully transferred or the operation has been diagnosed as unsuccessful.

An I/O operation has reached synchronized I/O data integrity completion when:

For reads, the operation has been completed or diagnosed if unsuccessful. The read is complete only when an image of the data has been successfully transferred to the requesting process. If there were any pending write requests affecting the data to be read at the time that the synchronized read operation was requested, these write requests are successfully transferred prior to reading the data.

For writes, the operation has been completed or diagnosed if unsuccessful. The write is complete only when the data specified in the write request is successfully transferred, and all file system information required to retrieve the data is successfully transferred.

File attributes that are not necessary for data retrieval (access time, modification time, status change time) are not successfully transferred prior to returning to the calling process.

Synchronized I/O file integrity completion is identical to synchronized I/O data integrity completion with the addition that all file attributes relative to the I/O operation (including access time, modification time, status change time) must be successfully transferred prior to returning to the calling process.

Synchronizing a File

The `fsync(3C)` and `fdatasync(3R)` functions explicitly synchronize a file to secondary storage:

```
int fsync (int fildes);
int fdatasync (int fildes);
```

The `fsync()` guarantees the function is synchronized at the I/O file integrity completion level, while The `fdatasync()` guarantees the function is synchronized at the I/O data integrity completion level.

Applications can arrange that each I/O operation is synchronized before the operation completes. Setting the `O_DSYNC` flag on the file description via `open(2)` or `fcntl(2)` ensures that all I/O writes (`write(2)`, `aiowrite(3)`) have reached I/O data completion before the operation is indicated as completed. Setting the `O_SYNC` flag on the file description ensures that all I/O writes have reached I/O file completion before the operation is indicated as completed. Setting the `O_RSYNC` flag on the file description ensures that all I/O reads (`read(2)`, `aioread(3)`) have reached the same level of completion as request for writes by the setting `O_DSYNC` or `O_SYNC` on the descriptor.

Interprocess Communication

This section describes the interprocess communication (IPC) functions of SunOS 5.x as they relate to realtime processing. Signals, pipes, FIFOs (named pipes), message queues, shared memory, file mapping, and semaphores are described here. For more information about the libraries, functions, and routines useful for interprocess communication, see chapter three, “Interprocess Communication,” and the *man Pages(3): Library Routines*.

Overview

Realtime processing often requires fast, high-bandwidth interprocess communication. The choice of which mechanisms should be used can be dictated by functional requirements, and the relative performance will depend upon application behavior.

The traditional method of interprocess communication in UNIX is the pipe. Unfortunately, pipes can have framing problems. Messages can become intermingled by multiple writers or can be torn apart by multiple readers.

IPC messages mimic the reading and writing of files. They are easier to use than pipes when more than two processes must communicate by using a single medium.

The IPC shared semaphore facility provides process synchronization. Shared memory is the fastest form of interprocess communication. The main advantage of shared memory is that the copying of message data is eliminated. The usual mechanism for synchronizing shared memory access is semaphores.

Signals

Signals may be used to send a small amount of information between processes. The sender can use the `sigqueue(3R)` function to send a signal together with a small amount of information to a target process:

```
int sigqueue(pid_t pid, int signo, const union sigval value);
union sigval {
    int      sival_int;      /* integer value */
    void     *sival_ptr;     /* pointer value */
};
```

The target process must have the `SA_SIGINFO` bit set for the given signal number (see `sigaction(2)`), in order that occurrences of the signal occurring when that signal is already pending will be queued.

The target process can receive the signals either synchronously or asynchronously. By leaving that signal blocked (`sigprocmask(2)`) and calling either `sigwaitinfo(3R)` or `sigtimedwait(3R)`, the signal will be received synchronously, with the value sent by the caller of `sigqueue()` being stored in the `si_value` member of the `siginfo_t` argument. By leaving the signal unblocked, the arrival will be delivered to the signal handler specified by `sigaction()`, with the value appearing in the `si_value` of the `siginfo_t` argument to the handler.

Only a fixed number of signals with associated values can be sent by a process and remain undelivered. Storage for `{SIGQUEUE_MAX}` signals is allocated at the first call to `sigqueue()`. Thereafter, a call to `sigqueue()` either successfully enqueues at the target process or fails within a bounded amount of time.

Pipes

Pipes provide one-way communication between processes. Pipes are created by a process using the `pipe(2)` system call. The `pipe(2)` system call returns two file descriptors, the first for reading and the second for writing. Once the pipe

is created, the process must create other processes with the `fork(2)` system call, which allows the processes to communicate among themselves. Processes must have a common ancestor in order to communicate with pipes.

Data passed through a pipe is treated as a conventional UNIX byte stream. Data is sent into the pipe by calls to `write(2V)` using the writing file descriptor.

Data is received from the pipe by calls to `read(2V)` using the reading file descriptor. The `read` call is usually a blocking function: it does not return to the caller until some data can be returned. To get a non-blocking read, the pipe can be set so that it doesn't block by using the `ioctl(2)` or `fcntl(2)` functions.

A read on an empty, non-blocking pipe returns with an indication that no data is available.

Named Pipes

SunOS 5.x provides *named pipes* or *FIFOs*. The FIFO is more flexible than the pipe because it is a named entity in a directory. Once created, a FIFO can be opened by any process that has legitimate access to it. Processes do not have to share a parent and there is no need for a parent to initiate the pipe and pass it to the descendants. A FIFO can be created with `mknod(2)`.

A process connects to a FIFO through a call to `open(2V)`. A process that opens a FIFO for a read is blocked until that FIFO has been opened by a process for writing. The decision about whether or not reads block is made in the `open` call or by using a subsequent call to `fcntl`.

As with pipes, data in a FIFO is treated as a byte stream. Input is obtained from a FIFO with calls to `read` and output is sent with calls to `write`. A process ends use of a FIFO through a call to `close(2)`.

IPC Message Queues

IPC message queues provide a powerful means of communicating between processes by allowing any number of processes to send and receive from the same message queue. Messages are passed as blocks of arbitrary size, not as byte streams. Each message includes an integer type, which can be used by application convention as a message priority, or as message categories. The

latter usage provides multiple flows of messages with a single message queue. This can be simpler than opening an arbitrary number of pipes or FIFOs when a large number are required. Note that IPC insertion is strictly FIFO.

IPC message queue structures are initiated by a call to `msgget(2)`. A message is sent by a call to `msgsnd(2)`, and `msgrcv(2)` is called to extract a message from the queue structure. The `msgctl(2)` system call controls various functions on a message queue structure, including removal.

IPC Semaphores

The IPC semaphore is a mechanism that synchronizes access to shared resources. IPC semaphores are created in arrays, each element of which can be used to control the execution of processes that call for operations on the array elements.

Create an array of IPC semaphores with a call to `semget(2)`. Query or set individual semaphores or the complete array of semaphores with calls to `semctl(2)`. Acquire and release a semaphore or the array of semaphores with calls to `semop(2)`. Look in `intro(2)` for more information about information structures and the operation of IPC semaphores.

Note that using IPC semaphores can cause priority inversions unless these are explicitly avoided by the techniques mentioned earlier in this chapter.

Shared Memory

The fastest way for processes to communicate is directly, through a shared segment of memory. A common memory area is added to the address space of processes wishing to communicate. Applications use stores to send data and fetches to receive communicated data. SunOS 5.x provides two mechanisms for shared memory: memory mapped files and IPC shared memory.

The major difficulty with shared memory is that results can be wrong when more than two processes are trying to read and write in it at the same time. See “Shared Memory Synchronization” on page 138 for more information.

Memory Mapped Files

The system call `mmap(2)` connects a shared memory segment to the caller's memory. The caller specifies the shared segment by address and length. The caller must also specify access protection flags and how the mapped pages are managed.

The `mmap(2)` system call can also be used to map a file or a segment of a file to a process's memory. While this technique is very convenient in some applications, it is easy to forget that any access to the mapped file segment might result in implicit I/O. This can make an otherwise bounded process have unpredictable response times. The function `msync(3)` forces immediate or eventual copies of the specified memory segment to its permanent storage location(s).

The process can later change the access protection of the segment by the system call `mprotect(2)`. The segment is specified by address and length.

The system call `munmap(2)` disconnects a mapped memory segment. The segment is specified by address and length.

Fileless Memory Mapping

The zero special file, `/dev/zero(4S)`, can be used to create an unnamed, zero initialized memory object. The length of the memory object is the least number of pages that contain the mapping. The object can be shared only by descendants of a common ancestor process.

IPC Shared Memory

A `shmget(2)` call can be used either to create and obtain a shared memory segment or to obtain an existing shared memory segment. The call specifies an identifying key, the size of the segment, and a flag parameter. The flags contain the usual access permission bits and can contain a flag to create a new segment. The `shmget` function returns an identifier that is analogous to a file identifier.

The shared memory segment is made accessible to the process by a call to `shmat(2)`. The shared memory segment becomes a virtual segment of the process memory space and can be freely written to and read from depending on creating permissions. The shared memory segment is detached from a

process's memory space by a call to `shmdt(2)`. The `shmctl` system call can be used to control a variety of functions on an IPC shared memory object, including removal.

Shared Memory Synchronization

In sharing memory, a portion of memory can be mapped into the address space of one or more processes. This allows shared access to that portion of memory by the attached processes. No method of coordinating access is automatically provided, so nothing prevents two processes from writing to the shared memory at the same time. For this reason, it is typically used with semaphores, which are used to synchronize processes.

Choice of IPC Mechanism

Applications can have specific functional requirements that determine which IPC mechanism to use. If one of several mechanisms can be used, the application writer determines which mechanism performs best for the application. The SunOS 5.x interprocess communication facilities are sensitive to application behavior. Determine which mechanism provides the best response capabilities by measuring the throughput capacity of each mechanism for the particular combination of message sizes used in the application.

Asynchronous Networking

This section discusses the techniques of asynchronous network communication using Transport-Level Interface (TLI) for realtime applications. SunOS provides support for asynchronous network processing of TLI events using a combination of STREAMS asynchronous features and the non-blocking mode of the TLI library routines.

For more information on the Transport-Level Interface, see the *Transport Interfaces Programming Guide* and the *man Pages(3): Library Routines*.

Modes of Networking

The Transport-Level Interface provides two modes of service: *connection-mode* and *connectionless-mode*.

Connection-Mode Service

The *connection-mode* is circuit-oriented and enables the transmission of data over an established connection in a reliable, sequenced manner. It also provides an identification procedure that avoids the overhead of address resolution and transmission during the data transfer phase. This service is attractive for applications that require relatively long-lived, datastream-oriented interactions.

Connectionless-Mode Service

Connectionless-mode is message-oriented and supports data transfer in self-contained units with no logical relationship required among multiple units. All information required to deliver a unit of data, including the destination address, is passed by the sender to the transport provider, together with the data, in a single service request. Connectionless-mode service is attractive for applications that involve short-term request/response interactions and do not require guaranteed, in-sequence delivery of data. It is generally assumed that connectionless transports are unreliable.

Networking Programming Models

Like file and device I/O, network transfers can be done synchronously or asynchronously with process service requests.

Synchronous Networking

Synchronous networking proceeds similarly to synchronous file and device I/O. Like the `write(2)` function, the request to send returns after buffering the message, but might suspend the calling process if buffer space is not immediately available. Like the `read(2)` function, a request to receive suspends execution of the calling process until data arrives to satisfy the request. Because SunOS 5.x provides no guaranteed bounds for transport services, synchronous networking is inappropriate for processes that must have realtime behavior with respect to other devices.

Asynchronous Networking

Asynchronous networking is provided by non-blocking service requests. Additionally, applications can request asynchronous notification when a connection might be established, when data might be sent, or when data might be received.

Asynchronous Connectionless-Mode Service

Asynchronous connectionless mode networking is conducted by configuring the endpoint for non-blocking service, and either polling for or receiving asynchronous notification when data might be transferred. If asynchronous notification is used, the actual receipt of data typically takes place within a signal handler.

Making the Endpoint Asynchronous

After the endpoint has been established using `t_open(3)`, and its identity established using `t_bind(3)`, the endpoint can be configured for asynchronous service. This is done by using the `fcntl(2)` function to set the `O_NONBLOCK` flag on the endpoint. Thereafter, calls to `t_sndudata(3)` for which no buffer space is immediately available return `-1` with `t_errno` set to `TFLOW`. Likewise, calls to `t_rcvudata(3)` for which no data are available return `-1` with `t_errno` set to `TNODATA`.

Asynchronous Network Transfers

Although an application can use the `poll(2)` function to wait for the receipt of data on an endpoint, it might be necessary to receive asynchronous notification when data has arrived. This can be done by using the `ioctl(2)` function with the `I_SETSIG` command to request that a `SIGPOLL` signal be sent to the process upon receipt of data at the endpoint. Applications should check for the possibility of multiple messages causing a single signal.

In the following example, `protocol` is the name of the application-chosen transport protocol.

Code Example 7-1 Asynchronous Network Transfers

```
#include <sys/types.h>
#include <tiuser.h>
#include <signal.h>
```

```
#include <stropts.h>

int          fd;
struct t_bind *bind;
void         sigpoll(int);

        fd = t_open(protocol, O_RDWR, (struct t_info *) NULL);

        bind = (struct t_bind *) t_alloc(fd, T_BIND, T_ADDR);
        ... /* set up binding address */
        t_bind(fd, bind, bin

/* make endpoint non-blocking */
fcntl(fd, F_SETFL, fcntl(fd, F_GETFL) | O_NONBLOCK);

/* establish signal handler for SIGPOLL */
signal(SIGPOLL, sigpoll);

/* request SIGPOLL signal when receive data is available */
ioctl(fd, I_SETSIG, S_INPUT | S_HIPRI);

...

void sigpoll(int sig)
{
    int          flags;
    struct t_unitdata ud;

    for (;;) {
        ... /* initialize ud */
        if (t_rcvudata(fd, &ud, &flags) < 0) {
            if (t_errno == TNOData)
                break; /* no more messages */
            ... /* process other error conditions */
        }
        ... /* process message in ud */
    }
}
```

Asynchronous Connection-Mode Service

For connection-mode service, an application can arrange for not only the data transfer, but for the establishment of the connection itself to be done asynchronously. The sequence of operations depends on whether the process is attempting to connect to another process or is awaiting connection attempts.

Asynchronously Establishing a Connection

A process can attempt a connection and asynchronously complete the connection. The process first creates the connecting endpoint, and, using `fcntl()`, configures the endpoint for non-blocking operation. As with connectionless data transfers, the endpoint can also be configured for asynchronous notification upon completion of the connection and subsequent data transfers. The connecting process then uses the `t_connect(3)` function to initiate setting up the transfer. Then the `t_rcvconnect(3)` function is used to confirm the establishment of the connection.

Asynchronous Use of a Connection

To asynchronously await connections, a process first establishes a non-blocking endpoint bound to a service address. When either the result of `poll()` or an asynchronous notification indicates that a connection request has arrived, the process can get the connection request by using the `t_listen(3)` function. To accept the connection, the process uses the `t_accept(3)` function. The responding endpoint must be separately configured for asynchronous data transfers.

The following example illustrates how to request a connection asynchronously

Code Example 7-2 Asynchronous Connection Request

```
#include <tiuser.h>

int          fd;
struct t_call *call;

    fd = .../* establish a non-blocking endpoint */

    call = (struct t_call *) t_alloc(fd, T_CALL, T_ADDR);
    .../* initialize call structure */
    t_connect(fd, call, call);

    /* connection request is now proceeding asynchronously */

    .../* receive indication that connection has been accepted */
    t_rcvconnect(fd, &call);
```

The following example illustrates listening for connections asynchronously

Code Example 7-3 Asynchronous Listening

```
#include <tiuser.h>

int          fd, res_fd;
struct t_call call;

fd = ... /* establish non-blocking endpoint */

.../*receive indication that connection request has arrived */
call = (struct t_call *) t_alloc(fd, T_CALL, T_ALL);
t_listen(fd, &call);

.../* determine whether or not to accept connection */
res_fd = ... /* establish non-blocking endpoint for response */
t_accept(fd, res_fd, call);
```

Asynchronous Open

Occasionally, an application might be required to dynamically open a regular file in a file system mounted from a remote host, or on a device whose initialization might be prolonged. However, while such an open is in progress, the application would be unable to achieve realtime response to other events. Fortunately, SunOS 5.x provides a means of solving this problem by having a second process perform the actual open and then pass the file descriptor to the realtime process.

Transferring a File Descriptor

The STREAMS interface under SunOS 5.x provides a mechanism for passing an open file descriptor from one process to another. The process with the open file descriptor uses the `ioctl(2)` function with a command argument of `I_SENDFD`. The second process obtains the file descriptor by calling the `ioctl()` function with a command argument of `I_RECVFD`.

In this example, the parent process prints out information about the test file, and creates a pipe. Next, the parent creates a child process, which opens the test file, and passes the open file descriptor back to the parent through the pipe. The parent process then displays the status information on the new file descriptor.

Code Example 7-4 File Descriptor Transfer

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stropts.h>
#include <stdio.h>

#define TESTFILE "/dev/null"
main(int argc, char * argv)
{
    int fd;
    int pipefd[2];
    struct stat statbuf;

    stat(TESTFILE, &statbuf);
    statout(TESTFILE, &statbuf);
    pipe(pipefd);
    if (fork() == 0) {
        close(pipefd[0]);
        sendfd(pipefd[1]);
    } else {
        close(pipefd[1])
        recvfd(pipefd[0]);
    }
}

sendfd(int p)
{
    int tfd;

    tfd = open(TESTFILE, O_RDWR);
    ioctl(p, I_SENDFD, tfd);
}

recvfd(int p)
{
    struct strrecvfd rfdbuf;
    struct stat statbuf;
    char fdbuf[32];

    ioctl(p, I_RECVFD, &rfdbuf);
    fstat(rfdbuf.fd, &statbuf);
    sprintf(fdbuf, "recvfd=%d", rfdbuf.fd);
    statout(fdbuf, &statbuf);
}
```

```
statout(char *f, struct stat *s)
{
    printf("stat: from=%s mode=%0o, ino=%d, dev=%d, rdev=%d\n",
        f, s->st_mode, s->st_ino, s->st_dev, s->st_rdev);
    fflush(stdout);
}
```

Timers

This section describes the timing facilities available for realtime applications under SunOS 5.x. Realtime applications that want to take advantage of these mechanisms will require detailed information from the manual pages of the routines listed in this section. These can be found in the *man Pages(3): Library Routines*.

The timing functions of SunOS 5.x fall into two separate areas of functionality: timestamps and interval timers. The timestamp functions provide a measure of elapsed time and allow the application to measure the duration of a state or the time between events. Interval timers allow an application to wake up at specified times and to schedule activities based on the passage of time. Although an application can poll a timestamp function to schedule itself, such an application would monopolize the processor to the detriment of other system functions.

Timestamp Functions

Two functions provide timestamps. The `gettimeofday(2)` function provides the current time in a *timeval* structure, representing the time in seconds and microseconds since midnight, Greenwich Mean Time, on January 1, 1970. The `clock_gettime(3R)` function, with a clockid of `CLOCK_REALTIME`, provides the current time in a *timespec* structure, representing in seconds and nanoseconds the same time interval returned by `gettimeofday()`.

SunOS 5.x uses a hardware periodic timer. For some workstations, this is the sole timing information, and the accuracy of timestamps is limited to the resolution of that periodic timer. For other platforms, a timer register with a resolution of one microsecond allows SunOS 5.x to provide timestamps accurate to one microsecond.

Interval Timer Functions

Realtime applications often schedule their activities through the use of interval timers. Interval timers can be either of two types: a “one-shot” type or a “periodic” type.

The one-shot is an armed timer that is set with an initial expiration time relative either to current time or to an absolute time. This timer expires once and is then disarmed. Such a timer might be useful for clearing buffers after the data has been transferred to storage, or to time-out an operation that should have finished.

The periodic timer is armed with the initial expiration time (either absolute or relative) and a repetition interval. Each time the interval timer expires it is reloaded with the repetition interval and the timer is automatically rearmed. This timer might be useful for data logging or for servo-control. In calls to interval timer functions, time values smaller than the resolution of the system hardware periodic timer are rounded up to the next multiple of the hardware periodic timer interval (typically 10 ms).

The IPC shared semaphore facility provides process synchronization. Shared memory is the fastest form of interprocess communication. The main advantage of shared memory is that the copying of message data is eliminated. The usual mechanism for synchronizing shared memory access is semaphores.

There are two sets of timers interfaces in SunOS 5.x. The `setitimer(2)` and `getitimer(2)` interfaces provide access to fixed set timers, called the BSD timers, using the `timeval` structure to specify time intervals. The POSIX timers are specifically related to POSIX clocks; the only POSIX clock currently supported is `CLOCK_REALTIME`. POSIX timer operations are expressed in terms of the `timespec` structure.

The functions `getitimer(2)` and `setitimer(2)` respectively retrieve and establish the value of the specified BSD interval timer. There are three BSD interval timers available to a process, including a realtime timer designated `ITIMER_REAL`. If a BSD timer is armed and allowed to expire, the system sends a signal appropriate to the timer to the process that set the timer.

The `timer_create(3R)` function can create up to `{TIMER_MAX}` POSIX timers. At the time of creation, the caller can specify what signal and what associated value will be sent to the process upon timer expiration. The `timer_gettime(3R)` and `timer_settime(3R)` functions respectively retrieve and establish the value of the specified POSIX interval timer.

Expirations of POSIX timers while the required signal is pending delivery are counted, and the function `timer_getoverrun(3R)` retrieves the count of such expirations. The function `timer_delete(3R)` deallocates a POSIX timer.

Code Example 7-5 illustrates how to use the `setitimer` interface to generate a periodic interrupt, and how to control the arrival of timer interrupts.

Code Example 7-5 Controlling Timer Interrupts

```
#include<unistd.h>
#include<signal.h>
#include<sys/time.h>

#define TIMERCNT 8

void timerhandler();
int timercnt;
struct timeval alarmtimes[TIMERCNT];

main()
{
    struct itimerval times;
    sigset_t sigset;
    int i, ret;
    struct sigaction act;

    /* block SIGALRM */
    sigemptyset(&sigset);
    sigaddset(&sigset, SIGALRM);
    sigprocmask(SIG_BLOCK, &sigset, NULL);

    /* set up handler for SIGALRM */
    act.sa_handler = timerhandler;
    sigemptyset(&act.sa_mask);
    act.sa_flags = SA_SIGINFO;
    sigaction(SIGALRM, &act, NULL);
    /*
     * set up interval timer, starting in three seconds,
     * then every 1/3 second
     */
    times.it_value.tv_sec = 3;
    times.it_value.tv_usec = 0;
    times.it_interval.tv_sec = 0;
    times.it_interval.tv_usec = 333333;
    ret = setitimer(ITIMER_REAL, &times, NULL);
    printf("main:setitimer ret = %d\n", ret);
}
```

```

        /* now wait for the alarms */
        sigemptyset(&sigset);
        timerhandler(0, 0, NULL, NULL);
        while (timercnt < TIMERCNT) {
            ret = sigsuspend(&sigset);
        }
        printtimes();
    }

void timerhandler(sig, siginfo, context)
    int    sig;
    siginfo_t siginfo;
    void    *context;
{
    printf("timerhandler:start\n");
    gettimeofday(&alarmtimes[timercnt], NULL);
    timercnt++;
    printf("timerhandler:timercnt = %d\n", timercnt);
}

printtimes()
{
    int i;

    for (i = 0; i < TIMERCNT; i++) {
        printf("%d.%06d\n", alarmtimes[i].tv_sec,
            alarmtimes[i].tv_usec);
    }
}

```

Index

Symbols

/dev/zero, mapping, 98

A

address space of a process, 92, 103
advisory locking, 56
asynchronous I/O
 behavior, 111
 endpoint service, 140
 guaranteeing buffer state, 111
 listen for network connection, 142
 making connection request, 142
 notification of data arrival, 140
 opening a file, 143
 using aio_result_t structure, 111
 waiting for completion, 129
atomic updates to semaphores, 76

B

blocking mode
 defined, 120
 finite time quantum, 116
 opening a FIFO, 135
 priority inversion, 120
 time-sharing process, 109
 using the read() function, 135

bounded, 112
brk(2), 106

C

chmod(1), 62
class
 definition, 115
 priority queue, 118
 scheduling algorithm, 117
 scheduling priorities, 115
connectionless-mode
 asynchronous network service, 140
 definition, 139
connection-mode
 asynchronous network service, 141
 asynchronously connecting, 142
 definition, 139
 using asynchronous connection, 142
context switch
 preempting a process, 119
control
 semaphore set, 77
creat(), 56
creation flags, IPC, 67 to ??

D

dispatch
 priorities, 116
dispatch latency, 112
 under realtime, 112
dispatch table
 configuring, 124
 kernel, 119

F

F_GETLK, 59
fcntl(2), 56, 57
FIFO
 using as byte stream, 135
file and record locking, 55 to ??
file descriptor
 passing to another process, 143
 transferring, 143
file system
 contiguous, 112
 opening dynamically, 143
 using pipes, 135
files
 lock, 55 to ??
 memory-mapped, *See* mapped files
fork(2), ?? to 10
fsync(2), 93
functions
 advanced I/O, 53
 basic I/O, 52
 IPC, 65 to ??
 list file system control, 54
 signals, 17 to ??
 terminal I/O, 63
 user processes, 10

G

GETALL, 81
GETNCNT, 81
GETPID, 80
GETVAL, 80

GETZCNT, 81

I

I/O, *See* asynchronous I/O, or
 synchronous I/O
init(1M), scheduler properties, 47
Interprocess Communication (IPC)
 administering, 138
 creating pipes, 134
 memory mapped files, 137
 using fileless memory mapping, 137
 using memory mapping, 137
 using messages, 135
 using named pipes, 135
 using pipes, 133
 using semaphores, 136
 using shared memory, 136
 using the open() call, 135
IPC (interprocess communication), 65 to
 ??
 creation flags, 67 to ??
 functions, 67 to ??
 message header, 71
 message queue, 69
 messages, 68 to ??
 permissions, 66 to 67
 semaphore set, 77
 semaphores, 75 to 83
 shared memory, 83 to ??
IPC_NOWAIT, 83
IPC_RMID, 72, 81, 87
IPC_SET, 72, 81, 87
IPC_STAT, 72, 81, 87

K

kernel
 class independent, 117
 context switch, 119
 dispatch table, 119
 preempting current process, 119
 queue, 111

L

- lockf(3C), 57, 60, 178
- locking
 - advisory, 56, 62
 - F_GETLK, 59
 - finding locks, 59
 - mandatory, 56, ?? to 62
 - memory in realtime, 127
 - opening a file for, 58
 - read, 56, 58, 178
 - record, 58
 - removing, 58 to ??
 - setting, 58 to ??
 - supported file systems, 55
 - testing locks, 59
 - with fcntl(2), 56 to 57
 - with lockf(3C), 57
 - write, 56, 58, 178
- ls(1), 62
- lseek(2), 57

M

- mandatory locking, 56
- mapped files, 94 to 100
 - private, 95
 - shared, 95
- memory
 - locking, 127
 - locking a page, 127
 - locking all pages, 128
 - number of locked pages, 127
 - sticky locks, 128
 - unlocking a page, 128
- memory management, 91 to 106
 - address spaces, 92
 - address-space layout, 103
 - coherence, 93
 - concepts, 91
 - functions, 94
 - heterogeneity, 93
 - mapping, 91
 - mlock(3C), 101 to ??
 - mlockall(3C), 102 to ??

- mmap(2), 94 to 100
- mprotect(2), 103
- msync(3C), 102 to 103
- munmap(2), 100
- networking, 93
- pagesize, 103
- virtual memory, 91

- memory-mapped files, *See* mapped files
- message queue, 69
- message, header, 71
- messages, 65, 68 to ??
- mlock(3C), 101 to ??
- mlockall(3C), 102 to ??
- mmap(2), 94 to 100
- mprotect(2), 103
- msgget(), 69
- msqid, 71
- msync(3C), 102 to 103
- munmap(2), 100

N

- named pipe
 - defined, 135
 - FIFO, 133
 - using, 135
- named pipes, limitations, 65
- network
 - asynchronous connection, 138
 - asynchronous service, 140
 - asynchronous transfers, 140
 - asynchronous use, 140
 - connectionless-mode service, 139
 - connection-mode service, 139
 - programming models for
 - realtime, 139
 - services under realtime, 138
 - synchronous use, 139
 - using STREAMS
 - asynchronously, 138
 - using Transport-Level Interface (TLI), 138
- nice(1), 47
- nice(2), 47

- non-blocking mode
 - configuring endpoint connections, 142
 - defined, 138
 - endpoint bound to service address, 142
 - network service, 140
 - polling for notification, 140
 - service requests, 140
 - Transport-Level Interface (TLI), 138
 - using the `t_connect()` function, 142

O

- `open()`, 56

P

- page 0, 106
- performance, scheduler effect on, 47
- permissions
 - IPC, 66 to 67
- pipe
 - defined, 135
 - non-blocking read, 135
- pipes, limitations, 65
- polling
 - for a connection request, 142
 - notification of data, 140
 - using the `poll(2)` function, 140
- `prctl(1)`, 30 to ??
- `prctl(2)`, ?? to 44
- `prctlset(2)`, ?? to 46
- priority inversion
 - defined, 109
 - synchronization, 120
- priority queue
 - linear linked list, 119
- process
 - defined for realtime, 107
 - dispatching, 119
 - highest priority, 108
 - preemption, 119
 - residence in memory, 127

- runaway, 111
- scheduling for realtime, 116
- setting priorities, 123

- process address space, 92, 103

- process priority
 - global, 25
 - real-time, 28
 - setting and retrieving, 30 to 46
 - system, 28
 - time-sharing, 28

- process, spawning, ?? to 10

- processes, cooperating, locking, 56

R

- read

- blocking, 135

- read lock, 56, 58, 178

- `read()`, 56

- real-time, scheduler class, 27

- removing record locks, 58 to ??

- response time

- blocking processes, 110

- bounds to I/O, 109

- degrading, 108

- inheriting priority, 109

- servicing interrupts, 109

- sharing libraries, 109

- sticky locks, 110

- reversing operations for semaphores, 77

S

- `sbrk(2)`, 106

- scheduler, 13 to 14, 23 to 50

- classes, 117

- configuring, 124

- effect on performance, 47

- priority, 115

- realtime, 112

- real-time policy, 27

- scheduling classes, 115

- system policy, 27

- time-sharing policy, 26

- using system calls, 121
 - using utilities, 122
- scheduler, class, 27
- SEM_UNDO, 83
- semaphores, 75 to 83
 - advantages, 65
 - arbitrary simultaneous updates, 76
 - atomic updates, 76
 - reversing operations and SEM_UNDO, 77
 - set structure, 77
 - undo structure, 76
- semget(), 75
- semop(), 76
- SETALL, 81
- setting record locks, 58 to ??
- SETVAL, 80
- shared memory, 65, 83 to ??
- SHM_LOCK, 87
- SHM_UNLOCK, 87
- shmget(), 84
- signals
 - code blocking, 20
 - handlers, 17, 20
 - limitations, 65
 - process control, 17
 - resource limits, 17
 - sending, 18 to 19
 - stacks, 21 to ??
- structure, semaphore set, 77
- synchronization, 93, 113
 - shared memory, 138
- synchronous I/O
 - blocking, 129
 - critical timing, 109

T

- time slice, real-time process, 40
- timers
 - for interval timing, 145
 - for realtime applications, 145
 - timestamping, 145

- using one-shot, 146
 - using periodic type, 146
- time-sharing
 - scheduler class, 26
 - scheduler parameter table, 27
- Transport-Level Interface (TLI)
 - asynchronous endpoint, 140
 - connectionless-mode, 138
 - connection-mode, 138

U

- undo structure for semaphores, 76
- updates, atomic for semaphores, 76
- user priority, 29

V

- virtual memory, 91 to 106
 - See also* memory management

W

- write lock, 56, 58, 178
- write(), 56

Z

- zero(7), 98

Copyright 1995 Sun Microsystems Inc., 2550 Garcia Avenue, Mountain View, Californie 94043-1100 U.S.A.

Tous droits réservés. Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, et la décompilation. Aucune partie de ce produit ou de sa documentation associée ne peuvent Être reproduits sous aucune forme, par quelque moyen que ce soit sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il en a.

Des parties de ce produit pourront être dérivées du système UNIX[®], licencié par UNIX System Laboratories, Inc., filiale entièrement détenue par Novell, Inc., ainsi que par le système 4.3. de Berkeley, licencié par l'Université de Californie. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

LEGENDE RELATIVE AUX DROITS RESTREINTS: l'utilisation, la duplication ou la divulgation par l'administration américaine sont soumises aux restrictions visées à l'alinéa (c)(1)(ii) de la clause relative aux droits des données techniques et aux logiciels informatiques du DFARS 252.227-7013 et FAR 52.227-19. Le produit décrit dans ce manuel peut Être protégé par un ou plusieurs brevet(s) américain(s), étranger(s) ou par des demandes en cours d'enregistrement.

MARQUES

Sun, Sun Microsystems, le logo Sun, SunSoft, le logo SunSoft, Solaris, SunOS, OpenWindows, DeskSet, ONC, ONC+ et NFS sont des marques déposées ou enregistrées par Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays, et exclusivement licenciée par X/Open Company Ltd. OPEN LOOK est une marque enregistrée de Novell, Inc. PostScript et Display PostScript sont des marques d'Adobe Systems, Inc.

Toutes les marques SPARC sont des marques déposées ou enregistrées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. SPARCcenter, SPARCcluster, SPARCcompiler, SPARCdesign, SPARC811, SPARCengine, SPARCprinter, SPARCserver, SPARCstation, SPARCstorage, SPARCworks, microSPARC, microSPARC-II, et UltraSPARC sont exclusivement licenciées à Sun Microsystems, Inc. Les produits portant les marques sont basés sur une architecture développée par Sun Microsystems, Inc.

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

Le système X Window est un produit du X Consortium, Inc.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" SANS GARANTIE D'AUCUNE SORTE, NI EXPRESSE NI IMPLICITE, Y COMPRIS, ET SANS QUE CETTE LISTE NE SOIT LIMITATIVE, DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DES PRODUITS A REPENDRE A UNE UTILISATION PARTICULIERE OU LE FAIT QU'ILS NE SOIENT PAS CONTREFAISANTS DE PRODUITS DE TIERS.

CETTE PUBLICATION PEUT CONTENIR DES MENTIONS TECHNIQUES ERRONEES OU DES ERREURS TYPOGRAPHIQUES. DES CHANGEMENTS SONT PERIODIQUEMENT APPORTES AUX INFORMATIONS CONTENUES AUX PRESENTES. CES CHANGEMENTS SERONT INCORPORES AUX NOUVELLES EDITIONS DE LA PUBLICATION. SUN MICROSYSTEMS INC. PEUT REALISER DES AMELIORATIONS ET/OU DES CHANGEMENTS DANS LE(S) PRODUIT(S) ET/OU LE(S) PROGRAMME(S) DECRITS DANS CETTE PUBLICATION A TOUS MOMENTS.

