



Multithreaded Programming Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A.

Part Number 806-5257-10
January 2001

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

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

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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

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

Copyright 2001 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, Californie 94303-4900 Etats-Unis. Tous droits réservés.

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

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

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



Contents

Preface 13

1. Covering Multithreading Basics 17

Defining Multithreading Terms 17

Meeting Multithreading Standards 19

Benefiting From Multithreading 19

Improving Application Responsiveness 19

Using Multiprocessors Efficiently 19

Improving Program Structure 20

Using Fewer System Resources 20

Combining Threads and RPC 20

Understanding Basic Multithreading Concepts 20

Concurrency and Parallelism 20

Looking at Multithreading Structure 21

Scheduling 23

Cancellation 24

Synchronization 24

Using the 64-bit Architecture 25

2. Basic Threads Programming 27

The Threads Library 27

Create a Default Thread	28
Wait for Thread Termination	29
A Simple Threads Example	31
Detaching a Thread	32
Create a Key for Thread-Specific Data	33
Delete the Thread-Specific Data Key	34
Set the Thread-Specific Data Key	35
Get the Thread-Specific Data Key	35
Get the Thread Identifier	38
Compare Thread IDs	39
Initializing Threads	39
Yield Thread Execution	40
Set the Thread Priority	40
Get the Thread Priority	41
Send a Signal to a Thread	42
Access the Signal Mask of the Calling Thread	43
Forking Safely	44
Terminate a Thread	44
Finishing Up	45
Cancellation	45
Cancel a Thread	47
Enable or Disable Cancellation	48
Set Cancellation Type	48
Create a Cancellation Point	49
Push a Handler Onto the Stack	50
Pull a Handler Off the Stack	50

3. Thread Create Attributes 53

Attributes	54
------------	----

Initialize Attributes	55
Destroy Attributes	56
Set Detach State	57
Get Detach State	58
Set Stack Guard Size	59
Get Stack Guard Size	60
Set Scope	60
Get Scope	61
Set Thread Concurrency Level	62
Get Thread Concurrency Level	63
Set Scheduling Policy	63
Get Scheduling Policy	65
Set Inherited Scheduling Policy	65
Get Inherited Scheduling Policy	66
Set Scheduling Parameters	67
Get Scheduling Parameters	68
Set Stack Size	69
Get Stack Size	70
About Stacks	71
Set Stack Address	72
Get Stack Address	74
4. Programming with Synchronization Objects	75
Mutual Exclusion Lock Attributes	76
Initialize a Mutex Attribute Object	78
Destroy a Mutex Attribute Object	79
Set the Scope of a Mutex	80
Get the Scope of a Mutex	81
Set the Mutex Type Attribute	81

Get the Mutex Type Attribute	83
Set Mutex Attribute's Protocol	83
Get Mutex Attribute's Protocol	86
Set Mutex Attribute's Priority Ceiling	87
Get Mutex Attribute's Priority Ceiling	88
Set Mutex's Priority Ceiling	89
Get Mutex's Priority Ceiling	90
Set Mutex's Robust Attribute	91
Get Mutex's Robust Attribute	92
Using Mutual Exclusion Locks	93
Initialize a Mutex	94
Make Mutex Consistent	95
Lock a Mutex	97
Unlock a Mutex	98
Lock With a Nonblocking Mutex	99
Destroy a Mutex	101
Mutex Lock Code Examples	101
Condition Variable Attributes	106
Initialize a Condition Variable Attribute	107
Remove a Condition Variable Attribute	108
Set the Scope of a Condition Variable	109
Get the Scope of a Condition Variable	110
Using Condition Variables	110
Initialize a Condition Variable	111
Block on a Condition Variable	112
Unblock a Specific Thread	114
Block Until a Specified Event	115
Unblock All Threads	116

Destroy Condition Variable State	118
The Lost Wake-Up Problem	118
The Producer/Consumer Problem	119
Semaphores	122
Counting Semaphores	123
Initialize a Semaphore	124
Named Semaphores	125
Increment a Semaphore	126
Block on a Semaphore Count	126
Decrement a Semaphore Count	127
Destroy the Semaphore State	128
The Producer/Consumer Problem, Using Semaphores	128
Read-Write Lock Attributes	130
Initialize a Read-Write Lock Attribute	131
Destroy a Read-Write Lock Attribute	131
Set a Read-Write Lock Attribute	132
Get a Read-Write Lock Attribute	132
Using Read-Write Locks	133
Initialize a Read-Write Lock	134
Read Lock on Read-Write Lock	135
Read Lock With a Nonblocking Read-Write Lock	136
Write Lock on Read-Write Lock	136
Write Lock With a Nonblocking Read-Write Lock	137
Unlock a Read-Write Lock	138
Destroy a Read-Write Lock	139
Synchronization Across Process Boundaries	139
Producer/Consumer Problem Example	140
Interprocess Locking Without the Threads Library	141

Comparing Primitives 141

5. Programming With the Operating Environment 143

Process Creation—Forking Issues 143

The Fork-One Model 144

The Fork-All Model 147

Choosing the Right Fork 147

Process Creation—exec(2) and exit(2) Issues 148

Timers, Alarms, and Profiling 148

Per-LWP POSIX Timers 148

Per-Thread Alarms 149

Profiling 149

Nonlocal Goto—setjmp(3C) and longjmp(3C) 150

Resource Limits 150

LWPs and Scheduling Classes 150

Timeshare Scheduling 151

Realtime Scheduling 152

LWP Scheduling and Thread Binding 152

SIGWAITING—Creating LWPs for Waiting Threads 153

Aging LWPs 154

Extending Traditional Signals 154

Synchronous Signals 155

Asynchronous Signals 155

Continuation Semantics 156

Operations on Signals 156

Thread-Directed Signals 158

Completion Semantics 159

Signal Handlers and Async-Signal Safety 160

Interrupted Waits on Condition Variables (Solaris Threads Only) 162

I/O Issues	164
I/O as a Remote Procedure Call	164
Tamed Asynchrony	164
Asynchronous I/O	165
Shared I/O and New I/O System Calls	166
Alternatives to getc(3S) and putc(3S)	167
6. Safe and Unsafe Interfaces	169
Thread Safety	169
MT Interface Safety Levels	170
Reentrant Functions for Unsafe Interfaces	172
Async-Signal-Safe Functions	173
MT Safety Levels for Libraries	173
Unsafe Libraries	174
7. Compiling and Debugging	175
Compiling a Multithreaded Application	175
Preparing for Compilation	175
Choosing Solaris or POSIX Semantics	176
Including <thread.h> or <pthread.h>	176
Defining _REENTRANT or _POSIX_C_SOURCE	177
Linking With libthread or libpthread	177
Linking With -lposix4 for POSIX Semaphores	179
Link Old With New	179
Linking With the Alternate One-level libthread Library	180
Debugging a Multithreaded Program	181
Common Oversights	181
Tracing and Debugging With the TNF Utilities	182
Using truss(1)	182
Using adb(1)	182

Using dbx 183

8. Programming With Solaris Threads 187

Comparing APIs for Solaris Threads and POSIX Threads 187

Major API Differences 188

Function Comparison Table 188

Unique Solaris Threads Functions 194

Suspend Thread Execution 194

Continue a Suspended Thread 195

Set Thread Concurrency Level 196

Get Thread Concurrency 197

Similar Synchronization Functions—Read-Write Locks 198

Initialize a Read-Write Lock 198

Acquire a Read Lock 200

Try to Acquire a Read Lock 200

Acquire a Write Lock 201

Try to Acquire a Write Lock 202

Unlock a Read-Write Lock 202

Destroy Read-Write Lock State 203

Similar Solaris Threads Functions 204

Create a Thread 205

Get the Minimal Stack Size 208

Get the Thread Identifier 209

Yield Thread Execution 209

Send a Signal to a Thread 209

Access the Signal Mask of the Calling Thread 209

Terminate a Thread 210

Wait for Thread Termination 210

Create a Thread-Specific Data Key 211

Set the Thread-Specific Data Key	211
Get the Thread-Specific Data Key	212
Set the Thread Priority	212
Get the Thread Priority	213
Similar Synchronization Functions—Mutual Exclusion Locks	213
Initialize a Mutex	213
Destroy a Mutex	215
Acquire a Mutex	215
Release a Mutex	215
Try to Acquire a Mutex	216
Similar Synchronization Functions—Condition Variables	216
Initialize a Condition Variable	216
Destroy a Condition Variable	217
Wait for a Condition	217
Wait for an Absolute Time	218
Signal One Condition Variable	218
Signal All Condition Variables	219
Similar Synchronization Functions—Semaphores	219
Initialize a Semaphore	219
Increment a Semaphore	220
Block on a Semaphore Count	220
Decrement a Semaphore Count	221
Destroy the Semaphore State	221
Synchronization Across Process Boundaries	221
Using LWPs Between Processes	222
Producer/Consumer Problem Example	222
Special Issues for fork() and Solaris Threads	224
9. Programming Guidelines	225

Rethinking Global Variables	225
Providing for Static Local Variables	226
Synchronizing Threads	227
Single-Threaded Strategy	228
Reentrance	228
Avoiding Deadlock	230
Deadlocks Related to Scheduling	231
Locking Guidelines	231
Following Some Basic Guidelines	232
Creating and Using Threads	232
Lightweight Processes	233
Unbound Threads	234
Bound Threads	234
Thread Concurrency (Solaris Threads Only)	235
Efficiency	236
Thread Creation Guidelines	236
Working With Multiprocessors	236
The Underlying Architecture	237
Summary	241
Further Reading	241
A. Sample Application—Multithreaded grep	243
Description of tgrep	243
Getting Online Source Code	244
B. Solaris Threads Example: barrier.c	267
C. MT Safety Levels: Library Interfaces	271
Index	379

Preface

The *Multithreaded Programming Guide* describes the multithreaded programming interfaces for POSIX and Solaris threads in the Solaris™ Operating Environment. This guide shows application programmers how to create new multithreaded programs and how to add multithreading to existing programs.

Although this guide covers both the POSIX and Solaris threads implementations, most topics assume a POSIX threads interest. Information applying to only Solaris threads is covered in a special chapter.

To understand this guide, a reader must be familiar with

- A UNIX® SVR4 system—preferably the Solaris Operating Environment.
- The C programming language—multithreading is implemented through the libthread library
- The principles of concurrent programming (as opposed to sequential programming)—multithreading requires a different way of thinking about function interactions. Some books you might want to read are:
 - *Algorithms for Mutual Exclusion* by Michel Raynal (MIT Press, 1986)
 - *Concurrent Programming* by Alan Burns & Geoff Davies (Addison-Wesley, 1993)
 - *Distributed Algorithms and Protocols* by Michel Raynal (Wiley, 1988)
 - *Operating System Concepts* by Silberschatz, Peterson, & Galvin (Addison-Wesley, 1991)
 - *Principles of Concurrent Programming* by M. Ben-Ari (Prentice-Hall, 1982)

Ordering Sun Documents

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

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at <http://www1.fatbrain.com/documentation/sun>.

How This Guide Is Organized

Chapter 1 gives a structural overview of threads implementation in this release.

Chapter 2 discusses the general POSIX threads library routines, emphasizing creating a thread with default attributes.

Chapter 3 covers creating a thread with nondefault attributes.

Chapter 4 covers the threads library synchronization routines.

Chapter 5 discusses changes to the operating environment to support multithreading.

Chapter 6 covers multithreading safety issues.

Chapter 7 covers the basics of compiling and debugging multithreaded applications.

Chapter 8 covers the Solaris threads (as opposed to POSIX threads) interfaces.

Chapter 9 discusses issues that affect programmers writing multithreaded applications.

Appendix A shows how code can be designed for POSIX threads.

Appendix B shows an example of building a barrier in Solaris threads.

Appendix C lists the safety levels of library routines.

Accessing Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is <http://docs.sun.com>.

What Typographic Conventions Mean

The following table describes the typographic changes used in this book.

TABLE P-1 Typographic Conventions

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

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell prompt	<code>machine_name%</code>
C shell superuser prompt	<code>machine_name#</code>

TABLE P-2 Shell Prompts *(continued)*

Shell	Prompt
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Covering Multithreading Basics

The word *multithreading* can be translated as *multiple threads of control* or *multiple flows of control*. While a traditional UNIX process always has contained and still does contain a single thread of control, multithreading (MT) separates a process into many execution threads, each of which runs independently.

Multithreading your code can

- Improve application responsiveness
- Use multiprocessors more efficiently
- Improve program structure
- Use fewer system resources

This chapter explains some multithreading terms, benefits, and concepts. If you are ready to start using multithreading, skip to Chapter 2.

- “Defining Multithreading Terms” on page 17
- “Meeting Multithreading Standards” on page 19
- “Benefiting From Multithreading” on page 19
- “Understanding Basic Multithreading Concepts” on page 20

Defining Multithreading Terms

Table 1–1 introduces some of the terms used in this book.

TABLE 1-1 Multithreading Terms

Term	Definition
Process	The UNIX environment (such as file descriptors, user ID, and so on) created with the <code>fork(2)</code> system call, which is set up to run a program.
Thread	A sequence of instructions executed within the context of a process.
pthread (POSIX threads)	A POSIX 1003.1c compliant threads interface.
Solaris threads	A Sun Microsystems™ threads interface that is not POSIX compliant. A predecessor of pthreads.
Single-threaded	Restricting access to a single thread.
Multithreaded	Allowing access to two or more threads.
User- or Application-level threads	Threads managed by the threads library routines in user (as opposed to kernel) space.
Lightweight processes	Threads in the kernel that execute kernel code and system calls (also called LWPs).
Bound threads	Threads that are permanently bound to LWPs.
Unbound threads	A default Solaris thread that context switches very quickly without kernel support.
Attribute object	Contains opaque data types and related manipulation functions used to standardize some of the configurable aspects of POSIX threads, mutual exclusion locks (mutexes), and condition variables.
Mutual exclusion locks	Functions that lock and unlock access to shared data.
Condition variables	Functions that block threads until a change of state.
Counting semaphore	A memory-based synchronization mechanism.
Parallelism	A condition that arises when at least two threads are <i>executing</i> simultaneously.
Concurrency	A condition that exists when at least two threads are <i>making progress</i> . A more generalized form of parallelism that can include time-slicing as a form of virtual parallelism.

Meeting Multithreading Standards

The concept of multithreaded programming goes back to at least the 1960s. Its development on UNIX systems began in the mid-1980s. While there is agreement about what multithreading is and the features necessary to support it, the interfaces used to implement multithreading have varied greatly.

For several years a group called POSIX (Portable Operating System Interface) 1003.4a has been working on standards for multithreaded programming. The standard has now been ratified. This Multithreaded Programming Guide is based on the POSIX standards: P1003.1b final draft 14 (realtime), and P1003.1c final draft 10 (multithreading).

This guide covers both POSIX threads (also called *pthread*s) and Solaris threads. Solaris threads were available in the Solaris 2.4 release, and are not functionally different from POSIX threads. However, because POSIX threads are more portable than Solaris threads, this guide covers multithreading from the POSIX perspective. Subjects specific to Solaris threads only are covered in the Chapter 8.

Benefiting From Multithreading

Improving Application Responsiveness

Any program in which many activities are not dependent upon each other can be redesigned so that each activity is defined as a thread. For example, the user of a multithreaded GUI does not have to wait for one activity to complete before starting another.

Using Multiprocessors Efficiently

Typically, applications that express concurrency requirements with threads need not take into account the number of available processors. The performance of the application improves transparently with additional processors.

Numerical algorithms and applications with a high degree of parallelism, such as matrix multiplications, can run much faster when implemented with threads on a multiprocessor.

Improving Program Structure

Many programs are more efficiently structured as multiple independent or semi-independent units of execution instead of as a single, monolithic thread. Multithreaded programs can be more adaptive to variations in user demands than single-threaded programs.

Using Fewer System Resources

Programs that use two or more processes that access common data through shared memory are applying more than one thread of control.

However, each process has a full address space and operating environment state. The cost of creating and maintaining this large amount of state information makes each process much more expensive than a thread in both time and space.

In addition, the inherent separation between processes can require a major effort by the programmer to communicate between the threads in different processes, or to synchronize their actions.

Combining Threads and RPC

By combining threads and a remote procedure call (RPC) package, you can exploit nonshared-memory multiprocessors (such as a collection of workstations). This combination distributes your application relatively easily and treats the collection of workstations as a multiprocessor.

For example, one thread might create child threads. Each of these children could then place a remote procedure call, invoking a procedure on another workstation. Although the original thread has merely created threads that are now running in parallel, this parallelism involves other computers.

Understanding Basic Multithreading Concepts

Concurrency and Parallelism

In a multithreaded process on a single processor, the processor can switch execution resources between threads, resulting in concurrent execution.

In the same multithreaded process in a shared-memory multiprocessor environment, each thread in the process can run on a separate processor at the same time, resulting in parallel execution.

When the process has fewer or as many threads as there are processors, the threads support system in conjunction with the operating environment ensure that each thread runs on a different processor.

For example, in a matrix multiplication that has the same number of threads and processors, each thread (and each processor) computes a row of the result.

Looking at Multithreading Structure

Traditional UNIX already supports the concept of threads—each process contains a single thread, so programming with multiple processes is programming with multiple threads. But a process is also an address space, and creating a process involves creating a new address space.

Creating a thread is much less expensive when compared to creating a new process, because the newly created thread uses the current process address space. The time it takes to switch between threads is much less than the time it takes to switch between processes, partly because switching between threads does not involve switching between address spaces.

Communicating between the threads of one process is simple because the threads share everything—address space, in particular. So, data produced by one thread is immediately available to all the other threads.

The interface to multithreading support is through a subroutine library, `libpthread` for POSIX threads, and `libthread` for Solaris threads. Multithreading provides flexibility by decoupling kernel-level and user-level resources.

User-Level Threads

Threads are the primary programming interface in multithreaded programming. User-level threads¹ are handled in user space and avoid kernel context switching penalties. An application can have hundreds of threads and still not consume many kernel resources. How many kernel resources the application uses is largely determined by the application.

Threads are visible only from within the process, where they share all process resources like address space, open files, and so on. The following state is unique to each thread.

■ Thread ID

1. User-level threads are named to distinguish them from kernel-level threads, which are the concern of systems programmers, only. Because this book is for application programmers, kernel-level threads are not discussed.

- Register state (including PC and stack pointer)
- Stack
- Signal mask
- Priority
- Thread-private storage

Because threads share the process instructions and most of the process data, a change in shared data by one thread can be seen by the other threads in the process. When a thread needs to interact with other threads in the same process, it can do so without involving the operating environment.

By default, threads are very lightweight. But, to get more control over a thread (for instance, to control scheduling policy more), the application can bind the thread. When an application binds threads to execution resources, the threads become kernel resources (see “System Scope (Bound Threads)” on page 24 for more information).

To summarize, user-level threads are:

- Inexpensive to create because they do not need to create their own address space. They are bits of virtual memory that are allocated from your address space at run time.
- Fast to synchronize because synchronization is done at the application level, not at the kernel level.
- Easily managed by the threads library; either `libpthread` or `libthread`.

Lightweight Processes

The threads library uses underlying threads of control called *lightweight processes* that are supported by the kernel. You can think of an LWP as a virtual CPU that executes code or system calls.

You usually do not need to concern yourself with LWPs to program with threads. The information here about LWPs is provided as background, so you can understand the differences in scheduling scope, described on “Process Scope (Unbound Threads)” on page 23.

Note - The LWPs in the Solaris 2, Solaris 7, and Solaris 8 operating environments are *not* the same as the LWPs in the SunOS™ 4.0 LWP library, which are not supported in the Solaris 2, Solaris 7, and Solaris 8 operating environments.

Much as the `stdio` library routines such as `fopen( )` and `fread( )` use the `open( )` and `read( )` functions, the threads interface uses the LWP interface, and for many of the same reasons.

Lightweight processes (LWPs) bridge the user level and the kernel level. Each process contains one or more LWP, each of which runs one or more user threads.

(See Figure 1–1.) The creation of a thread usually involves just the creation of some user context, but not the creation of an LWP.




Figure 1–1 User-level Threads and Lightweight Processes

Each LWP is a kernel resource in a kernel pool, and is allocated (attached) and de-allocated (detached) to a thread on a per thread basis. This happens as threads are scheduled or created and destroyed.

Scheduling

POSIX specifies three scheduling policies: first-in-first-out (`SCHED_FIFO`), round-robin (`SCHED_RR`), and custom (`SCHED_OTHER`). `SCHED_FIFO` is a queue-based scheduler with different queues for each priority level. `SCHED_RR` is like FIFO except that each thread has an execution time quota.

Both `SCHED_FIFO` and `SCHED_RR` are POSIX Realtime extensions. `SCHED_OTHER` is the default scheduling policy.

See “LWPs and Scheduling Classes” on page 150 for information about the `SCHED_OTHER` policy, and about emulating some properties of the POSIX `SCHED_FIFO` and `SCHED_RR` policies.

Two scheduling scopes are available: process scope for unbound threads and system scope for bound threads. Threads with differing scope states can coexist on the same system and even in the same process. In general, the scope sets the range in which the threads scheduling policy is in effect.

Process Scope (Unbound Threads)

Unbound threads are created `PTHREAD_SCOPE_PROCESS`. These threads are scheduled in user space to attach and detach from available LWPs in the LWP pool.

LWPs are available to threads in this process only; that is threads are scheduled on these LWPs.

In most cases, threads should be `PTHREAD_SCOPE_PROCESS`. This allows the threads to float among the LWPs, and this improves threads performance (and is equivalent to creating a Solaris thread in the `THR_UNBOUND` state). The threads library decides, with regard to other threads, which threads get serviced by the kernel.

System Scope (Bound Threads)

Bound threads are created `PTHREAD_SCOPE_SYSTEM`. A boundthread is permanently attached to an LWP.

Each bound thread is bound to an LWP for the lifetime of the thread. This is equivalent to creating a Solaris thread in the `THR_BOUND` state. You can bind a thread to give it an alternate signal stack or to use special scheduling attributes with Realtime scheduling. All scheduling is done by the operating environment.

Note - In neither case, bound or unbound, can a thread be directly accessed by or moved to another process.

Cancellation

Thread cancellation allows a thread to terminate the execution of any other thread in the process. The target thread (the one being cancelled) can keep cancellation requests pending and can perform application-specific cleanup when it acts upon the cancellation notice.

The pthreads cancellation feature permits either asynchronous or deferred termination of a thread. Asynchronous cancellation can occur at any time; deferred cancellation can occur only at defined points. Deferred cancellation is the default type.

Synchronization

Synchronization allows you to control program flow and access to shared data for concurrently executing threads.

The four synchronization models are mutex locks, read / write locks, condition variables, and semaphores.

- Mutex locks allow only one thread at a time to execute a specific section of code, or to access specific data.

- Read/write locks permit concurrent reads and exclusive writes to a protected shared resource. To modify a resource, a thread must first acquire the exclusive write lock. An exclusive write lock is not permitted until all read locks have been released.
- Condition variables block threads until a particular condition is true.
- Counting semaphores typically coordinate access to resources. The count is the limit on how many threads can have access to a semaphore. When the count is reached, the semaphore blocks.

Using the 64-bit Architecture

For application developers, the major difference between the Solaris 64-bit and 32-bit operating environments is the C-language data type model used. The 64-bit data type uses the LP64 model where longs and pointers are 64-bits wide. All other fundamental data types remain the same as the in the 32-bit implementation. The 32-bit data type uses the ILP32 model where ints, longs, and pointers are 32-bits.

The following summary briefly describes the major features and considerations for using the 64-bit environment:

- **Large Virtual Address Space**

In the 64-bit environment, a process can have up to 64 bits of virtual address space, or 18 exabytes. This is 4 billion times the current 4 Gbyte maximum of a 32-bit process. Because of hardware restrictions, however, some platforms might not support the full 64 bits of address space.

Large address space increases the number of threads that can be created with the default stack size (1 Mbyte on 32 bits, 2 Mbytes on 64 bits). The number of threads with the default stack size is approximately 2000 threads on a 32-bit system and 8000 billion on a 64-bit system.

- **Kernel Memory Readers**

Because the kernel is an LP64 object that uses 64-bit data structures internally, existing 32-bit applications that use `libkvm`, `/dev/mem`, or `/dev/kmem` do not work properly and must be converted to 64-bit programs.

- **/proc Restrictions**

A 32-bit program that uses `/proc` is able to look at 32-bit processes, but is unable to understand a 64-bit process; the existing interfaces and data structures that describe the process are not large enough to contain the 64-bit quantities involved. Such programs must be recompiled as 64-bit programs to work for both 32-bit and 64-bit processes.

- **64-bit Libraries**

32-bit applications are required to link with 32-bit libraries, and 64-bit applications are required to link with 64-bit libraries. With the exception of those libraries that have become obsolete, all of the system libraries are provided in both 32-bit and 64-bit versions. However, no 64-bit libraries are provided in static form.

- 64-bit Arithmetic

Though 64-bit arithmetic has long been available in previous 32-bit Solaris releases, the 64-bit implementation now provides full 64-bit machine registers for integer operations and parameter passing.

- Large Files

If an application requires only large file support, then it can remain 32-bit and use the Large Files interface. It is, however, recommended that the application be converted to 64-bit to take full advantage of 64-bit capabilities.

Basic Threads Programming

The Threads Library

This chapter introduces the basic threads programming routines from the POSIX threads library, `libpthread(3T)`. This chapter covers *default threads*, or threads with default attribute values, which are the kind of threads that are most often used in multithreaded programming.

Chapter 3, Thread Create Attributes, explains how to create and use threads with nondefault attributes.

Note - Attributes are specified only at thread creation time; they cannot be altered when the thread is being used.

The POSIX (`libpthread`) routines introduced here have programming interfaces that are similar to the original (`libthread`) Solaris multithreading library.

The following brief roadmap directs you to the discussion of a particular task and its associated man page.

- “Create a Default Thread” on page 28
- “Wait for Thread Termination” on page 29
- “Detaching a Thread” on page 32
- “Create a Key for Thread-Specific Data” on page 33
- “Delete the Thread-Specific Data Key” on page 34
- “Set the Thread-Specific Data Key” on page 35
- “Get the Thread-Specific Data Key” on page 35
- “Get the Thread Identifier” on page 38

- “Compare Thread IDs” on page 39
- “Initializing Threads” on page 39
- “Yield Thread Execution” on page 40
- “Set the Thread Priority” on page 40
- “Get the Thread Priority” on page 41
- “Send a Signal to a Thread” on page 42
- “Access the Signal Mask of the Calling Thread” on page 43
- “Forking Safely” on page 44
- “Terminate a Thread” on page 44
- “Cancel a Thread” on page 47
- “Enable or Disable Cancellation” on page 48
- “Set Cancellation Type” on page 48
- “Create a Cancellation Point” on page 49
- “Push a Handler Onto the Stack” on page 50
- “Pull a Handler Off the Stack” on page 50

Create a Default Thread

When an attribute object is not specified, it is `NULL`, and the default thread is created with the following attributes:

- Unbound
- Nondetached
- With a default stack and stack size
- With the parent’s priority

You can also create a default attribute object with `pthread_attr_init()`, and then use this attribute object to create a default thread. See the section “Initialize Attributes” on page 55 for details.

`pthread_create(3THR)`

Use `pthread_create(3THR)` to add a new thread of control to the current process.

Prototype:

```
int pthread_create(pthread_t *tid, const pthread_attr_t *tattr,
    void*(*start_routine)(void *), void *arg);
```

```
#include <pthread.h>
```

```
pthread_attr_t ( )tattr;
pthread_t tid;
extern void *start_routine(void *arg);
void *arg;
int ret;

/* default behavior*/
ret = pthread_create(&tid, NULL, start_routine, arg);

/* initialized with default attributes */
ret = pthread_attr_init(&tattr);
/* default behavior specified*/
ret = pthread_create(&tid, &tattr, start_routine, arg);
```

The `pthread_create()` function is called with *attr* having the necessary state behavior. *start_routine* is the function with which the new thread begins execution. When *start_routine* returns, the thread exits with the exit status set to the value returned by *start_routine* (see “`pthread_create(3THR)`” on page 28).

When `pthread_create()` is successful, the ID of the thread created is stored in the location referred to as *tid*.

Creating a thread using a NULL attribute argument has the same effect as using a default attribute; both create a default thread. When *tattr* is initialized, it acquires the default behavior.

Return Values

`pthread_create()` returns a zero and exits when it completes successfully. Any other returned value indicates that an error occurred. When any of the following conditions are detected, `pthread_create()` fails and returns the corresponding value.

EAGAIN

A system limit is exceeded, such as when too many LWPs have been created.

EINVAL

The value of *tattr* is invalid.

Wait for Thread Termination

pthread_join(3THR)

Use `pthread_join(3THR)` to wait for a thread to terminate.

Prototype:

```
int pthread_join(thread_t tid, void **status);
```

```
#include <pthread.h>

pthread_t tid;
int ret;
int status;

/* waiting to join thread "tid" with status */
ret = pthread_join(tid, &status);

/* waiting to join thread "tid" without status */
ret = pthread_join(tid, NULL);
```

The `pthread_join( )` function blocks the calling thread until the specified thread terminates.

The specified thread must be in the current process and must not be detached. For information on thread detachment, see “Set Detach State” on page 57.

When *status* is not `NULL`, it points to a location that is set to the exit status of the terminated thread when `pthread_join( )` returns successfully.

Multiple threads cannot wait for the same thread to terminate. If they try to, one thread returns successfully and the others fail with an error of `ESRCH`.

After `pthread_join( )` returns, any stack storage associated with the thread can be reclaimed by the application.

Return Values

Returns a zero when it completes successfully. Any other returned value indicates that an error occurred. When any of the following conditions are detected, `pthread_join( )` fails and returns the corresponding value.

`ESRCH`

tid is not a valid, undetached thread in the current process.

`EDEADLK`

tid specifies the calling thread.

`EINVAL`

The value of *tid* is invalid.

The `pthread_join( )` routine takes two arguments, giving you some flexibility in its use. When you want the caller to wait until a specific thread terminates, supply that thread’s ID as the first argument.

If you are interested in the exit code of the defunct thread, supply the address of an area to receive it.

Remember that `pthread_join()` works only for target threads that are `nondetached`. When there is no reason to synchronize with the termination of a particular thread, then that thread should be detached.

Think of a detached thread as being the thread you use in most instances and reserve `nondetached` threads for only those situations that require them.

A Simple Threads Example

In Code Example 2-1, one thread executes the procedure at the top, creating a helper thread that executes the procedure `fetch()`, which involves a complicated database lookup and might take some time.

The main thread wants the results of the lookup but has other work to do in the meantime. So it does those other things and then waits for its helper to complete its job by executing `pthread_join()`.

An argument, *pbe*, to the new thread is passed as a stack parameter. This can be done here because the main thread waits for the spun-off thread to terminate. In general, though, it is better to use `malloc(3C)` to allocate storage from the heap instead of passing an address to thread stack storage, because this address might disappear or be reassigned if the thread terminated.

CODE EXAMPLE 2-1 A Simple Threads Program

```
void mainline (...)
{
    struct phonebookentry *pbe;
    pthread_attr_t tattr;
    pthread_t helper;
    int status;

    pthread_create(&helper, NULL, fetch, &pbe);

    /* do something else for a while */

    pthread_join(helper, &status);
    /* it's now safe to use result */
}

void fetch(struct phonebookentry *arg)
{
    struct phonebookentry *npbe;
    /* fetch value from a database */

    npbe = search (prog_name)
    if (npbe != NULL)
        *arg = *npbe;
    pthread_exit(0);
}

struct phonebookentry {
    char name[64];
```

```

        char phonenumber[32];
        char flags[16];
    }

```

Detaching a Thread

pthread_detach(3THR)

`pthread_detach(3THR)` is an alternative to `pthread_join(3THR)` to reclaim storage for a thread that is created with a *detachstate* attribute set to `PTHREAD_CREATE_JOINABLE`.

Prototype:

```
int pthread_detach(thread_t tid);
```

```
#include <pthread.h>
```

```
pthread_t tid;
int ret;
```

```
/* detach thread tid */
ret = pthread_detach(tid);
```

The `pthread_detach( )` function is used to indicate to the implementation that storage for the thread *tid* can be reclaimed when the thread terminates. If *tid* has not terminated, `pthread_detach( )` does not cause it to terminate. The effect of multiple `pthread_detach( )` calls on the same target thread is unspecified.

Return Values

`pthread_detach( )` returns a zero when it completes successfully. Any other returned value indicates that an error occurred. When any of the following conditions is detected, `pthread_detach( )` fails and returns the corresponding value.

EINVAL

tid is not a valid thread.

ESRCH

tid is not a valid, undetached thread in the current process.

Create a Key for Thread-Specific Data

Single-threaded C programs have two basic classes of data—local data and global data. For multithreaded C programs a third class is added—*thread-specific data* (TSD). This is very much like global data, except that it is private to a thread.

Thread-specific data is maintained on a per-thread basis. TSD is the only way to define and refer to data that is private to a thread. Each thread-specific data item is associated with a *key* that is global to all threads in the process. Using the *key*, a thread can access a pointer (*void **) that is maintained per-thread.

pthread_key_create(3THR)

Use `pthread_key_create(3THR)` to allocate a *key* that is used to identify thread-specific data in a process. The *key* is global to all threads in the process, and all threads initially have the value `NULL` associated with the *key* when it is created.

Call `pthread_key_create( )` is once for each *key* before using the *key*. There is no implicit synchronization.

Once a *key* has been created, each thread can bind a value to the *key*. The values are specific to the thread and are maintained for each thread independently. The per-thread binding is deallocated when a thread terminates if the *key* was created with a destructor function.

Prototype:

```
int pthread_key_create(pthread_key_t *key,
    void (*destructor) (void *));
```

```
#include <pthread.h>
```

```
pthread_key_t key;
int ret;
```

```
/* key create without destructor */
ret = pthread_key_create(&key, NULL);
```

```
/* key create with destructor */
ret = pthread_key_create(&key, destructor);
```

When `pthread_key_create( )` returns successfully, the allocated key is stored in the location pointed to by *key*. The caller must ensure that the storage and access to this key are properly synchronized.

An optional destructor function, *destructor*, can be used to free stale storage. When a key has a non-NULL destructor function and the thread has a non-NULL value associated with that key, the destructor function is called with the current associated value when the thread exits. The order in which the destructor functions are called is unspecified.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occur, `pthread_keycreate( )` fails and returns the corresponding value.

EAGAIN

The *key* name space is exhausted.

ENOMEM

Not enough virtual memory is available in this process to create a new key.

Delete the Thread-Specific Data Key

`pthread_key_delete(3THR)`

Use `pthread_key_delete(3THR)` to destroy an existing thread-specific data key. Any memory associated with the key can be freed because the key has been invalidated and will return an error if ever referenced. There is no comparable function in Solaris threads.

Prototype:

```
int pthread_key_delete(pthread_key_t key);
```

```
#include <pthread.h>
```

```
pthread_key_t key;  
int ret;
```

```
/* key previously created */  
ret = pthread_key_delete(key);
```

Once a *key* has been deleted, any reference to it with the `pthread_setspecific( )` or `pthread_getspecific( )` call results in the EINVAL error.

It is the responsibility of the programmer to free any thread-specific resources before calling the delete function. This function does not invoke any of the destructors.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, `pthread_key_create( )` fails and returns the corresponding value.

EINVAL

The *key* value is invalid.

Set the Thread-Specific Data Key

pthread_setspecific(3THR)

Use `pthread_setspecific(3THR)` to set the thread-specific binding to the specified thread-specific data key.

Prototype:

```
int pthread_setspecific(pthread_key_t key, const void *value);
```

```
#include <pthread.h>
```

```
pthread_key_t key;  
void *value;  
int ret;
```

```
/* key previously created */  
ret = pthread_setspecific(key, value);
```

Return Values

`pthread_setspecific()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occur, `pthread_setspecific()` fails and returns the corresponding value.

ENOMEM

Not enough virtual memory is available.

EINVAL

key is invalid.

Note - `pthread_setspecific()` does not free its storage. If a new binding is set, the existing binding must be freed; otherwise, a memory leak can occur.

Get the Thread-Specific Data Key

pthread_getspecific(3THR)

Use `pthread_getspecific(3THR)` to get the calling thread's binding for *key*, and store it in the location pointed to by *value*.

Prototype:
`int pthread_getspecific(pthread_key_t key);`

```
#include <pthread.h>

pthread_key_t key;
void *value;

/* key previously created */
value = pthread_getspecific(key);
```

Return Values

No errors are returned.

Global and Private Thread-Specific Data Example

Code Example 2–2 shows an excerpt from a multithreaded program. This code is executed by any number of threads, but it has references to two global variables, *errno* and *mywindow*, that really should be references to items private to each thread.

CODE EXAMPLE 2–2 Thread-Specific Data—Global but Private

```
body() {
    ...

    while (write(fd, buffer, size) == -1) {
        if (errno != EINTR) {
            fprintf(mywindow, "%s\n", strerror(errno));
            exit(1);
        }
    }

    ...
}
```

References to *errno* should get the system error code from the routine called by this thread, not by some other thread. So, references to *errno* by one thread refer to a different storage location than references to *errno* by other threads.

The *mywindow* variable is intended to refer to a *stdio* stream connected to a window that is private to the referring thread. So, as with *errno*, references to *mywindow* by one thread should refer to a different storage location (and, ultimately, a different window) than references to *mywindow* by other threads. The only difference here is that the threads library takes care of *errno*, but the programmer must somehow make this work for *mywindow*.

The next example shows how the references to *mywindow* work. The preprocessor converts references to *mywindow* into invocations of the `_mywindow( )` procedure.

This routine in turn invokes `pthread_getspecific()`, passing it the `mywindow_key` global variable (it really is a global variable) and an output parameter, `win`, that receives the identity of this thread's window.

CODE EXAMPLE 2-3 Turning Global References Into Private References

```
thread_key_t mywin_key;

FILE *_mywindow(void) {
    FILE *win;

    pthread_getspecific(mywin_key, &win);
    return(win);
}

#define mywindow _mywindow()

void routine_uses_win( FILE *win) {
    ...
}

void thread_start(...) {
    ...
    make_mywin();
    ...
    routine_uses_win( mywindow )
    ...
}
```

The `mywin_key` variable identifies a class of variables for which each thread has its own private copy; that is, these variables are thread-specific data. Each thread calls `make_mywin()` to initialize its window and to arrange for its instance of `mywindow` to refer to it.

Once this routine is called, the thread can safely refer to `mywindow` and, after `_mywindow()`, the thread gets the reference to its private window. So, references to `mywindow` behave as if they were direct references to data private to the thread.

Code Example 2-4 shows how to set this up.

CODE EXAMPLE 2-4 Initializing the Thread-Specific Data

```
void make_mywindow(void) {
    FILE **win;
    static pthread_once_t mykeycreated = PTHREAD_ONCE_INIT;

    pthread_once(&mykeycreated, mykeycreate);

    win = malloc(sizeof(*win));
    create_window(win, ...);

    pthread_setspecific(mywindow_key, win);
}

void mykeycreate(void) {
    pthread_keycreate(&mywindow_key, free_key);
}
```

```

}

void free_key(void *win) {
    free(win);
}

```

First, get a unique value for the key, *mywin_key*. This key is used to identify the thread-specific class of data. So, the first thread to call `make_mywin( )` eventually calls `pthread_keycreate( )`, which assigns to its first argument a unique *key*. The second argument is a **destructor** function that is used to deallocate a thread's instance of this thread-specific data item once the thread terminates.

The next step is to allocate the storage for the caller's instance of this thread-specific data item. Having allocated the storage, a call is made to the `create_window( )` routine, which sets up a window for the thread and sets the storage pointed to by *win* to refer to it. Finally, a call is made to `pthread_setspecific( )`, which associates the value contained in *win* (that is, the location of the storage containing the reference to the window) with the key.

After this, whenever this thread calls `pthread_getspecific( )`, passing the global *key*, it gets the value that was associated with this key by this thread when it called `pthread_setspecific( )`.

When a thread terminates, calls are made to the **destructor** functions that were set up in `pthread_key_create( )`. Each destructor function is called only if the terminating thread established a value for the *key* by calling `pthread_setspecific( )`.

Get the Thread Identifier

`pthread_self(3THR)`

Use `pthread_self(3THR)` to get the thread identifier of the calling thread.

Prototype:
`pthread_t pthread_self(void);`

```

#include <pthread.h>

pthread_t tid;

tid = pthread_self();

```

Return Values

Returns the thread identifier of the calling thread.

Compare Thread IDs

pthread_equal(3THR)

Use `pthread_equal(3THR)` to compare the thread identification numbers of two threads.

Prototype:

```
int pthread_equal(pthread_t tid1, pthread_t tid2);
```

```
#include <pthread.h>
```

```
pthread_t tid1, tid2;  
int ret;
```

```
ret = pthread_equal(tid1, tid2);
```

Return Values

Returns a nonzero value when *tid1* and *tid2* are equal; otherwise, zero is returned. When either *tid1* or *tid2* is an invalid thread identification number, the result is unpredictable.

Initializing Threads

pthread_once(3THR)

Use `pthread_once(3THR)` to call an initialization routine the first time `pthread_once(3T)` is called. Subsequent calls to `pthread_once( )` have no effect.

Prototype:

```
int pthread_once(pthread_once_t *once_control,  
void (*init_routine)(void));
```

```
#include <pthread.h>
```

```
pthread_once_t once_control = PTHREAD_ONCE_INIT;  
int ret;
```

```
ret = pthread_once(&once_control, init_routine);
```

The *once_control* parameter determines whether the associated initialization routine has been called.

Return Values

`pthread_once( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, `pthread_once( )` fails and returns the corresponding value.

EINVAL

once_control or *init_routine* is NULL.

Yield Thread Execution

`sched_yield(3RT)`

Use `sched_yield(3RT)` to cause the current thread to yield its execution in favor of another thread with the same or greater priority.

Prototype:
`int sched_yield(void);`

```
#include <sched.h>
```

```
int ret;
```

```
ret = sched_yield();
```

Return Values

Returns zero after completing successfully. Otherwise -1 is returned and *errno* is set to indicate the error condition.

ENOSYS

`sched_yield(3R)` is not supported in this implementation.

Set the Thread Priority

`pthread_setschedparam(3THR)`

Use `pthread_setschedparam(3THR)` to modify the priority of an existing thread. This function has no effect on scheduling policy.

Prototype:
`int pthread_setschedparam(pthread_t tid, int policy,
 const struct sched_param *param);`

```
#include <pthread.h>

pthread_t tid;
int ret;
struct sched_param param;
int priority;

/* sched_priority will be the priority of the thread */
sched_param.sched_priority = priority;

/* only supported policy, others will result in ENOTSUP */
policy = SCHED_OTHER;

/* scheduling parameters of target thread */
ret = pthread_setschedparam(tid, policy, &param);
```

Return Values

`pthread_setschedparam( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When either of the following conditions occurs, the `pthread_setschedparam( )` function fails and returns the corresponding value.

EINVAL

The value of the attribute being set is not valid.

ENOTSUP

An attempt was made to set the attribute to an unsupported value.

Get the Thread Priority

`pthread_getschedparam(3THR)`

`pthread_getschedparam(3THR)` gets the priority of the existing thread.

Prototype:

```
int pthread_getschedparam(pthread_t tid, int policy,
    struct schedparam *param);
```

```
#include <pthread.h>
```

```
pthread_t tid;
sched_param param;
int priority;
int policy;
int ret;
```

```
/* scheduling parameters of target thread */
ret = pthread_getschedparam (tid, &policy, &param);
```

```
/* sched_priority contains the priority of the thread */
priority = param.sched_priority;
```

Return Values

`pthread_getschedparam( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

ESRCH

The value specified by *tid* does not refer to an existing thread.

Send a Signal to a Thread

`pthread_kill(3THR)`

Use `pthread_kill(3THR)` to send a signal to a thread.

Prototype:

```
int pthread_kill(thread_t tid, int sig);
```

```
#include <pthread.h>
#include <signal.h>
```

```
int sig;
pthread_t tid;
int ret;
```

```
ret = pthread_kill(tid, sig);
```

`pthread_kill( )` sends the signal *sig* to the thread specified by *tid*. *tid* must be a thread within the same process as the calling thread. The *sig* argument must be from the list given in `signal(5)`.

When *sig* is zero, error checking is performed but no signal is actually sent. This can be used to check the validity of *tid*.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. When either of the following conditions occurs, `pthread_kill( )` fails and returns the corresponding value.

EINVAL

sig is not a valid signal number.

ESRCH

tid cannot be found in the current process.

Access the Signal Mask of the Calling Thread

pthread_sigmask(3THR)

Use `pthread_sigmask(3THR)` to change or examine the signal mask of the calling thread.

Prototype:

```
int pthread_sigmask(int how, const sigset_t *new, sigset_t *old);
```

```
#include <pthread.h>
#include <signal.h>
```

```
int ret;
sigset_t old, new;
```

```
ret = pthread_sigmask(SIG_SETMASK, &new, &old); /* set new mask */
ret = pthread_sigmask(SIG_BLOCK, &new, &old); /* blocking mask */
ret = pthread_sigmask(SIG_UNBLOCK, &new, &old); /* unblocking */
```

how determines how the signal set is changed. It can have one of the following values:

- **SIG_BLOCK**—Add *new* to the current signal mask, where *new* indicates the set of signals to block.
- **SIG_UNBLOCK**—Delete *new* from the current signal mask, where *new* indicates the set of signals to unblock.
- **SIG_SETMASK**—Replace the current signal mask with *new*, where *new* indicates the new signal mask.

When the value of *new* is `NULL`, the value of *how* is not significant and the signal mask of the thread is unchanged. So, to inquire about currently blocked signals, assign a `NULL` value to the *new* argument.

The *old* variable points to the space where the previous signal mask is stored, unless it is `NULL`.

Return Values

Returns a zero when it completes successfully. Any other returned value indicates that an error occurred. When the following condition occurs, `pthread_sigmask()` fails and returns the corresponding value.

`EINVAL`

The value of *how* is not defined.

Forking Safely

`pthread_atfork(3THR)`

See the discussion about `pthread_atfork(3THR)` in “The Solution—`pthread_atfork(3T)`” on page 146.

Prototype:

```
int pthread_atfork(void (*prepare) (void), void (*parent) (void),
    void (*child) (void) );
```

Terminate a Thread

`pthread_exit(3THR)`

Use `pthread_exit(3THR)` to terminate a thread.

Prototype:

```
void pthread_exit(void *status);
```

```
#include <pthread.h>
```

```
int status;
```

```
pthread_exit(&status); /* exit with status */
```

The `pthread_exit()` function terminates the calling thread. All thread-specific data bindings are released. If the calling thread is not detached, then the thread’s ID and the exit status specified by *status* are retained until the thread is waited for (blocked). Otherwise, *status* is ignored and the thread’s ID can be reclaimed immediately. For information on thread detachment, see “Set Detach State” on page 57.

Return Values

The calling thread terminates with its exit status set to the contents of *status* if *status* is not NULL.

Finishing Up

A thread can terminate its execution in the following ways:

- By returning from its first (outermost) procedure, the threads start routine; see `pthread_create(3T)`
- By calling `pthread_exit( )`, supplying an exit status
- By termination with POSIX cancel functions; see `pthread_cancel( )`

The default behavior of a thread is to linger until some other thread has acknowledged its demise by “joining” with it. This is the same as the default `pthread_create( )` attribute being non-detached; see `pthread_detach(3T)`. The result of the join is that the joining thread picks up the exit status of the dying thread and the dying thread vanishes.

An important special case arises when the initial thread — the one calling `main( )`,— returns from calling `main( )` or calls `exit(3C)`. This action causes the entire process to be terminated, along with all its threads. So take care to ensure that the initial thread does not return from `main( )` prematurely.

Note that when the main thread merely calls `pthread_exit(3T)`, it terminates only itself—the other threads in the process, as well as the process, continue to exist. (The process terminates when all threads terminate.)

Cancellation

The POSIX threads library introduces the ability to cancel threads to threads programming. Cancellation allows a thread to terminate the execution of any other thread, or all threads, in the process. Cancellation is an option when all further operations of a related set of threads are undesirable or unnecessary. A good method is to cancel all threads, restore the process to a consistent state, and then return to the point of origin.

One example of thread cancellation is an asynchronously generated cancel condition, such as, when a user requesting to close or exit some running application. Another example is the completion of a task undertaken by a number of threads. One of the threads might ultimately complete the task while the others continue to operate. Since they are serving no purpose at that point, they all should be cancelled.

There are dangers in performing cancellations. Most deal with properly restoring invariants and freeing shared resources. A thread that is cancelled without care might

leave a mutex in a locked state, leading to a deadlock. Or it might leave a region of memory allocated with no way to identify it and therefore no way to free it.

The `pthread`s library specifies a cancellation interface that permits or forbids cancellation programmatically. The library defines the set of points at which cancellation can occur (*cancellation points*). It also allows the scope of *cancellation handlers*, which provide clean up services, to be defined so that they are sure to operate when and where intended.

Placement of cancellation points and the effects of cancellation handlers must be based on an understanding of the application. A mutex is explicitly not a cancellation point and should be held only the minimal essential time.

Limit regions of asynchronous cancellation to sequences with no external dependencies that could result in dangling resources or unresolved state conditions. Take care to restore cancellation state when returning from some alternate, nested cancellation state. The interface provides features to facilitate restoration: `pthread_setcancelstate(3T)` preserves the current cancel state in a referenced variable; `pthread_setcanceltype(3T)` preserves the current cancel type in the same way.

Cancellations can occur under three different circumstances:

- Asynchronously
- At various points in the execution sequence as defined by the standard
- At discrete points specified by the application

By default, cancellation can occur only at well-defined points as defined by the POSIX standard.

In all cases, take care that resources and state are restored to a condition consistent with the point of origin.

Cancellation Points

Be careful to cancel a thread only when cancellation is safe. The `pthread`s standard specifies several cancellation points, including:

- Programmatically establish a thread cancellation point through a `pthread_testcancel(3T)` call.
- Threads waiting for the occurrence of a particular condition in `pthread_cond_wait(3T)` or `pthread_cond_timedwait(3T)`.
- Threads waiting for termination of another thread in `pthread_join(3T)`.
- Threads blocked on `sigwait(2)`.
- Some standard library calls. In general, these are functions in which threads can block; see the man page `cancellation(3T)` for a list.

Cancellation is enabled by default. At times you might want an application to disable cancellation. This has the result of deferring all cancellation requests until they are enabled again.

See “pthread_setcancelstate(3THR)” on page 48 for information about disabling cancellation.

Cancel a Thread

pthread_cancel(3THR)

Use pthread_cancel(3THR) to cancel a thread.

Prototype:

```
int pthread_cancel(pthread_t thread);
```

```
#include <pthread.h>
```

```
pthread_t thread;  
int ret;
```

```
ret = pthread_cancel(thread);
```

How the cancellation request is treated depends on the state of the target thread. Two functions, pthread_setcancelstate(3T) and pthread_setcanceltype(3T), determine that state.

Return Values

pthread_cancel() returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

ESRCH

No thread could be found corresponding to that specified by the given thread ID.

Enable or Disable Cancellation

`pthread_setcancelstate(3THR)`

Use `pthread_setcancelstate(3THR)` to enable or disable thread cancellation. When a thread is created, thread cancellation is enabled by default.

Prototype:

```
int pthread_setcancelstate(int state, int *oldstate);

#include <pthread.h>

int oldstate;
int ret;

/* enabled */
ret = pthread_setcancelstate(PTHREAD_CANCEL_ENABLE, &oldstate);

/* disabled */
ret = pthread_setcancelstate(PTHREAD_CANCEL_DISABLE, &oldstate);
```

Return Values

`pthread_setcancelstate()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the `pthread_setcancelstate()` function fails and returns the corresponding value.

EINVAL

The state is not `PTHREAD_CANCEL_ENABLE` or `PTHREAD_CANCEL_DISABLE`.

Set Cancellation Type

`pthread_setcanceltype(3THR)`

Use `pthread_setcanceltype(3THR)` to set the cancellation type to either deferred or asynchronous mode. When a thread is created, the cancellation type is set to deferred mode by default. In deferred mode, the thread can be cancelled only at cancellation points. In asynchronous mode, a thread can be cancelled at any point during its execution. Using asynchronous mode is discouraged.

Prototype:

```
int pthread_setcanceltype(int type, int *oldtype);
```

```
#include <pthread.h>

int oldtype;
int ret;

/* deferred mode */
ret = pthread_setcanceltype(PTHREAD_CANCEL_DEFERRED, &oldtype);

/* async mode*/
ret = pthread_setcanceltype(PTHREAD_CANCEL_ASYNCHRONOUS, &oldtype);
```

Return Values

`pthread_setcanceltype( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The type is not `PTHREAD_CANCEL_DEFERRED` or `PTHREAD_CANCEL_ASYNCHRONOUS`.

Create a Cancellation Point

`pthread_testcancel(3THR)`

Use `pthread_testcancel(3THR)` to establish a cancellation point for a thread.

Prototype:

```
void pthread_testcancel(void);
```

```
#include <pthread.h>

pthread_testcancel();
```

The `pthread_testcancel( )` function is effective when thread cancellation is enabled and in deferred mode. Calling this function while cancellation is disabled has no effect.

Be careful to insert `pthread_testcancel( )` only in sequences where it is safe to cancel a thread. In addition to programmatically establishing cancellation points through the `pthread_testcancel( )` call, the pthreads standard specifies several cancellation points. See “Cancellation Points” on page 46 for more details.

There is no return value.

Push a Handler Onto the Stack

Use cleanup handlers to restore conditions to a state consistent with that at the point of origin, such as cleaning up allocated resources and restoring invariants. Use the `pthread_cleanup_push(3T)` and `pthread_cleanup_pop(3T)` functions to manage the handlers.

Cleanup handlers are pushed and popped in the same lexical scope of a program. They should always match; otherwise compiler errors will be generated.

`pthread_cleanup_push(3THR)`

Use `pthread_cleanup_push(3THR)` to push a cleanup handler onto a cleanup stack (LIFO).

Prototype:

```
void pthread_cleanup_push(void(*routine)(void *), void *args);
```

```
#include <pthread.h>
```

```
/* push the handler "routine" on cleanup stack */
pthread_cleanup_push (routine, arg);
```

Pull a Handler Off the Stack

`pthread_cleanup_pop(3THR)`

Use `pthread_cleanup_pop(3THR)` to pull the cleanup handler off the cleanup stack.

A nonzero argument in the pop function removes the handler from the stack and executes it. An argument of zero pops the handler without executing it.

`pthread_cleanup_pop( )` is effectively called with a nonzero argument if a thread either explicitly or implicitly calls `pthread_exit(3T)` or if the thread accepts a cancel request.

Prototype:

```
void pthread_cleanup_pop(int execute);
```

```
#include <pthread.h>
```

```
/* pop the "func" out of cleanup stack and execute "func" */
pthread_cleanup_pop (1);
```

```
/* pop the "func" and DONT execute "func" */
pthread_cleanup_pop (0);
```

There are no return values.

Thread Create Attributes

The previous chapter covered the basics of threads creation using default attributes. This chapter discusses setting attributes at thread creation time.

Note that only pthreads uses attributes and cancellation, so the API covered in this chapter is for POSIX threads only. Otherwise, the *functionality* for Solaris threads and pthreads is largely the same. (See Chapter 9, Programming With Solaris Threads, for more information about similarities and differences.)

- “Initialize Attributes” on page 55
- “Destroy Attributes” on page 56
- “Set Detach State” on page 57
- “Get Detach State” on page 58
- “Set Stack Guard Size” on page 59
- “Get Stack Guard Size” on page 60
- “Set Scope” on page 60
- “Get Scope” on page 61
- “Set Thread Concurrency Level” on page 62
- “Get Thread Concurrency Level” on page 63
- “Set Scheduling Policy” on page 63
- “Get Scheduling Policy” on page 65
- “Set Inherited Scheduling Policy” on page 65
- “Get Inherited Scheduling Policy” on page 66
- “Set Scheduling Parameters” on page 67
- “Get Scheduling Parameters” on page 68
- “Set Stack Size” on page 69

- “Get Stack Size” on page 70
- “Set Stack Address” on page 72
- “Get Stack Address” on page 74

Attributes

Attributes are a way to specify behavior that is different from the default. When a thread is created with `pthread_create(3T)` or when a synchronization variable is initialized, an attribute object can be specified. The defaults are usually sufficient.

Note - Attributes are specified only at thread creation time; they cannot be altered while the thread is being used.

An attribute object is opaque, and cannot be directly modified by assignments. A set of functions is provided to initialize, configure, and destroy each object type.

Once an attribute is initialized and configured, it has process-wide scope. The suggested method for using attributes is to configure all required state specifications at one time in the early stages of program execution. The appropriate attribute object can then be referred to as needed.

Using attribute objects has two primary advantages.

- First, it adds to code portability.

Even though supported attributes might vary between implementations, you need not modify function calls that create thread entities because the attribute object is hidden from the interface.

If the target port supports attributes that are not found in the current port, provision must be made to manage the new attributes. This is an easy porting task though, because attribute objects need only be initialized once in a well-defined location.

- Second, state specification in an application is simplified.

As an example, consider that several sets of threads might exist within a process, each providing a separate service, and each with its own state requirements.

At some point in the early stages of the application, a thread attribute object can be initialized for each set. All future thread creations will then refer to the attribute object initialized for that type of thread. The initialization phase is simple and localized, and any future modifications can be made quickly and reliably.

Attribute objects require attention at process exit time. When the object is initialized, memory is allocated for it. This memory must be returned to the system. The `pthreads` standard provides function calls to destroy attribute objects.

Initialize Attributes

pthread_attr_init(3THR)

Use `pthread_attr_init(3THR)` to initialize object attributes to their default values. The storage is allocated by the thread system during execution.

Prototype:

```
int pthread_attr_init(pthread_attr_t *tattr);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;  
int ret;
```

```
/* initialize an attribute to the default value */  
ret = pthread_attr_init(&tattr);
```

Table 3–1 shows the default values for attributes (*tattr*) .

TABLE 3–1 Default Attribute Values for *tattr*

Attribute	Value	Result
<i>scope</i>	PTHREAD_SCOPE_PROCESS	New thread is unbound – not permanently attached to LWP.
<i>detachstate</i>	PTHREAD_CREATE_JOINABLE	Exit status and thread are preserved after the thread terminates.
<i>stackaddr</i>	NULL	New thread has system-allocated stack address.
<i>stacksize</i>	1 megabyte	New thread has system-defined stack size.
<i>priority</i>		New thread inherits parent thread priority.

TABLE 3-1 Default Attribute Values for *tattr* (continued)

Attribute	Value	Result
<i>inheritsched</i>	PTHREAD_INHERIT_SCHED	New thread inherits parent thread scheduling priority.
<i>schedpolicy</i>	SCHED_OTHER	New thread uses Solaris-defined fixed priority scheduling; threads run until preempted by a higher-priority thread or until they block or yield.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

ENOMEM

Returned when there is not enough memory to initialize the thread attributes object.

Destroy Attributes

pthread_attr_destroy(3THR)

Use `pthread_attr_destroy(3THR)` to remove the storage allocated during initialization. The attribute object becomes invalid.

Prototype:

```
int pthread_attr_destroy(pthread_attr_t *tattr);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;
int ret;
```

```
/* destroy an attribute */
ret = pthread_attr_destroy(&tattr);
```

Return Values

`pthread_attr_destroy()` returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

Indicates that the value of *tattr* was not valid.

Set Detach State

`pthread_attr_setdetachstate(3THR)`

When a thread is created detached (`PTHREAD_CREATE_DETACHED`), its thread *ID* and other resources can be reused as soon as the thread terminates. Use `pthread_attr_setdetachstate(3THR)` when the calling thread does not want to wait for the thread to terminate.

When a thread is created nondetached (`PTHREAD_CREATE_JOINABLE`), it is assumed that you will be waiting for it. That is, it is assumed that you will be executing a `pthread_join(3T)()` on the thread.

Whether a thread is created detached or nondetached, the process does not exit until all threads have exited. See “Finishing Up” on page 45 for a discussion of process termination caused by premature exit from `main()`.

Prototype:

```
int pthread_attr_setdetachstate(pthread_attr_t *tattr, int detachstate);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;  
int ret;
```

```
/* set the thread detach state */  
ret = pthread_attr_setdetachstate(&tattr, PTHREAD_CREATE_DETACHED);
```

Note - When there is no explicit synchronization to prevent it, a newly created, detached thread can die and have its thread ID reassigned to another new thread before its creator returns from `pthread_create()`.

For nondetached (`PTHREAD_CREATE_JOINABLE`) threads, it is very important that some thread join with it after it terminates—otherwise the resources of that thread

are not released for use by new threads. This commonly results in a memory leak. So when you do not want a thread to be joined, create it as a detached thread.

CODE EXAMPLE 3-1 Creating a Detached Thread

```
#include <pthread.h>

pthread_attr_t tattr;
pthread_t tid;
void *start_routine;
void arg
int ret;

/* initialized with default attributes */
ret = pthread_attr_init(&tattr);
ret = pthread_attr_setdetachstate(&tattr, PTHREAD_CREATE_DETACHED);
ret = pthread_create(&tid, &tattr, start_routine, arg);
```

Return Values

`pthread_attr_setdetachstate()` returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

Indicates that the value of *detachstate* or *tattr* was not valid.

Get Detach State

pthread_attr_getdetachstate(3THR)

Use `pthread_attr_getdetachstate(3THR)` to retrieve the thread create state, which can be either detached or joined.

Prototype:

```
int pthread_attr_getdetachstate(const pthread_attr_t *tattr,
                               int *detachstate);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;
int detachstate;
int ret;
```

```
/* get detachstate of thread */
ret = pthread_attr_getdetachstate (&tattr, &detachstate);
```

Return Values

`pthread_attr_getdetachstate()` returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

Indicates that the value of *detachstate* is NULL or *tattr* is invalid.

Set Stack Guard Size

`pthread_attr_setguardsize(3THR)`

`pthread_attr_setguardsize(3THR)` sets the *guardsize* of the *attr* object.

The *guardsize* argument provides protection against overflow of the stack pointer. If a thread's stack is created with guard protection, the implementation allocates extra memory at the overflow end of the stack as a buffer against stack overflow of the stack pointer. If an application overflows into this buffer an error results (possibly in a SIGSEGV signal being delivered to the thread).

The *guardsize* attribute is provided to the application for two reasons:

1. Overflow protection can potentially result in wasted system resources. An application that creates a large number of threads, and knows its threads will never overflow their stack, can save system resources by turning off guard areas.
2. When threads allocate large data structures on stack, a large guard area may be needed to detect stack overflow.

If *guardsize* is zero, a guard area will not be provided for threads created with *attr*. If *guardsize* is greater than zero, a guard area of at least size *guardsize* bytes is provided for each thread created with *attr*. By default, a thread has an implementation-defined, non-zero guard area.

A conforming implementation is permitted to round up the value contained in *guardsize* to a multiple of the configurable system variable `PAGESIZE` (see `PAGESIZE` in `sys/mman.h`). If an implementation rounds up the value of *guardsize* to a multiple of `PAGESIZE`, a call to `pthread_attr_getguardsize()` specifying *attr* will store, in *guardsize*, the guard size specified in the previous call to `pthread_attr_setguardsize()`.

```
#include <pthread.h>
```

```
int pthread_attr_setguardsize(pthread_attr_t *attr, size_t guardsize);
```

Return Value

`pthread_attr_setguardsize()` fails if:

EINVAL

The argument *attr* is invalid, the argument *guardsize* is invalid, or the argument *guardsize* contains an invalid value.

Get Stack Guard Size

pthread_attr_getguardsize(3THR)

`pthread_attr_getguardsize(3THR)` gets the *guardsize* of the *attr* object.

A conforming implementation is permitted to round up the value contained in *guardsize* to a multiple of the configurable system variable `PAGESIZE` (see `PAGESIZE` in `sys/mman.h`). If an implementation rounds up the value of *guardsize* to a multiple of `PAGESIZE`, a call to `pthread_attr_getguardsize( )` specifying *attr* will store, in *guardsize*, the guard size specified in the previous call to `pthread_attr_setguardsize( )`.

```
#include <pthread.h>
```

```
int pthread_attr_getguardsize(const pthread_attr_t *attr, size_t *guardsize);
```

Return Value

`pthread_attr_getguardsize( )` fails if:

EINVAL

The argument *attr* is invalid, the argument *guardsize* is invalid, or the argument *guardsize* contains an invalid value,.

Set Scope

pthread_attr_setscope(3THR)

Use `pthread_attr_setscope(3THR)` to create a bound thread (`PTHREAD_SCOPE_SYSTEM`) or an unbound thread (`PTHREAD_SCOPE_PROCESS`).

Note - Both thread types are accessible only within a given process.

Prototype:

```
int pthread_attr_setscope(pthread_attr_t *tattr, int scope);
```

```

#include <pthread.h>

pthread_attr_t tattr;
int ret;

/* bound thread */
ret = pthread_attr_setscope(&tattr, PTHREAD_SCOPE_SYSTEM);

/* unbound thread */
ret = pthread_attr_setscope(&tattr, PTHREAD_SCOPE_PROCESS);

```

Notice that there are three function calls in this example: one to initialize the attributes, one to set any variations from the default attributes, and one to create the pthreads.

```

#include <pthread.h>

pthread_attr_t attr;
pthread_t tid;
void start_routine;
void arg;
int ret;

/* initialized with default attributes */
ret = pthread_attr_init (&tattr);

/* BOUND behavior */
ret = pthread_attr_setscope(&tattr, PTHREAD_SCOPE_SYSTEM);
ret = pthread_create (&tid, &tattr, start_routine, arg);

```

Return Values

`pthread_attr_setscope( )` returns zero after completing *successfully*. Any other returned value indicates that an error occurred. If the following conditions occur, the function fails and returns the corresponding value.

EINVAL

An attempt was made to set *tattr* to a value that is not valid.

Get Scope

`pthread_attr_getscope(3THR)`

Use `pthread_attr_getscope(3THR)` to retrieve the thread scope, which indicates whether the thread is bound or unbound.

Prototype:

```
int pthread_attr_getscope(pthread_attr_t *tattr, int *scope);
```

```
#include <pthread.h>

pthread_attr_t tattr;
int scope;
int ret;

/* get scope of thread */
ret = pthread_attr_getscope(&tattr, &scope);
```

Return Values

`pthread_attr_getscope( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value of *scope* is NULL or *tattr* is invalid.

Set Thread Concurrency Level

`pthread_setconcurrency(3THR)`

Unbound threads in a process might or might not be required to be simultaneously active. By default, the threads implementation ensures that a sufficient number of threads are active so that the process can continue to make progress. While this conserves system resources, it might not produce the most effective level of concurrency.

`pthread_setconcurrency(3THR)` allows an application to inform the threads implementation of its desired concurrency level, *new_level*. The actual level of concurrency provided by the implementation as a result of this function call is unspecified. (For Solaris threads, see “`thr_setconcurrency(3THR)`” on page 196.)

If *new_level* is zero, the implementation maintains the concurrency level at its discretion as if `pthread_setconcurrency( )` was never called.

When an application calls `pthread_setconcurrency( )`, it is informing the implementation of its desired concurrency level. The implementation uses this as a hint, not a requirement.

```
#include <pthread.h>

int pthread_setconcurrency(int new_level);
```

Return Value

`pthread_setconcurrency( )` fails if:

EINVAL

The value specified by *new_level* is negative.

EAGAIN

The value specified by *new_level* would cause a system resource to be exceeded.

Get Thread Concurrency Level

pthread_getconcurrency(3THR)

pthread_getconcurrency(3THR) returns the value set by a previous call to pthread_setconcurrency(). If the pthread_setconcurrency() function was not previously called, pthread_getconcurrency() returns zero to indicate that the implementation is maintaining the concurrency level. (For Solaris threads, see “thr_getconcurrency(3THR)” on page 197.)

```
#include <pthread.h>
```

```
int pthread_getconcurrency(void);
```

Return Value

pthread_getconcurrency() always returns the concurrency level set by a previous call to pthread_setconcurrency(). If pthread_setconcurrency() has never been called, pthread_getconcurrency() returns zero.

Set Scheduling Policy

pthread_attr_setschedpolicy(3THR)

Use pthread_attr_setschedpolicy(3THR) to set the scheduling policy. The POSIX draft standard specifies scheduling policy attributes of SCHED_FIFO (first-in-first-out), SCHED_RR (round-robin), or SCHED_OTHER (an implementation-defined method).

■ SCHED_FIFO

First-In-First-Out; threads scheduled to this policy, if not preempted by a higher priority, will proceed until completion. Threads whose contention scope is system (PTHREAD_SCOPE_SYSTEM) are in real-time (RT) scheduling class and the calling process must have an effective user id of 0. SCHED_FIFO for threads that have a

contention scope of process (PTHREAD_SCOPE_PROCESS) is based on the TS scheduling class.

■ SCHED_RR

Round-Robin; threads scheduled to this policy, if not preempted by a higher priority, will execute for a time period determined by the system. Threads whose contention scope is system (PTHREAD_SCOPE_SYSTEM) are in real-time (RT) scheduling class and the calling process must have an effective user id of 0. SCHED_RR for threads that have a contention scope of process (PTHREAD_SCOPE_PROCESS) is based on the TS scheduling class.

SCHED_FIFO and SCHED_RR are optional in POSIX, and are supported for real time bound threads only.

Currently, only the Solaris SCHED_OTHER, time-sharing, default value is supported in pthreads. For a discussion of scheduling, see the section “Scheduling” on page 23.

Prototype:

```
int pthread_attr_setschedpolicy(pthread_attr_t *tattr, int policy);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;  
int policy;  
int ret;
```

```
/* set the scheduling policy to SCHED_OTHER */  
ret = pthread_attr_setschedpolicy(&tattr, SCHED_OTHER);
```

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. When either of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

An attempt was made to set *tattr* to a value that is not valid.

ENOTSUP

An attempt was made to set the attribute to an unsupported value.

Get Scheduling Policy

pthread_attr_getschedpolicy(3THR)

Use `pthread_attr_getschedpolicy(3THR)` to retrieve the scheduling policy. Currently, only the Solaris-based `SCHED_OTHER` default value is supported in pthreads.

Prototype:

```
int pthread_attr_getschedpolicy(pthread_attr_t *tattr, int *policy);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;  
int policy;  
int ret;
```

```
/* get scheduling policy of thread */  
ret = pthread_attr_getschedpolicy (&tattr, &policy);
```

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The parameter *policy* is NULL or *tattr* is invalid.

Set Inherited Scheduling Policy

pthread_attr_setinheritsched(3THR)

Use `pthread_attr_setinheritsched(3THR)` to set the inherited scheduling policy.

An *inherit* value of `PTHREAD_INHERIT_SCHED` (the default) means that the scheduling policies defined in the creating thread are to be used, and any scheduling attributes defined in the `pthread_create( )` call are to be ignored. If `PTHREAD_EXPLICIT_SCHED` is used, the attributes from the `pthread_create( )` call are to be used.

Prototype:

```

int pthread_attr_setinheritsched(pthread_attr_t *tattr, int inherit);

#include <pthread.h>

pthread_attr_t tattr;
int inherit;
int ret;

/* use the current scheduling policy */
ret = pthread_attr_setinheritsched(&tattr, PTHREAD_EXPLICIT_SCHED);

```

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. When either of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

An attempt was made to set *tattr* to a value that is not valid.

ENOTSUP

An attempt was made to set the attribute to an unsupported value.

Get Inherited Scheduling Policy

pthread_attr_getinheritsched(3THR)

`pthread_attr_getinheritsched(3THR)` returns the scheduling policy set by `pthread_attr_setinheritsched()`.

Prototype:

```

int pthread_attr_getinheritsched(pthread_attr_t *tattr, int *inherit);

#include <pthread.h>

pthread_attr_t tattr;
int inherit;
int ret;

/* get scheduling policy and priority of the creating thread */
ret = pthread_attr_getinheritsched (&tattr, &inherit);

```

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The parameter *inherit* is NULL or *tattr* is invalid.

Set Scheduling Parameters

pthread_attr_setschedparam(3THR)

pthread_attr_setschedparam(3THR) sets the scheduling parameters.

Scheduling parameters are defined in the `param` structure; only priority is supported. Newly created threads run with this priority.

■ SCHED_FIFO

First-In-First-Out; threads scheduled to this policy, if not preempted by a higher priority, will proceed until completion. Threads that have a contention scope of system (PTHREAD_SCOPE_SYSTEM) are in real-time (RT) scheduling class and the calling process must have an effective user id of 0. SCHED_FIFO for threads that have a contention scope of process (PTHREAD_SCOPE_PROCESS) is based on the TS scheduling class.

■ SCHED_RR

Round-Robin; threads scheduled to this policy, if not pre-empted by a higher priority, will execute for a time period determined by the system. Threads that have a contention scope of system (PTHREAD_SCOPE_SYSTEM) are in real-time (RT) scheduling class and the calling process must have an effective user id of 0. SCHED_RR for threads that have a contention scope of process (PTHREAD_SCOPE_PROCESS) is based on the TS scheduling class.

Prototype:

```
int pthread_attr_setschedparam(pthread_attr_t *tattr,  
    const struct sched_param *param);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;  
int newprio;  
sched_param param;  
newprio = 30;
```

```

/* set the priority; others are unchanged */
param.sched_priority = newprio;

/* set the new scheduling param */
ret = pthread_attr_setschedparam (&tattr, &param);

```

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following conditions occur, the function fails and returns the corresponding value.

EINVAL

The value of *param* is NULL or *tattr* is invalid.

You can manage pthreads priority two ways. You can set the priority attribute before creating a child thread, or you can change the priority of the parent thread and then change it back.

Get Scheduling Parameters

pthread_attr_getschedparam(3THR)

pthread_attr_getschedparam(3THR) returns the scheduling parameters defined by pthread_attr_setschedparam().

Prototype:

```

int pthread_attr_getschedparam(pthread_attr_t *tattr,
    const struct sched_param *param);

```

```

#include <pthread.h>

```

```

pthread_attr_t attr;
struct sched_param param;
int ret;

```

```

/* get the existing scheduling param */
ret = pthread_attr_getschedparam (&tattr, &param);

```

Return Values

pthread_attr_getschedparam() returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value of *param* is NULL or *tattr* is invalid.

Creating a Thread With a Specified Priority

You can set the priority attribute before creating the thread. The child thread is created with the new priority that is specified in the `sched_param` structure (this structure also contains other scheduling information).

It is always a good idea to get the existing parameters, change the priority, xxx the thread, and then reset the priority.

Code Example 3–2 shows an example of this.

CODE EXAMPLE 3–2 Creating a Prioritized Thread

```
#include <pthread.h>
#include <sched.h>

pthread_attr_t tattr;
pthread_t tid;
int ret;
int newprio = 20;
sched_param param;

/* initialized with default attributes */
ret = pthread_attr_init (&tattr);

/* safe to get existing scheduling param */
ret = pthread_attr_getschedparam (&tattr, &param);

/* set the priority; others are unchanged */
param.sched_priority = newprio;

/* setting the new scheduling param */
ret = pthread_attr_setschedparam (&tattr, &param);

/* with new priority specified */
ret = pthread_create (&tid, &tattr, func, arg);
```

Set Stack Size

`pthread_attr_setstacksize(3THR)`

`pthread_attr_setstacksize(3THR)` sets the thread stack size.

The *stacksize* attribute defines the size of the stack (in bytes) that the system will allocate. The size should not be less than the system-defined minimum stack size. See “About Stacks” on page 71 for more information.

Prototype:

```
int pthread_attr_setstacksize(pthread_attr_t *tattr, int size);
```

```
#include <pthread.h>

pthread_attr_t tattr;
int size;
int ret;

size = (PTHREAD_STACK_MIN + 0x4000);

/* setting a new size */
ret = pthread_attr_setstacksize(&tattr, size);
```

In the example above, *size* contains the number of bytes for the stack that the new thread uses. If *size* is zero, a default size is used. In most cases, a zero value works best.

PTHREAD_STACK_MIN is the amount of stack space required to start a thread. This does not take into consideration the threads routine requirements that are needed to execute application code.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value returned is less than the value of PTHREAD_STACK_MIN, or exceeds a system-imposed limit, or *tattr* is not valid.

Get Stack Size

pthread_attr_getstacksize(3THR)

pthread_attr_getstacksize(3THR) returns the stack size set by pthread_attr_setstacksize().

Prototype:

```
int pthread_attr_getstacksize(pthread_attr_t *tattr, size_t *size);
```

```
#include <pthread.h>

pthread_attr_t tattr;
int size;
int ret;
```

```
/* getting the stack size */
ret = pthread_attr_getstacksize(&tattr, &size);
```

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value returned is less than the value of PTHREAD_STACK_MIN, or exceeds a system-imposed limit.

About Stacks

Typically, thread stacks begin on page boundaries and any specified size is rounded up to the next page boundary. A page with no access permission is appended to the top of the stack so that most stack overflows result in sending a SIGSEGV signal to the offending thread. Thread stacks allocated by the caller are used as is.

When a stack is specified, the thread should also be created PTHREAD_CREATE_JOINABLE. That stack cannot be freed until the pthread_join(3T) call for that thread has returned, because the thread's stack cannot be freed until the thread has terminated. The only reliable way to know if such a thread has terminated is through pthread_join(3T).

Generally, you do not need to allocate stack space for threads. The threads library allocates 1 Mbyte of virtual memory for each thread's stack with no swap space reserved. (The library uses the MAP_NORESERVE option of mmap() to make the allocations.)

Each thread stack created by the threads library has a red zone. The library creates the red zone by appending a page to the top of a stack to catch stack overflows. This page is invalid and causes a memory fault if it is accessed. Red zones are appended to all automatically allocated stacks whether the size is specified by the application or the default size is used.

Note - Because runtime stack requirements vary, you should be absolutely certain that the specified stack will satisfy the runtime requirements needed for library calls and dynamic linking.

There are very few occasions when it is appropriate to specify a stack, its size, or both. It is difficult even for an expert to know if the right size was specified. This is because even a program compliant with ABI standards cannot determine its stack size statically. Its size is dependent on the needs of the particular runtime environment in which it executes.

Building Your Own Stack

When you specify the size of a thread stack, be sure to account for the allocations needed by the invoked function and by each function called. The accounting should include calling sequence needs, local variables, and information structures.

Occasionally you want a stack that is a bit different from the default stack. An obvious situation is when the thread needs more than one megabyte of stack space. A less obvious situation is when the default stack is too large. You might be creating thousands of threads and not have enough virtual memory to handle the gigabytes of stack space that this many default stacks require.

The limits on the maximum size of a stack are often obvious, but what about the limits on its minimum size? There must be enough stack space to handle all of the stack frames that are pushed onto the stack, along with their local variables, and so on.

You can get the absolute minimum limit on stack size by calling the macro `PTHREAD_STACK_MIN`, which returns the amount of stack space required for a thread that executes a `NULL` procedure. Useful threads need more than this, so be very careful when reducing the stack size.

```
#include <pthread.h>

pthread_attr_t tattr;
pthread_t tid;
int ret;

int size = PTHREAD_STACK_MIN + 0x4000;

/* initialized with default attributes */
ret = pthread_attr_init(&tattr);

/* setting the size of the stack also */
ret = pthread_attr_setstacksize(&tattr, size);

/* only size specified in tattr */
ret = pthread_create(&tid, &tattr, start_routine, arg);
```

When you allocate your own stack, be sure to append a red zone to its end by calling `mprotect(2)`.

Set Stack Address

`pthread_attr_setstackaddr(3THR)`

`pthread_attr_setstackaddr(3THR)` sets the thread stack address.

The *stackaddr* attribute defines the base of the thread's stack. If this is set to non-null (`NULL` is the default) the system initializes the stack at that address.

Prototype:

```
int pthread_attr_setstackaddr(pthread_attr_t *tattr, void *stackaddr);
```

```
#include <pthread.h>
```

```
pthread_attr_t tattr;
void *base;
int ret;
```

```
base = (void *) malloc(PTHREAD_STACK_MIN + 0x4000);
```

```
/* setting a new address */
ret = pthread_attr_setstackaddr(&tattr, base);
```

In the previous example, *base* contains the address for the stack that the new thread uses. If *base* is NULL, then `pthread_create(3T)` allocates a stack for the new thread with at least `PTHREAD_STACK_MIN` bytes.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value of *base* or *tattr* is incorrect.

This example shows how to create a thread with a custom stack address.

```
#include <pthread.h>
```

```
pthread_attr_t tattr;
pthread_t tid;
int ret;
void *stackbase;
```

```
stackbase = (void *) malloc(size);
```

```
/* initialized with default attributes */
ret = pthread_attr_init(&tattr);
```

```
/* setting the base address in the attribute */
ret = pthread_attr_setstackaddr(&tattr, stackbase);
```

```
/* only address specified in attribute tattr */
ret = pthread_create(&tid, &tattr, func, arg);
```

This example shows how to create a thread with both a custom stack address and a custom stack size.

```
#include <pthread.h>
```

```

pthread_attr_t tattr;
pthread_t tid;
int ret;
void *stackbase;

int size = PTHREAD_STACK_MIN + 0x4000;
stackbase = (void *) malloc(size);

/* initialized with default attributes */
ret = pthread_attr_init(&tattr);

/* setting the size of the stack also */
ret = pthread_attr_setstacksize(&tattr, size);

/* setting the base address in the attribute */
ret = pthread_attr_setstackaddr(&tattr, stackbase);

/* address and size specified */
ret = pthread_create(&tid, &tattr, func, arg);

```

Get Stack Address

pthread_attr_getstackaddr(3THR)

pthread_attr_getstackaddr(3THR) returns the thread stack address set by pthread_attr_setstackaddr().

Prototype:

```
int pthread_attr_getstackaddr(pthread_attr_t *tattr, void * *stackaddr);
```

```
#include <pthread.h>
```

```

pthread_attr_t tattr;
void *base;
int ret;

/* getting a new address */
ret = pthread_attr_getstackaddr (&tattr, *base);

```

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value or *base* or *tattr* is incorrect.

Programming with Synchronization Objects

This chapter describes the synchronization types available with threads and discusses when and how to use synchronization.

- “Mutual Exclusion Lock Attributes” on page 76
- “Using Mutual Exclusion Locks” on page 93
- “Condition Variable Attributes” on page 106
- “Using Condition Variables” on page 110
- “Semaphores” on page 122
- “Read-Write Lock Attributes” on page 130
- “Set Mutex Attribute’s Protocol” on page 83
- “Synchronization Across Process Boundaries” on page 139
- “Interprocess Locking Without the Threads Library” on page 141
- “Comparing Primitives” on page 141

Synchronization objects are variables in memory that you access just like data. Threads in different processes can communicate with each other through synchronization objects placed in threads-controlled shared memory, even though the threads in different processes are generally invisible to each other.

Synchronization objects can also be placed in files and can have lifetimes beyond that of the creating process.

The available types of synchronization objects are:

- Mutex Locks
- Condition Variables
- Semaphores

Here are situations that can benefit from the use of synchronization:

- When synchronization is the only way to ensure consistency of shared data.
- When threads in two or more processes can use a single synchronization object jointly. Note that the synchronization object should be initialized by only one of the cooperating processes, because reinitializing a synchronization object sets it to the *unlocked* state.
- When synchronization can ensure the safety of mutable data.
- When a process can map a file and have a thread in this process get a record's lock. Once the lock is acquired, any other thread in any process mapping the file that tries to acquire the lock is blocked until the lock is released.
- Even when accessing a single primitive variable, such as an integer. On machines where the integer is not aligned to the bus data width or is larger than the data width, a single memory load can use more than one memory cycle. While this cannot happen on the SPARC[®] Platform Edition architecture, portable programs cannot rely on this.

Note - On 32-bit architectures a `long long` is not atomic¹ and is read and written as two 32-bit quantities. The types `int`, `char`, `float`, and pointers are atomic on SPARC Platform Edition machines and Intel Architecture machines.

Mutual Exclusion Lock Attributes

Use mutual exclusion locks (mutexes) to serialize thread execution. Mutual exclusion locks synchronize threads, usually by ensuring that only one thread at a time executes a critical section of code. Mutex locks can also preserve single-threaded code.

To change the default mutex attributes, you can declare and initialize an attribute object. Often, the mutex attributes are set in one place at the beginning of the application so they can be located quickly and modified easily. Table 4–1 lists the functions discussed in this section that manipulate mutex attributes.

1. An *atomic* operation cannot be divided into smaller operations.

TABLE 4-1 Mutex Attributes Routines

Operation	Destination Discussion
Initialize a mutex mtribute object	“pthread_mutexattr_init(3THR)” on page 78
Destroy a mutex attribute object	“pthread_mutexattr_destroy(3THR)” on page 79
Set the scope of a mutex	“pthread_mutexattr_setpshared(3THR)” on page 80
Get the scope of a mutex	“pthread_mutexattr_getpshared(3THR)” on page 81
Set the mutex type attribute	“pthread_mutexattr_settype(3THR)” on page 81
Get the mutex type attribute	“pthread_mutexattr_gettype(3THR)” on page 83
Set mutex attribute’s protocol	“pthread_mutexattr_setprotocol(3T)” on page 83
Get mutex attribute’s protocol	“pthread_mutexattr_getprotocol(3T)” on page 86
Set mutex attribute’s priority ceiling	“pthread_mutexattr_setprioceiling(3T)” on page 87
Get mutex attribute’s priority ceiling	“pthread_mutexattr_getprioceiling(3T)” on page 88
Set mutex’s priority ceiling	“pthread_mutex_setprioceiling(3T)” on page 89
Get mutex’s priority ceiling	“pthread_mutex_getprioceiling(3T)” on page 90
Set mutex’s robust attribute	“pthread_mutexattr_setrobust_np(3T)” on page 91
Get mutex’s robust attribute	“pthread_mutexattr_getrobust_np(3T)” on page 92

The differences between Solaris threads and POSIX threads, when defining the scope of a mutex, are shown in Table 4-2.

TABLE 4-2 Mutex Scope Comparison

Solaris	POSIX	Definition
USYNC_PROCESS	PTHREAD_PROCESS_SHARED	Use to synchronize threads in this and other processes
USYNC_PROCESS_ROBUST	No POSIX equivalent	Use to <i>robustly</i> synchronize threads between processes
USYNC_THREAD	PTHREAD_PROCESS_PRIVATE	Use to synchronize threads in this process only

Initialize a Mutex Attribute Object

pthread_mutexattr_init(3THR)

Use `pthread_mutexattr_init(3THR)` to initialize attributes associated with this object to their default values. Storage for each attribute object is allocated by the threads system during execution.

The default value of the *pshared* attribute when this function is called is `PTHREAD_PROCESS_PRIVATE`, which means that the initialized mutex can be used within a process.

Prototype:

```
int pthread_mutexattr_init(pthread_mutexattr_t *mattr);
```

```
#include <pthread.h>
```

```
pthread_mutexattr_t mattr;  
int ret;
```

```
/* initialize an attribute to default value */  
ret = pthread_mutexattr_init(&mattr);
```

mattr is an opaque type that contains a system-allocated attribute object. The possible values of *mattr*'s scope are `PTHREAD_PROCESS_PRIVATE` (the default) and `PTHREAD_PROCESS_SHARED`.

Before a mutex attribute object can be reinitialized, it must first be destroyed by a call to `pthread_mutexattr_destroy(3T)`. The `pthread_mutexattr_init()` call results in the allocation of an opaque object. If the object is not destroyed, a memory leak will result.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If either of the following conditions occurs, the function fails and returns the corresponding value.

ENOMEM

There is not enough memory to initialize the mutex attributes object.

Destroy a Mutex Attribute Object

`pthread_mutexattr_destroy(3THR)`

`pthread_mutexattr_destroy(3THR)` deallocates the storage space used to maintain the attribute object created by `pthread_mutexattr_init()`.

Prototype:

```
int pthread_mutexattr_destroy(pthread_mutexattr_t *mattr)
```

```
#include <pthread.h>
```

```
pthread_mutexattr_t mattr;  
int ret;
```

```
/* destroy an attribute */  
ret = pthread_mutexattr_destroy(&mattr);
```

Return Values

`pthread_mutexattr_destroy()` returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value specified by *mattr* is invalid.

Set the Scope of a Mutex

`pthread_mutexattr_setpshared(3THR)`

`pthread_mutexattr_setpshared(3THR)` sets the scope of the mutex variable.

The scope of a mutex variable can be either process private (intraprocess) or system wide (interprocess). If the mutex is created with the `pshared` attribute set to the `PTHREAD_PROCESS_SHARED` state, and it exists in shared memory, it can be shared among threads from more than one process. This is equivalent to the `USYNC_PROCESS` flag in `mutex_init()` in the original Solaris threads.

Prototype:

```
int pthread_mutexattr_setpshared(pthread_mutexattr_t *mattr,
    int pshared);
```

```
#include <pthread.h>
```

```
pthread_mutexattr_t mattr;
int ret;
```

```
ret = pthread_mutexattr_init(&mattr);
/*
 * resetting to its default value: private
 */
ret = pthread_mutexattr_setpshared(&mattr,
    PTHREAD_PROCESS_PRIVATE);
```

If the mutex `pshared` attribute is set to `PTHREAD_PROCESS_PRIVATE`, only those threads created by the same process can operate on the mutex.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value specified by *mattr* is invalid.

Get the Scope of a Mutex

`pthread_mutexattr_getpshared(3THR)`

`pthread_mutexattr_getpshared(3THR)` returns the scope of the mutex variable defined by `pthread_mutexattr_setpshared()`.

Prototype:

```
int pthread_mutexattr_getpshared(pthread_mutexattr_t *mattr,
                                int *pshared);
```

```
#include <pthread.h>
```

```
pthread_mutexattr_t mattr;
int pshared, ret;
```

```
/* get pshared of mutex */
ret = pthread_mutexattr_getpshared(&mattr, &pshared);
```

Get the current value of *pshared* for the attribute object *mattr*. It is either `PTHREAD_PROCESS_SHARED` or `PTHREAD_PROCESS_PRIVATE`.

Return Values

Returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value specified by *mattr* is invalid.

Set the Mutex Type Attribute

`pthread_mutexattr_settype(3THR)`

```
#include <pthread.h>
```

```
int pthread_mutexattr_settype(pthread_mutexattr_t *attr, int type);
```

`pthread_mutexattr_settype(3THR)` sets the mutex *type* attribute. The default value of the *type* attribute is `PTHREAD_MUTEX_DEFAULT`.

The *type* argument specifies the type of mutex. Valid mutex types include:

`PTHREAD_MUTEX_NORMAL`

This type of mutex does not detect deadlock. A thread attempting to relock this mutex without first unlocking it will deadlock. Attempting to unlock a mutex locked by a different thread results in undefined behavior. Attempting to unlock an unlocked mutex results in undefined behavior.

PTHREAD_MUTEX_ERRORCHECK

This type of mutex provides error checking. A thread attempting to relock this mutex without first unlocking it will return with an error. A thread attempting to unlock a mutex which another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error.

PTHREAD_MUTEX_RECURSIVE

A thread attempting to relock this mutex without first unlocking it will succeed in locking the mutex. The relocking deadlock which can occur with mutexes of type PTHREAD_MUTEX_NORMAL cannot occur with this type of mutex. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread attempting to unlock a mutex which another thread has locked will return with an error. A thread attempting to unlock an unlocked mutex will return with an error. This type of mutex is only supported for mutexes whose process shared attribute is PTHREAD_PROCESS_PRIVATE.

PTHREAD_MUTEX_DEFAULT

Attempting to recursively lock a mutex of this type results in undefined behaviour. Attempting to unlock a mutex of this type which was not locked by the calling thread results in undefined behaviour. Attempting to unlock a mutex of this type which is not locked results in undefined behaviour. An implementation is allowed to map this mutex to one of the other mutex types. (For Solaris threads, PTHREAD_PROCESS_DEFAULT is mapped to PTHREAD_PROCESS_NORMAL.)

Return Values

If successful, the `pthread_mutexattr_settype` function returns zero. Otherwise, an error number is returned to indicate the error.

EINVAL

The value *type* is invalid.

EINVAL

The value specified by *attr* is invalid.

Get the Mutex Type Attribute

`pthread_mutexattr_gettype(3THR)`

```
#include <pthread.h>
```

```
int pthread_mutexattr_gettype(pthread_mutexattr_t *attr, int *type);
```

`pthread_mutexattr_gettype(3THR)` gets the mutex *type* attribute set by `pthread_mutexattr_settype( )`. The default value of the type attribute is `PTHREAD_MUTEX_DEFAULT`.

The *type* argument specifies the type of mutex. Valid mutex types include:

- `PTHREAD_MUTEX_NORMAL`
- `PTHREAD_MUTEX_ERRORCHECK`
- `PTHREAD_MUTEX_RECURSIVE`
- `PTHREAD_MUTEX_DEFAULT`

For a description of each type, see “`pthread_mutexattr_settype(3THR)`” on page 81.

Set Mutex Attribute’s Protocol

`pthread_mutexattr_setprotocol(3T)`

`pthread_mutexattr_setprotocol(3T)` sets the protocol attribute of a mutex attribute object.

```
#include <pthread.h>
```

```
int pthread_mutexattr_setprotocol(pthread_mutexattr_t *attr, int protocol);
```

attr points to a mutex attribute object created by an earlier call to `pthread_mutexattr_init( )`.

protocol defines the protocol applied to the mutex attribute object.

The value of *protocol*, defined in `pthread.h`, can be: `PTHREAD_PRIO_NONE`, `PTHREAD_PRIO_INHERIT`, or `PTHREAD_PRIO_PROTECT`.

- `PTHREAD_PRIO_NONE`

A thread’s priority and scheduling are not affected by the mutex ownership.

- `PTHREAD_PRIO_INHERIT`

This protocol value affects a thread's (such as `thrd1`) priority and scheduling when higher-priority threads block on one or more mutexes owned by `thrd1` where those mutexes are initialized with `PTHREAD_PRIO_INHERIT`. `thrd1` runs with the higher of its priority or the highest priority of any thread waiting on any of the mutexes owned by `thrd1`.

If `thrd1` blocks on a mutex owned by another thread, `thrd3`, the same priority inheritance effect recursively propagates to `thrd3`.

Use `PTHREAD_PRIO_INHERIT` to avoid priority inversion. Priority inversion occurs when a low-priority thread holds a lock that a higher-priority thread wants. Because the higher-priority thread cannot continue until the lower-priority thread releases the lock, each thread is treated as if it had the inverse of its intended priority.

If the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined, for a mutex initialized with the protocol attribute value `PTHREAD_PRIO_INHERIT`, the following actions occur in the Solaris Operating Environment when the owner of that mutex dies:

Note - The behavior on owner death depends on the value of the *robustness* argument of `pthread_mutexattr_setrobust_np()`.

- The mutex is unlocked.
- The next owner of the mutex acquires it with an error return of `EOWNERDEAD`.
- The next owner of the mutex should try to make the state protected by the mutex consistent—the state might have been left inconsistent when the previous owner died. If the owner is successful in making the state consistent, call `pthread_mutex_init()` for the mutex and unlock the mutex.

Note - If `pthread_mutex_init()` is called on a previously initialized, but not yet destroyed mutex, the mutex is not reinitialized.

- If the owner is unable to make the state consistent, do *not* call `pthread_mutex_init()`, but unlock the mutex. In this event, all waiters will be woken up and all subsequent calls to `pthread_mutex_lock()` will fail to acquire the mutex and return an error code of `ENOTRECOVERABLE`. You can now make the mutex state consistent by calling `pthread_mutex_destroy()` to uninitialize the mutex and calling `pthread_mutex_init()` to reinitialize it.
 - If the thread that acquired the lock with `EOWNERDEAD` dies, the next owner acquires the lock with an error code of `EOWNERDEAD`.
-
- **PTHREAD_PRIO_PROTECT**

This protocol value affects a thread's (such as `thrd2`) priority and scheduling when the thread owns one or more mutexes initialized with `PTHREAD_PRIO_PROTECT`. `thrd2` runs with the higher of its priority or the

highest-priority ceiling of all mutexes it owns. Higher-priority threads blocked on any of the mutexes, owned by `thrd2`, have no effect on the scheduling of `thrd2`.

When a thread owns a mutex that is initialized with `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT`, and that thread's original priority changes, such as by a call to `sched_setparam()`, the scheduler does not move the thread to the tail of the scheduling queue at its new priority. Similarly, when a thread unlocks a mutex that is initialized with `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT`, and that thread's original priority changes, the scheduler does not move the thread to the tail of the scheduling queue at its new priority.

If a thread simultaneously owns several mutexes initialized with a mix of `PTHREAD_PRIO_INHERIT` and `PTHREAD_PRIO_PROTECT`, it executes at the highest priority obtained by either of these protocols.

Return Values

On successful completion, `pthread_mutexattr_setprotocol()` returns 0. Any other returned value indicates that an error occurred.

If either of the following conditions occurs, `pthread_mutexattr_setprotocol()` fails and returns the corresponding value.

ENOSYS

Neither of the options `_POSIX_THREAD_PRIO_INHERIT` and `_POSIX_THREAD_PRIO_PROTECT` is defined and the implementation does not support the function.

ENOTSUP

The value specified by *protocol* is an unsupported value.

If either of the following conditions occurs, `pthread_mutexattr_setprotocol()` might fail and return the corresponding value.

EINVAL

The value specified by *attr* or *protocol* is not valid.

EPERM

The caller does not have the privilege to perform the operation.

Get Mutex Attribute's Protocol

`pthread_mutexattr_getprotocol(3T)`

`pthread_mutexattr_getprotocol(3T)` gets the protocol attribute of a mutex attribute object.

```
#include <pthread.h>
```

```
int pthread_mutexattr_getprotocol(const pthread_mutexattr_t *attr, int *protocol);
```

attr points to a mutex attribute object created by an earlier call to `pthread_mutexattr_init( )`.

protocol contains the protocol attribute: `PTHREAD_PRIO_NONE`, `PTHREAD_PRIO_INHERIT`, or `PTHREAD_PRIO_PROTECT`.

Return Values

On successful completion, `pthread_mutexattr_getprotocol( )` returns 0. Any other returned value indicates that an error occurred.

If the following condition occurs, `pthread_mutexattr_getprotocol( )` fails and returns the corresponding value.

ENOSYS

Neither of the options, `_POSIX_THREAD_PRIO_INHERIT` nor `_POSIX_THREAD_PRIO_PROTECT` is defined and the implementation does not support the function.

If either of the following conditions occurs, `pthread_mutexattr_getprotocol( )` might fail and return the corresponding value.

EINVAL

The value specified by *attr* is invalid.

EPERM

The caller does not have the privilege to perform the operation.

Set Mutex Attribute's Priority Ceiling

pthread_mutexattr_setprioceiling(3T)

pthread_mutexattr_setprioceiling(3T) sets the priority ceiling attribute of a mutex attribute object.

```
#include <pthread.h>
```

```
int pthread_mutexattr_setprioceiling(pthread_mutexattr_t *attr, int prioceiling, int *oldceiling);
```

attr points to a mutex attribute object created by an earlier call to *pthread_mutexattr_init()*.

Note - The *attr* mutex attribute object includes the priority ceiling attribute only if the symbol `_POSIX_THREAD_PRIO_PROTECT` is defined.

prioceiling specifies the priority ceiling of initialized mutexes. The ceiling defines the minimum priority level at which the critical section guarded by the mutex is executed. *prioceiling* will be within the maximum range of priorities defined by `SCHED_FIFO`. To avoid priority inversion, *prioceiling* will be set to a priority higher than or equal to the highest priority of all the threads that might lock the particular mutex.

oldceiling contains the old priority ceiling value.

Return Values

On successful completion, *pthread_mutexattr_setprioceiling()* returns 0. Any other returned value indicates that an error occurred.

If any of the following conditions occurs, *pthread_mutexattr_setprioceiling()* fails and returns the corresponding value.

ENOSYS

The option `_POSIX_THREAD_PRIO_PROTECT` is not defined and the implementation does not support the function.

If either of the following conditions occurs, *pthread_mutexattr_setprioceiling()* might fail and return the corresponding value.

EINVAL

The value specified by *attr* or *prioceiling* is invalid.

EPERM

The caller does not have the privilege to perform the operation.

Get Mutex Attribute's Priority Ceiling

`pthread_mutexattr_getprioceiling(3T)`

pthread_mutexattr_getprioceiling(3T) gets the priority ceiling attribute of a mutex attribute object.

```
#include <pthread.h>
```

```
int pthread_mutexattr_getprioceiling(const pthread_mutexattr_t *attr, int *prioceiling);
```

attr designates the attribute object created by an earlier call to `pthread_mutexattr_init()`.

Note - The *attr* mutex attribute object includes the priority ceiling attribute only if the symbol `_POSIX_THREAD_PRIO_PROTECT` is defined.

`pthread_mutexattr_getprioceiling()` returns the priority ceiling of initialized mutexes in *prioceiling*. The ceiling defines the minimum priority level at which the critical section guarded by the mutex is executed. *prioceiling* will be within the maximum range of priorities defined by `SCHED_FIFO`. To avoid priority inversion, *prioceiling* will be set to a priority higher than or equal to the highest priority of all the threads that might lock the particular mutex.

Return Values

On successful completion, `pthread_mutexattr_getprioceiling()` returns 0. Any other returned value indicates that an error occurred.

If any of the following conditions occurs, `pthread_mutexattr_getprioceiling()` fails and returns the corresponding value.

ENOSYS

The option `_POSIX_THREAD_PRIO_PROTECT` is not defined and the implementation does not support the function.

If either of the following conditions occurs, `pthread_mutexattr_getprioceiling()` might fail and return the corresponding value.

EINVAL

The value specified by *attr* is invalid.

EPERM

The caller does not have the privilege to perform the operation.

Set Mutex's Priority Ceiling

pthread_mutex_setprioceiling(3T)

pthread_mutex_setprioceiling(3T) sets the priority ceiling of a mutex.

```
#include <pthread.h>
```

```
int pthread_mutex_setprioceiling(pthread_mutexatt_t *mutex, int prioceiling, int *old_ceiling);
```

pthread_mutex_setprioceiling() changes the priority ceiling, *prioceiling*, of a mutex, *mutex*. *pthread_mutex_setprioceiling()* locks the mutex if it is unlocked, or blocks until it can successfully lock the mutex, changes the priority ceiling of the mutex and releases the mutex. The process of locking the mutex need not adhere to the priority protect protocol.

Note - The mutex attribute object, *mutex*, includes the priority ceiling attribute only if the symbol `_POSIX_THREAD_PRIO_PROTECT` is defined.

If *pthread_mutex_setprioceiling()* succeeds, the previous value of the priority ceiling is returned in *old_ceiling*. If *pthread_mutex_setprioceiling()* fails, the mutex priority ceiling remains unchanged.

Return Values

On successful completion, *pthread_mutex_setprioceiling()* returns 0. Any other returned value indicates that an error occurred.

If the following condition occurs, *pthread_mutexatt_setprioceiling()* fails and returns the corresponding value.

ENOSYS

The option `_POSIX_THREAD_PRIO_PROTECT` is not defined and the implementation does not support the function.

If any of the following conditions occurs, *pthread_mutex_setprioceiling()* might fail and return the corresponding value.

EINVAL

The priority requested by *prioceiling* is out of range.

EINVAL

The value specified by *mutex* does not refer to a currently existing mutex.

ENOSYS

The implementation does not support the priority ceiling protocol for mutexes.

EPERM

The caller does not have the privilege to perform the operation.

Get Mutex's Priority Ceiling

`pthread_mutex_getprioceiling(3T)`

pthread_mutex_getprioceiling(3T) gets the priority ceiling of a mutex.

```
#include <pthread.h>
```

```
int pthread_mutex_getprioceiling(const pthread_mutexatt_t *mutex, int *prioceiling);
```

`pthread_mutex_getprioceiling( )` returns the priority ceiling, *prioceiling* of a mutex *mutex*.

Return Values

On successful completion, `pthread_mutex_getprioceiling( )` returns 0. Any other returned value indicates that an error occurred.

If any of the following conditions occurs, `pthread_mutexatt_getprioceiling( )` fails and returns the corresponding value.

ENOSYS

The option `_POSIX_THREAD_PRIO_PROTECT` is not defined and the implementation does not support the function.

If any of the following conditions occurs, `pthread_mutex_getprioceiling( )` might fail and return the corresponding value.

EINVAL

The value specified by *mutex* does not refer to a currently existing mutex.

ENOSYS

The implementation does not support the priority ceiling protocol for mutexes.

EPERM

The caller does not have the privilege to perform the operation.

Set Mutex's Robust Attribute

`pthread_mutexattr_setrobust_np(3T)`

`pthread_mutexattr_setrobust_np(3T)` sets the robust attribute of a mutex attribute object.

```
#include <pthread.h>
```

```
int pthread_mutexattr_setrobust_np(pthread_mutexattr_t *attr, int *robustness);
```

Note - `pthread_mutexattr_setrobust_np( )` applies only if the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined.

attr points to the mutex attribute object previously created by a call to `pthread_mutexattr_init( )`.

robustness defines the behavior when the owner of the mutex dies. The value of *robustness*, defined in `pthread.h`, is `PTHREAD_MUTEX_ROBUST_NP` or `PTHREAD_MUTEX_STALLED_NP`. The default value is `PTHREAD_MUTEX_STALLED_NP`.

■ **PTHREAD_MUTEX_ROBUST_NP**

When the owner of the mutex dies, all subsequent calls to `pthread_mutex_lock( )` are blocked from progress in an unspecified manner.

■ **PTHREAD_MUTEX_STALLED_NP**

When the owner of the mutex dies, the mutex is unlocked. The next owner of this mutex acquires it with an error return of `EOWNERDEAD`.

Note - Your application must check the return code from `pthread_mutex_lock( )` for a mutex of this type.

- The new owner of this mutex should make the state protected by the mutex consistent; this state might have been left inconsistent when the previous owner died.
- If the new owner is able to make the state consistent, call `pthread_mutex_consistent_np( )` for the mutex, and unlock the mutex.

- If the new owner is *not* able to make the state consistent, do *not* call `pthread_mutex_consistent_np()` for the mutex, but unlock the mutex.

All waiters are woken up and all subsequent calls to `pthread_mutex_lock()` fail to acquire the mutex. The return code is `ENOTRECOVERABLE`. The mutex can be made consistent by calling `pthread_mutex_destroy()` to uninitialize the mutex, and calling `pthread_mutex_init()` to reinitialize it.

If the thread that acquire the lock with `EOWNERDEAD` died, the next owner acquires the lock with an `EOWNERDEAD` return code.

Return Values

On successful completion, `pthread_mutexattr_setrobust_np()` returns 0. Any other returned value indicates that an error occurred.

If any of the following conditions occurs, `pthread_mutexattr_setrobust_np()` fails and returns the corresponding value.

ENOSYS

The option `_POSIX_THREAD_PRIO__INHERIT` is not defined or the implementation does not support `pthread_mutexattr_setrobust_np()`.

ENOTSUP

The value specified by *robustness* is not supported.

`pthread_mutexattr_setrobust_np()` might fail if:

EINVAL

The value specified by *attr* or *robustness* is invalid.

Get Mutex's Robust Attribute

`pthread_mutexattr_getrobust_np(3T)`

`pthread_mutexattr_getrobust_np(3T)` gets the robust attribute of a mutex attribute object.

```
#include <pthread.h>
```

```
int pthread_mutexattr_getrobust_np(pthread_mutexattr_t *attr, int *robustness);
```

Note - `pthread_mutexattr_getrobust_np()` applies only if the symbol `_POSIX_THREAD_PRIORITY_INHERIT` is defined.

attr points to the mutex attribute object previously created by a call to `pthread_mutexattr_init()`.

robustness is the value of the robust attribute of a mutex attribute object.

Return Values

On successful completion, `pthread_mutexattr_getrobust_np()` returns 0. Any other returned value indicates that an error occurred.

If any of the following conditions occurs, `pthread_mutexattr_getrobust_np()` fails and returns the corresponding value.

ENOSYS

The option `_POSIX_THREAD_PRIORITY_INHERIT` is not defined or the implementation does not support `pthread_mutexattr_getrobust_np()`.

ENOTSUP

The value specified by *robustness* is not supported.

`pthread_mutexattr_getrobust_np()` might fail if:

EINVAL

The value specified by *attr* or *robustness* is invalid.

Using Mutual Exclusion Locks

Table 4–3 lists the functions discussed in this chapter that manipulate mutex locks.

TABLE 4–3 Routines for Mutual Exclusion Locks

Operation	Destination Discussion
Initialize a mutex	“ <code>pthread_mutex_init(3THR)</code> ” on page 94
Make mutex consistent	“ <code>pthread_mutex_consistent_np(3T)</code> ” on page 95

TABLE 4-3 Routines for Mutual Exclusion Locks *(continued)*

Operation	Destination Discussion
Lock a mutex	“pthread_mutex_lock(3THR)” on page 97
Unlock a mutex	“pthread_mutex_unlock(3THR)” on page 98
Lock with a nonblocking mutex	“pthread_mutex_trylock(3THR)” on page 99
Destroy a mutex	“pthread_mutex_destroy(3THR)” on page 101

The default scheduling policy, `SCHED_OTHER`, does not specify the order in which threads can acquire a lock. When multiple threads are waiting for a mutex, the order of acquisition is undefined. When there is contention, the default behavior is to unblock threads in priority order.

Initialize a Mutex

pthread_mutex_init(3THR)

Use `pthread_mutex_init(3THR)` to initialize the mutex pointed at by `mp` to its default value (`mattr` is `NULL`), or to specify mutex attributes that have already been set with `pthread_mutexattr_init( )`. (For Solaris threads, see “mutex_init(3THR)” on page 213.)

Prototype:

```
int pthread_mutex_init(pthread_mutex_t *mp,
    const pthread_mutexattr_t *mattr);
```

```
#include <pthread.h>
```

```
pthread_mutex_t mp = PTHREAD_MUTEX_INITIALIZER;
pthread_mutexattr_t mattr;
int ret;
```

```
/* initialize a mutex to its default value */
ret = pthread_mutex_init(&mp, NULL);
```

```
/* initialize a mutex */
ret = pthread_mutex_init(&mp, &mattr);
```

When the mutex is initialized, it is in an unlocked state. The mutex can be in memory shared between processes or in memory private to a process.

Note - The mutex memory must be cleared to zero before initialization.

The effect of *mattr* being NULL is the same as passing the address of a default mutex attribute object, but without the memory overhead.

Statically defined mutexes can be initialized directly to have default attributes with the macro `PTHREAD_MUTEX_INITIALIZER`.

A mutex lock must not be reinitialized or destroyed while other threads might be using it. Program failure will result if either action is not done correctly. If a mutex is reinitialized or destroyed, the application must be sure the mutex is not currently in use.

Return Values

`pthread_mutex_init()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EBUSY

The implementation has detected an attempt to reinitialize the object referenced by *mp* (a previously initialized, but not yet destroyed mutex).

EINVAL

The *mattr* attribute value is invalid. The mutex has not been modified.

EFAULT

The address for the mutex pointed at by *mp* is invalid.

Make Mutex Consistent

`pthread_mutex_consistent_np(3T)`

```
#include <pthread.h>
int pthread_mutex_consistent_np(pthread_mutex_t *mutex);
```

Note - `pthread_mutex_consistent_np( )` applies only if the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined and for mutexes that are initialized with the protocol attribute value `PTHREAD_PRIO_INHERIT`.

If the owner of a mutex dies, the mutex can become inconsistent.

`pthread_mutex_consistent_np` makes the mutex object, *mutex*, consistent after the death of its owner.

Call `pthread_mutex_lock( )` to acquire the inconsistent mutex. The `EOWNERDEAD` return value indicates an inconsistent mutex.

Call `pthread_mutex_consistent_np( )` while holding the mutex acquired by a previous call to `pthread_mutex_lock( )`.

Because the critical section protected by the mutex might have been left in an inconsistent state by the dead owner, make the mutex consistent only if you are able to make the critical section protected by the mutex consistent.

Calls to `pthread_mutex_lock( )`, `pthread_mutex_unlock( )`, and `pthread_mutex_trylock( )` for a consistent mutex behave in the normal manner.

The behavior of `pthread_mutex_consistent_np( )` for a mutex that is *not* inconsistent, or that is not held, is undefined.

Return Values

`pthread_mutex_consistent_np( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

`pthread_mutex_consistent_np( )` fails if:

ENOSYS

The option `_POSIX_THREAD_PRIO_INHERIT` is not defined or the implementation does not support `pthread_mutex_consistent_np( )`.

`pthread_mutex_consistent_np( )` might fail if:

EINVAL

The value specified by *mutex* is invalid.

Lock a Mutex

pthread_mutex_lock(3THR)

Prototype:

```
int pthread_mutex_lock(pthread_mutex_t *mutex);
```

```
#include <pthread.h>
```

```
pthread_mutex_t mutex;  
int ret;
```

```
ret = pthread_mutex_lock(&mp); /* acquire the mutex */
```

Use `pthread_mutex_lock(3THR)` to lock the mutex pointed to by *mutex*. When `pthread_mutex_lock( )` returns, the mutex is locked and the calling thread is the owner. If the mutex is already locked and owned by another thread, the calling thread blocks until the mutex becomes available. (For Solaris threads, see “`mutex_lock(3THR)`” on page 215.)

If the mutex type is `PTHREAD_MUTEX_NORMAL`, deadlock detection is not provided. Attempting to relock the mutex causes deadlock. If a thread attempts to unlock a mutex that it has not locked or a mutex that is unlocked, undefined behaviour results.

If the mutex type is `PTHREAD_MUTEX_ERRORCHECK`, then error checking is provided. If a thread attempts to relock a mutex that it has already locked, an error will be returned. If a thread attempts to unlock a mutex that it has not locked or a mutex that is unlocked, an error will be returned.

If the mutex type is `PTHREAD_MUTEX_RECURSIVE`, then the mutex maintains the concept of a lock count. When a thread successfully acquires a mutex for the first time, the lock count is set to one. Every time a thread relocks this mutex, the lock count is incremented by one. Each time the thread unlocks the mutex, the lock count is decremented by one. When the lock count reaches zero, the mutex becomes available for other threads to acquire. If a thread attempts to unlock a mutex that it has not locked or a mutex which is unlocked, an error will be returned.

If the mutex type is `PTHREAD_MUTEX_DEFAULT`, attempting to recursively lock the mutex results in undefined behavior. Attempting to unlock the mutex if it was not locked by the calling thread results in undefined behavior. Attempting to unlock the mutex if it is not locked results in undefined behavior.

Return Values

`pthread_mutex_lock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EAGAIN

The mutex could not be acquired because the maximum number of recursive locks for mutex has been exceeded.

EDEADLK

The current thread already owns the mutex.

If the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined, the mutex is initialized with the protocol attribute value `PTHREAD_PRIO_INHERIT`, and the *robustness* argument of `pthread_mutexattr_setrobust_np()` is `PTHREAD_MUTEX_ROBUST_NP` the function fails and returns:

EOWNERDEAD

The last owner of this mutex died while holding the mutex. This mutex is now owned by the caller. The caller must attempt to make the state protected by the mutex consistent.

If the caller is able to make the state consistent, call `pthread_mutex_consistent_np()` for the mutex and unlock the mutex. Subsequent calls to `pthread_mutex_lock()` will behave normally.

If the caller is unable to make the state consistent, do not call `pthread_mutex_init()` for the mutex, but unlock the mutex. Subsequent calls to `pthread_mutex_lock()` fail to acquire the mutex and return an `ENOTRECOVERABLE` error code.

If the owner that acquired the lock with `EOWNERDEAD` dies, the next owner acquires the lock with `EOWNERDEAD`.

ENOTRECOVERABLE

The mutex you are trying to acquire is protecting state left irrecoverable by the mutex's previous owner that died while holding the lock. The mutex has not been acquired. This condition can occur when the lock was previously acquired with `EOWNERDEAD` and the owner was unable to cleanup the state and had unlocked the mutex without making the mutex state consistent.

ENOMEM

The limit on the number of simultaneously held mutexes has been exceeded.

Unlock a Mutex

`pthread_mutex_unlock(3THR)`

Use `pthread_mutex_unlock(3THR)` to unlock the mutex pointed to by *mutex*. (For Solaris threads, see “`mutex_unlock(3THR)`” on page 215.)

```

Prototype:
int pthread_mutex_unlock(pthread_mutex_t *mutex);

#include <pthread.h>

pthread_mutex_t mutex;
int ret;

ret = pthread_mutex_unlock(&mutex); /* release the mutex */

```

`pthread_mutex_unlock( )` releases the mutex object referenced by *mutex*. The manner in which a mutex is released is dependent upon the mutex's type attribute. If there are threads blocked on the mutex object referenced by *mutex* when `pthread_mutex_unlock( )` is called, resulting in the mutex becoming available, the scheduling policy is used to determine which thread shall acquire the mutex. (In the case of `PTHREAD_MUTEX_RECURSIVE` mutexes, the mutex becomes available when the count reaches zero and the calling thread no longer has any locks on this mutex).

Return Values

`pthread_mutex_unlock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EPERM

The current thread does not own the mutex.

Lock With a Nonblocking Mutex

`pthread_mutex_trylock(3THR)`

Use `pthread_mutex_trylock(3THR)` to attempt to lock the mutex pointed to by *mutex*. (For Solaris threads, see “`mutex_trylock(3THR)`” on page 216.)

```

Prototype:
int pthread_mutex_trylock(pthread_mutex_t *mutex);

#include <pthread.h>

pthread_mutex_t mutex;
int ret;

ret = pthread_mutex_trylock(&mutex); /* try to lock the mutex */

```

`pthread_mutex_trylock( )` is a nonblocking version of `pthread_mutex_lock( )`. If the mutex object referenced by *mutex* is currently locked (by any thread, including the current thread), the call returns immediately. Otherwise, the mutex is locked and the calling thread is the owner.

Return Values

`pthread_mutex_trylock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EBUSY

The mutex could not be acquired because the mutex pointed to by *mutex* was already locked.

EAGAIN

The mutex could not be acquired because the maximum number of recursive locks for *mutex* has been exceeded.

If the symbol `_POSIX_THREAD_PRIO_INHERIT` is defined, the mutex is initialized with the protocol attribute value `PTHREAD_PRIO_INHERIT`, and the *robustness* argument of `pthread_mutexattr_setrobust_np( )` is `PTHREAD_MUTEX_ROBUST_NP` the function fails and returns:

EOWNERDEAD

The last owner of this mutex died while holding the mutex. This mutex is now owned by the caller. The caller must attempt to make the state protected by the mutex consistent.

If the caller is able to make the state consistent, call `pthread_mutex_consistent_np( )` for the mutex and unlock the mutex. Subsequent calls to `pthread_mutex_lock( )` will behave normally.

If the caller is unable to make the state consistent, do not call `pthread_mutex_init( )` for the mutex, but unlock the mutex. Subsequent calls to `pthread_mutex_trylock( )` fail to acquire the mutex and return an `ENOTRECOVERABLE` error code.

If the owner that acquired the lock with `EOWNERDEAD` dies, the next owner acquires the lock with `EOWNERDEAD`.

ENOTRECOVERABLE

The mutex you are trying to acquire is protecting state left irrecoverable by the mutex's previous owner that died while holding the lock. The mutex has not been acquired. This condition can occur when the lock was previously acquired with

EOWNERDEAD and the owner was unable to cleanup the state and had unlocked the mutex without making the mutex state consistent.

ENOMEM

The limit on the number of simultaneously held mutexes has been exceeded.

Destroy a Mutex

pthread_mutex_destroy(3THR)

Use `pthread_mutex_destroy(3THR)` to destroy any state associated with the mutex pointed to by *mp*. (For Solaris threads, see “`mutex_destroy(3THR)`” on page 215.)

Prototype:

```
int pthread_mutex_destroy(pthread_mutex_t *mp);
```

```
#include <pthread.h>
```

```
pthread_mutex_t mp;  
int ret;
```

```
ret = pthread_mutex_destroy(&mp); /* mutex is destroyed */
```

Note that the space for storing the mutex is not freed.

Return Values

`pthread_mutex_destroy( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occur, the function fails and returns the corresponding value.

EINVAL

The value specified by *mp* does not refer to an initialized mutex object.

Mutex Lock Code Examples

Code Example 4–1 shows some code fragments with mutex locking.

CODE EXAMPLE 4–1 Mutex Lock Example

```
#include <pthread.h>
```

```

pthread_mutex_t count_mutex;
long long count;

void
increment_count()
{
    pthread_mutex_lock(&count_mutex);
    count = count + 1;
    pthread_mutex_unlock(&count_mutex);
}

long long
get_count()
{
    long long c;

    pthread_mutex_lock(&count_mutex);
    c = count;
    pthread_mutex_unlock(&count_mutex);
    return (c);
}

```

The two functions in Code Example 4–1 use the mutex lock for different purposes. The `increment_count()` function uses the mutex lock simply to ensure an atomic update of the shared variable. The `get_count()` function uses the mutex lock to guarantee that the 64-bit quantity *count* is read atomically. On a 32-bit architecture, a `long long` is really two 32-bit quantities.

Reading an integer value is an atomic operation because integer is the common word size on most machines.

Using Locking Hierarchies

You will occasionally want to access two resources at once. Perhaps you are using one of the resources, and then discover that the other resource is needed as well. There could be a problem if two threads attempt to claim both resources but lock the associated mutexes in different orders. For example, if the two threads lock mutexes 1 and 2 respectively, then a deadlock occurs when each attempts to lock the other mutex. Code Example 4–2 shows possible deadlock scenarios.

Thread 1	Thread 2
<pre>pthread_mutex_lock(&m1); /* use resource 1 */ pthread_mutex_lock(&m2); /* use resources 1 and 2 */ pthread_mutex_unlock(&m2); pthread_mutex_unlock(&m1);</pre>	<pre>pthread_mutex_lock(&m2); /* use resource 2 */ pthread_mutex_lock(&m1); /* use resources 1 and 2 */ pthread_mutex_unlock(&m1); pthread_mutex_unlock(&m2);</pre>

The best way to avoid this problem is to make sure that whenever threads lock multiple mutexes, they do so in the same order. When locks are always taken in a prescribed order, deadlock should not occur. This technique is known as lock hierarchies: order the mutexes by logically assigning numbers to them.

Also, honor the restriction that you cannot take a mutex that is assigned n when you are holding any mutex assigned a number greater than n .

However, this technique cannot always be used—sometimes you must take the mutexes in an order other than prescribed. To prevent deadlock in such a situation, use `pthread_mutex_trylock( )`. One thread must release its mutexes when it discovers that deadlock would otherwise be inevitable.

Thread 1	Thread 2
pthread_mutex_lock(&m1); pthread_mutex_lock(&m2);	for (; ;) { pthread_mutex_lock(&m2); if(pthread_mutex_trylock(&m1)==0) /* got it! */ break; /* didn't get it */ pthread_mutex_unlock(&m2); } /* get locks; no processing */ pthread_mutex_unlock(&m1); pthread_mutex_unlock(&m2);
/* no processing */ pthread_mutex_unlock(&m2); pthread_mutex_unlock(&m1);	

In Code Example 4-3, thread 1 locks mutexes in the prescribed order, but thread 2 takes them out of order. To make certain that there is no deadlock, thread 2 has to take mutex 1 very carefully; if it were to block waiting for the mutex to be released, it is likely to have just entered into a deadlock with thread 1.

To ensure this does not happen, thread 2 calls `pthread_mutex_trylock()`, which takes the mutex if it is available. If it is not, thread 2 returns immediately, reporting failure. At this point, thread 2 must release mutex 2, so that thread 1 can lock it, and then release both mutex 1 and mutex 2.

Nested Locking With a Singly Linked List

Code Example 4-4 and Code Example 4-5 show how to take three locks at once, but prevent deadlock by taking the locks in a prescribed order.

CODE EXAMPLE 4-4 Singly Linked List Structure

```
typedef struct node1 {
    int value;
    struct node1 *link;
    pthread_mutex_t lock;
} node1_t;

node1_t ListHead;
```

This example uses a singly linked list structure with each node containing a mutex. To remove a node from the list, first search the list starting at *ListHead* (which itself is never removed) until the desired node is found.

To protect this search from the effects of concurrent deletions, lock each node before any of its contents are accessed. Because all searches start at *ListHead*, there is never a deadlock because the locks are always taken in list order.

When the desired node is found, lock both the node and its predecessor since the change involves both nodes. Because the predecessor's lock is always taken first, you are again protected from deadlock. Code Example 4–5 shows the C code to remove an item from a singly linked list.

CODE EXAMPLE 4–5 Singly Linked List With Nested Locking

```
node1_t *delete(int value)
{
    node1_t *prev, *current;

    prev = &ListHead;
    pthread_mutex_lock(&prev->lock);

    while ((current = prev->link) != NULL) {
        pthread_mutex_lock(&current->lock);
        if (current->value == value) {
            prev->link = current->link;
            pthread_mutex_unlock(&current->lock);

            pthread_mutex_unlock(&prev->lock);
            current->link = NULL;
            return(current);
        }
        pthread_mutex_unlock(&prev->lock);
        prev = current;
    }
    pthread_mutex_unlock(&prev->lock);
    return(NULL);
}
```

Nested Locking With a Circular Linked List

Code Example 4–6 modifies the previous list structure by converting it into a circular list. There is no longer a distinguished head node; now a thread might be associated with a particular node and might perform operations on that node and its neighbor. Note that lock hierarchies do not work easily here because the obvious hierarchy (following the links) is circular.

CODE EXAMPLE 4–6 Circular Linked List Structure

```
typedef struct node2 {
    int value;
    struct node2 *link;
    pthread_mutex_t lock;
} node2_t;
```

Here is the C code that acquires the locks on two nodes and performs an operation involving both of them.

CODE EXAMPLE 4-7 Circular Linked List With Nested Locking

```
void Hit Neighbor(node2_t *me) {
    while (1) {
        pthread_mutex_lock(&me->lock);
        if (pthread_mutex_lock(&me->link->lock) != 0) {
            /* failed to get lock */
            pthread_mutex_unlock(&me->lock);
            continue;
        }
        break;
    }
    me->link->value += me->value;
    me->value /= 2;
    pthread_mutex_unlock(&me->link->lock);
    pthread_mutex_unlock(&me->lock);
}
```

Condition Variable Attributes

Use condition variables to atomically block threads until a particular condition is true. Always use condition variables together with a mutex lock.

With a condition variable, a thread can atomically block until a condition is satisfied. The condition is tested under the protection of a mutual exclusion lock (mutex).

When the condition is false, a thread usually blocks on a condition variable and atomically releases the mutex waiting for the condition to change. When another thread changes the condition, it can signal the associated condition variable to cause one or more waiting threads to wake up, acquire the mutex again, and reevaluate the condition.

Condition variables can be used to synchronize threads among processes when they are allocated in memory that can be written to and is shared by the cooperating processes.

The scheduling policy determines how blocking threads are awakened. For the default `SCHED_OTHER`, threads are awakened in priority order.

The attributes for condition variables must be set and initialized before the condition variables can be used. The functions that manipulate condition variable attributes are listed in Table 4-4.

TABLE 4-4 Condition Variable Attributes

Operation	Destination Discussion
Initialize a condition variable attribute	“pthread_condattr_init(3THR)” on page 107
Remove a condition variable attribute	“pthread_condattr_destroy(3THR)” on page 108
Set the scope of a condition variable	“pthread_condattr_setpshared(3THR)” on page 109
Get the scope of a condition variable	“pthread_condattr_getpshared(3THR)” on page 110

The differences between Solaris and POSIX threads, when defining the scope of a condition variable, are shown in Table 4-5.

TABLE 4-5 Condition Variable Scope Comparison

Solaris	POSIX	Definition
USYNC_PROCESS	PTHREAD_PROCESS_SHARED	Use to synchronize threads in this and other processes
USYNC_THREAD	PTHREAD_PROCESS_PRIVATE	Use to synchronize threads in this process only

Initialize a Condition Variable Attribute

pthread_condattr_init(3THR)

Use `pthread_condattr_init(3THR)` to initialize attributes associated with this object to their default values. Storage for each attribute object is allocated by the threads system during execution. The default value of the *pshared* attribute when this

function is called is `PTHREAD_PROCESS_PRIVATE`, which means that the initialized condition variable can be used within a process.

Prototype:

```
int pthread_condattr_init(pthread_condattr_t *cattr);
```

```
#include <pthread.h>
pthread_condattr_t cattr;
int ret;
```

```
/* initialize an attribute to default value */
ret = pthread_condattr_init(&cattr);
```

cattr is an opaque data type that contains a system-allocated attribute object. The possible values of *cattr*'s scope are `PTHREAD_PROCESS_PRIVATE` (the default) and `PTHREAD_PROCESS_SHARED`.

Before a condition variable attribute can be reused, it must first be reinitialized by `pthread_condattr_destroy(3T)`. The `pthread_condattr_init( )` call returns a pointer to an opaque object. If the object is not destroyed, a memory leak will result.

Return Values

`pthread_condattr_init( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When either of the following conditions occurs, the function fails and returns the corresponding value.

ENOMEM

There is not enough memory to initialize the thread attributes object.

EINVAL

The value specified by *cattr* is invalid.

Remove a Condition Variable Attribute

`pthread_condattr_destroy(3THR)`

Use `pthread_condattr_destroy(3THR)` to remove storage and render the attribute object invalid.

Prototype:

```
int pthread_condattr_destroy(pthread_condattr_t *cattr);
```

```
#include <pthread.h>
pthread_condattr_t cattr;
```

```

int ret;

/* destroy an attribute */
ret
= pthread_condattr_destroy(&cattr);

```

Return Values

`pthread_condattr_destroy( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value specified by *cattr* is invalid.

Set the Scope of a Condition Variable

pthread_condattr_setpshared(3THR)

`pthread_condattr_setpshared(3THR)` sets the scope of a condition variable to either process private (intraprocess) or system wide (interprocess). If the condition variable is created with the `pshared` attribute set to the `PTHREAD_PROCESS_SHARED` state, and it exists in shared memory, it can be shared among threads from more than one process. This is equivalent to the `USYNC_PROCESS` flag in `mutex_init( )` in the original Solaris threads.

If the mutex `pshared` attribute is set to `PTHREAD_PROCESS_PRIVATE` (default value), only those threads created by the same process can operate on the mutex. Using `PTHREAD_PROCESS_PRIVATE` results in the same behavior as with the `USYNC_THREAD` flag in the original Solaris threads `cond_init( )` call, which is that of a local condition variable. `PTHREAD_PROCESS_SHARED` is equivalent to a global condition variable.

Prototype:

```

int pthread_condattr_setpshared(pthread_condattr_t *cattr,
    int pshared);

```

```

#include <pthread.h>

```

```

pthread_condattr_t cattr;
int ret;

```

```

/* all processes */
ret = pthread_condattr_setpshared(&cattr, PTHREAD_PROCESS_SHARED);

/* within a process */
ret = pthread_condattr_setpshared(&cattr, PTHREAD_PROCESS_PRIVATE);

```

Return Values

`pthread_condattr_setpshared( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. If the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value of *cattr* is invalid, or the *pshared* value is invalid.

Get the Scope of a Condition Variable

`pthread_condattr_getpshared(3THR)`

`pthread_condattr_getpshared(3THR)` gets the current value of *pshared* for the attribute object *cattr*. The value is either `PTHREAD_PROCESS_SHARED` or `PTHREAD_PROCESS_PRIVATE`.

Prototype:

```
int pthread_condattr_getpshared(const pthread_condattr_t *cattr,  
    int *pshared);
```

```
#include <pthread.h>
```

```
pthread_condattr_t cattr;  
int pshared;  
int ret;
```

```
/* get pshared value of condition variable */  
ret = pthread_condattr_getpshared(&cattr, &pshared);
```

Return Values

`pthread_condattr_getpshared( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value of *cattr* is invalid.

Using Condition Variables

This section explains using condition variables. Table 4–6 lists the functions that are available.

TABLE 4-6 Condition Variables Functions

Operation	Destination Discussion
Initialize a condition variable	“pthread_cond_init(3THR)” on page 111
Block on a condition variable	“pthread_cond_wait(3THR)” on page 112
Unblock a specific thread	“pthread_cond_signal(3THR)” on page 114
Block until a specified event	“pthread_cond_timedwait(3THR)” on page 115
Unblock all threads	“pthread_cond_broadcast(3THR)” on page 116
Destroy condition variable state	“pthread_cond_destroy(3THR)” on page 118

Initialize a Condition Variable

pthread_cond_init(3THR)

Use `pthread_cond_init(3THR)` to initialize the condition variable pointed at by `cv` to its default value (`cattr` is `NULL`), or to specify condition variable attributes that are already set with `pthread_condattr_init( )`. The effect of `cattr` being `NULL` is the same as passing the address of a default condition variable attribute object, but without the memory overhead. (For Solaris threads, see “`cond_init(3THR)`” on page 216.)

Prototype:

```
int pthread_cond_init(pthread_cond_t *cv,
    const pthread_condattr_t *cattr);
```

```
#include <pthread.h>
```

```
pthread_cond_t cv;
pthread_condattr_t cattr;
int ret;
```

```
/* initialize a condition variable to its default value */
```

```
ret = pthread_cond_init(&cv, NULL);

/* initialize a condition variable */
ret = pthread_cond_init(&cv, &cattr);
```

Statically defined condition variables can be initialized directly to have default attributes with the macro `PTHREAD_COND_INITIALIZER`. This has the same effect as dynamically allocating `pthread_cond_init()` with null attributes. No error checking is done.

Multiple threads must not simultaneously initialize or reinitialize the same condition variable. If a condition variable is reinitialized or destroyed, the application must be sure the condition variable is not in use.

Return Values

`pthread_cond_init()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

The value specified by *cattr* is invalid.

EBUSY

The condition variable is being used.

EAGAIN

The necessary resources are not available.

ENOMEM

There is not enough memory to initialize the condition variable.

Block on a Condition Variable

`pthread_cond_wait(3THR)`

Use `pthread_cond_wait(3THR)` to atomically release the mutex pointed to by *mp* and to cause the calling thread to block on the condition variable pointed to by *cv*. (For Solaris threads, see “`cond_wait(3THR)`” on page 217.)

Prototype:

```
int pthread_cond_wait(pthread_cond_t *cv, pthread_mutex_t *mutex);
```

```
#include <pthread.h>

pthread_cond_t cv;
pthread_mutex_t mp;
int ret;

/* wait on condition variable */
ret = pthread_cond_wait(&cv, &mp);
```

The blocked thread can be awakened by a `pthread_cond_signal()`, a `pthread_cond_broadcast()`, or when interrupted by delivery of a signal.

Any change in the value of a condition associated with the condition variable cannot be inferred by the return of `pthread_cond_wait()`, and any such condition must be reevaluated.

The `pthread_cond_wait()` routine always returns with the mutex locked and owned by the calling thread, even when returning an error.

This function blocks until the condition is signaled. It atomically releases the associated mutex lock before blocking, and atomically acquires it again before returning.

In typical use, a condition expression is evaluated under the protection of a mutex lock. When the condition expression is false, the thread blocks on the condition variable. The condition variable is then signaled by another thread when it changes the condition value. This causes one or all of the threads waiting on the condition to unblock and to try to acquire the mutex lock again.

Because the condition can change before an awakened thread returns from `pthread_cond_wait()`, the condition that caused the wait must be retested before the mutex lock is acquired. The recommended test method is to write the condition check as a `while()` loop that calls `pthread_cond_wait()`.

```
pthread_mutex_lock();
while(condition_is_false)
    pthread_cond_wait();
pthread_mutex_unlock();
```

No specific order of acquisition is guaranteed when more than one thread blocks on the condition variable.

Note - `pthread_cond_wait()` is a cancellation point. If a cancel is pending and the calling thread has cancellation enabled, the thread terminates and begins executing its cleanup handlers while continuing to hold the lock.

Return Values

`pthread_cond_wait()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

EINVAL

The value specified by *cv* or *mp* is invalid.

Unblock a Specific Thread

pthread_cond_signal(3THR)

Use `pthread_cond_signal(3THR)` to unblock one thread that is blocked on the condition variable pointed to by *cv*. (For Solaris threads, see “`cond_signal(3THR)`” on page 218.)

Prototype:

```
int pthread_cond_signal(pthread_cond_t *cv);
```

```
#include <pthread.h>
```

```
pthread_cond_t cv;  
int ret;
```

```
/* one condition variable is signaled */  
ret = pthread_cond_signal(&cv);
```

Call `pthread_cond_signal( )` under the protection of the same mutex used with the condition variable being signaled. Otherwise, the condition variable could be signaled between the test of the associated condition and blocking in `pthread_cond_wait( )`, which can cause an infinite wait.

The scheduling policy determines the order in which blocked threads are awakened. For `SCHED_OTHER`, threads are awakened in priority order.

When no threads are blocked on the condition variable, calling `pthread_cond_signal( )` has no effect.

Return Values

`pthread_cond_signal( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

EINVAL

cv points to an illegal address.

Code Example 4-8 shows how to use `pthread_cond_wait( )` and `pthread_cond_signal( )`.

CODE EXAMPLE 4-8 Using `pthread_cond_wait()` and `pthread_cond_signal()`

```
pthread_mutex_t count_lock;
pthread_cond_t count_nonzero;
unsigned count;

decrement_count()
{
    pthread_mutex_lock(&count_lock);
    while (count == 0)
        pthread_cond_wait(&count_nonzero, &count_lock);
    count = count - 1;
    pthread_mutex_unlock(&count_lock);
}

increment_count()
{
    pthread_mutex_lock(&count_lock);
    if (count == 0)
        pthread_cond_signal(&count_nonzero);
    count = count + 1;
    pthread_mutex_unlock(&count_lock);
}
```

Block Until a Specified Event

`pthread_cond_timedwait(3THR)`

Prototype:

```
int pthread_cond_timedwait(pthread_cond_t *cv,
    pthread_mutex_t *mp, const struct timespec *abstime);
```

```
#include <pthread.h>
#include <time.h>
```

```
pthread_cond_t cv;
pthread_mutex_t mp;
timestruct_t abstime;
int ret;
```

```
/* wait on condition variable */
ret = pthread_cond_timedwait(&cv, &mp, &abstime);
```

Use `pthread_cond_timedwait(3THR)` as you would use `pthread_cond_wait()`, except that `pthread_cond_timedwait()` does not block past the time of day specified by *abstime*. `pthread_cond_timedwait()` always returns with the mutex locked and owned by the calling thread, even when it is returning an error. (For Solaris threads, see “`cond_timedwait(3THR)`” on page 218.)

The `pthread_cond_timedwait()` function blocks until the condition is signaled or until the time of day, specified by the last argument, has passed.

Note - `pthread_cond_timedwait()` is also a cancellation point.

Return Values

`pthread_cond_timedwait()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When either of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

cv or *abstime* points to an illegal address.

ETIMEDOUT

The time specified by *abstime* has passed.

The timeout is specified as a time of day so that the condition can be retested efficiently without recomputing the value, as shown in Code Example 4–9.

CODE EXAMPLE 4–9 Timed Condition Wait

```
pthread_timestruc_t to;
pthread_mutex_t m;
pthread_cond_t c;
...
pthread_mutex_lock(&m);
to.tv_sec = time(NULL) + TIMEOUT;
to.tv_nsec = 0;
while (cond == FALSE) {
    err = pthread_cond_timedwait(&c, &m, &to);
    if (err == ETIMEDOUT) {
        /* timeout, do something */
        break;
    }
}
pthread_mutex_unlock(&m);
```

Unblock All Threads

pthread_cond_broadcast(3THR)

Prototype:
`int pthread_cond_broadcast(pthread_cond_t *cv);`

```
#include <pthread.h>

pthread_cond_t cv;
```

```

int ret;

/* all condition variables are signaled */
ret = pthread_cond_broadcast(&cv);

```

Use `pthread_cond_broadcast(3THR)` to unblock all threads that are blocked on the condition variable pointed to by `cv`, specified by `pthread_cond_wait( )`. When no threads are blocked on the condition variable, `pthread_cond_broadcast( )` has no effect. (For Solaris threads, see “`cond_broadcast(3THR)`” on page 219.)

Return Values

`pthread_cond_broadcast( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

EINVAL

`cv` points to an illegal address.

Condition Variable Broadcast Example

Since `pthread_cond_broadcast( )` causes all threads blocked on the condition to contend again for the mutex lock, use it with care. For example, use `pthread_cond_broadcast( )` to allow threads to contend for varying resource amounts when resources are freed, as shown in Code Example 4–10.

CODE EXAMPLE 4–10 Condition Variable Broadcast

```

pthread_mutex_t rsrc_lock;
pthread_cond_t rsrc_add;
unsigned int resources;

get_resources(int amount)
{
    pthread_mutex_lock(&rsrc_lock);
    while (resources < amount) {
        pthread_cond_wait(&rsrc_add, &rsrc_lock);
    }
    resources -= amount;
    pthread_mutex_unlock(&rsrc_lock);
}

add_resources(int amount)
{
    pthread_mutex_lock(&rsrc_lock);
    resources += amount;
    pthread_cond_broadcast(&rsrc_add);
    pthread_mutex_unlock(&rsrc_lock);
}

```

Note that in `add_resources()` it does not matter whether *resources* is updated first or if `pthread_cond_broadcast()` is called first inside the mutex lock.

Call `pthread_cond_broadcast()` under the protection of the same mutex that is used with the condition variable being signaled. Otherwise, the condition variable could be signaled between the test of the associated condition and blocking in `pthread_cond_wait()`, which can cause an infinite wait.

Destroy Condition Variable State

`pthread_cond_destroy(3THR)`

Use `pthread_cond_destroy(3THR)` to destroy any state associated with the condition variable pointed to by *cv*. (For Solaris threads, see “`cond_destroy(3THR)`” on page 217.)

Prototype:

```
int pthread_cond_destroy(pthread_cond_t *cv);
```

```
#include <pthread.h>
```

```
pthread_cond_t cv;  
int ret;
```

```
/* Condition variable is destroyed */  
ret = pthread_cond_destroy(&cv);
```

Note that the space for storing the condition variable is not freed.

Return Values

`pthread_cond_destroy()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occur, the function fails and returns the corresponding value.

EINVAL

The value specified by *cv* is invalid.

The Lost Wake-Up Problem

Calling `pthread_cond_signal()` or `pthread_cond_broadcast()` when the thread does not hold the mutex lock associated with the condition can lead to *lost wake-up* bugs.

A lost wake-up occurs when:

- A thread calls `pthread_cond_signal( )` or `pthread_cond_broadcast( )`.
- *And* another thread is between the test of the condition and the call to `pthread_cond_wait( )`.
- *And* no threads are waiting.

The signal has no effect, and therefore is lost.

The Producer/Consumer Problem

This problem is one of the small collection of standard, well-known problems in concurrent programming: a finite-size buffer and two classes of threads, producers and consumers, put items into the buffer (producers) and take items out of the buffer (consumers).

A producer must wait until the buffer has space before it can put something in, and a consumer must wait until something is in the buffer before it can take something out.

A condition variable represents a queue of threads waiting for some condition to be signaled.

Code Example 4–11 has two such queues, one (*less*) for producers waiting for a slot in the buffer, and the other (*more*) for consumers waiting for a buffer slot containing information. The example also has a mutex, as the data structure describing the buffer must be accessed by only one thread at a time.

CODE EXAMPLE 4–11 The Producer/Consumer Problem and Condition Variables

```
typedef struct {
    char buf[BSIZE];
    int occupied;
    int nextin;
    int nextout;
    pthread_mutex_t mutex;
    pthread_cond_t more;
    pthread_cond_t less;
} buffer_t;

buffer_t buffer;
```

As Code Example 4–12 shows, the producer thread acquires the mutex protecting the `buffer` data structure and then makes certain that space is available for the item being produced. If not, it calls `pthread_cond_wait( )`, which causes it to join the queue of threads waiting for the condition *less*, representing *there is room in the buffer*, to be signaled.

At the same time, as part of the call to `pthread_cond_wait( )`, the thread releases its lock on the mutex. The waiting producer threads depend on consumer threads to signal when the condition is true (as shown in Code Example 4–12). When the

condition is signaled, the first thread waiting on *less* is awakened. However, before the thread can return from `pthread_cond_wait()`, it must acquire the lock on the mutex again.

This ensures that it again has mutually exclusive access to the buffer data structure. The thread then must check that there really is room available in the buffer; if so, it puts its item into the next available slot.

At the same time, consumer threads might be waiting for items to appear in the buffer. These threads are waiting on the condition variable *more*. A producer thread, having just deposited something in the buffer, calls `pthread_cond_signal()` to wake up the next waiting consumer. (If there are no waiting consumers, this call has no effect.)

Finally, the producer thread unlocks the mutex, allowing other threads to operate on the buffer data structure.

CODE EXAMPLE 4-12 The Producer/Consumer Problem—the Producer

```
void producer(buffer_t *b, char item)
{
    pthread_mutex_lock(&b->mutex);

    while (b->occupied >= BSIZE)
        pthread_cond_wait(&b->less, &b->mutex);

    assert(b->occupied < BSIZE);

    b->buf[b->nextin++] = item;

    b->nextin %= BSIZE;
    b->occupied++;

    /* now: either b->occupied < BSIZE and b->nextin is the index
       of the next empty slot in the buffer, or
       b->occupied == BSIZE and b->nextin is the index of the
       next (occupied) slot that will be emptied by a consumer
       (such as b->nextin == b->nextout) */

    pthread_cond_signal(&b->more);

    pthread_mutex_unlock(&b->mutex);
}
```

Note the use of the `assert()` statement; unless the code is compiled with `NDEBUG` defined, `assert()` does nothing when its argument evaluates to true (that is, nonzero), but causes the program to abort if the argument evaluates to false (zero). Such assertions are especially useful in multithreaded programs—they immediately point out runtime problems if they fail, and they have the additional effect of being useful comments.

The comment that begins `/* now: either b->occupied ...` could better be expressed as an assertion, but it is too complicated as a Boolean-valued expression and so is given in English.

Both the assertion and the comments are examples of invariants. These are logical statements that should not be falsified by the execution of the program, except during brief moments when a thread is modifying some of the program variables mentioned in the invariant. (An assertion, of course, should be true whenever any thread executes it.)

Using invariants is an extremely useful technique. Even if they are not stated in the program text, think in terms of invariants when you analyze a program.

The invariant in the producer code that is expressed as a comment is always true whenever a thread is in the part of the code where the comment appears. If you move this comment to just after the `mutex_unlock()`, this does not necessarily remain true. If you move this comment to just after the `assert()`, this is still true.

The point is that this invariant expresses a property that is true at all times, except when either a producer or a consumer is changing the state of the buffer. While a thread is operating on the buffer (under the protection of a mutex), it might temporarily falsify the invariant. However, once the thread is finished, the invariant should be true again.

Code Example 4–13 shows the code for the consumer. Its flow is symmetric with that of the producer.

CODE EXAMPLE 4–13 The Producer/Consumer Problem—the Consumer

```
char consumer(buffer_t *b)
{
    char item;
    pthread_mutex_lock(&b->mutex);
    while(b->occupied <= 0)
        pthread_cond_wait(&b->more, &b->mutex);

    assert(b->occupied > 0);

    item = b->buf[b->nextout++];
    b->nextout %= BSIZE;
    b->occupied--;

    /* now: either b->occupied > 0 and b->nextout is the index
       of the next occupied slot in the buffer, or
       b->occupied == 0 and b->nextout is the index of the next
       (empty) slot that will be filled by a producer (such as
       b->nextout == b->nextin) */

    pthread_cond_signal(&b->less);
    pthread_mutex_unlock(&b->mutex);

    return(item);
}
```

Semaphores

Semaphores are a programming construct designed by E. W. Dijkstra in the late 1960s. Dijkstra's model was the operation of railroads: consider a stretch of railroad in which there is a single track over which only one train at a time is allowed.

Guarding this track is a semaphore. A train must wait before entering the single track until the semaphore is in a state that permits travel. When the train enters the track, the semaphore changes state to prevent other trains from entering the track. A train that is leaving this section of track must again change the state of the semaphore to allow another train to enter.

In the computer version, a semaphore appears to be a simple integer. A thread waits for permission to proceed and then signals that it has proceeded by performing a P operation on the semaphore.

The semantics of the operation are such that the thread must wait until the semaphore's value is positive, then change the semaphore's value by subtracting one from it. When it is finished, the thread performs a V operation, which changes the semaphore's value by adding one to it. It is crucial that these operations take place atomically—they cannot be subdivided into pieces between which other actions on the semaphore can take place. In the P operation, the semaphore's value must be positive just before it is decremented (resulting in a value that is guaranteed to be nonnegative and one less than what it was before it was decremented).

In both P and V operations, the arithmetic must take place without interference. If two V operations are performed simultaneously on the same semaphore, the net effect should be that the semaphore's new value is two greater than it was.

The mnemonic significance of P and V is unclear to most of the world, as Dijkstra is Dutch. However, in the interest of true scholarship: P stands for *proberen te verlagen*, which means *try to decrease*. V stands for *verhogen*, which means *increase*. This is discussed in one of Dijkstra's technical notes, EWD 74.

`sem_wait(3R)` and `sem_post(3R)` correspond to Dijkstra's P and V operations. `sem_trywait(3R)` is a conditional form of the P operation: if the calling thread cannot decrement the value of the semaphore without waiting, the call returns immediately with a nonzero value.

There are two basic sorts of semaphores: binary semaphores, which never take on values other than zero or one, and counting semaphores, which can take on arbitrary nonnegative values. A binary semaphore is logically just like a mutex.

However, although it is not enforced, mutexes should be unlocked only by the thread holding the lock. There is no notion of “the thread holding the semaphore,” so any thread can perform a V (or `sem_post(3R)`) operation.

Counting semaphores are about as powerful as conditional variables (used in conjunction with mutexes). In many cases, the code might be simpler when it is implemented with counting semaphores rather than with condition variables (as shown in the next few examples).

However, when a mutex is used with condition variables, there is an implied bracketing—it is clear which part of the program is being protected. This is not necessarily the case for a semaphore, which might be called the *go to* of concurrent programming—it is powerful but too easy to use in an unstructured, indeterminate way.

Counting Semaphores

Conceptually, a semaphore is a nonnegative integer count. Semaphores are typically used to coordinate access to resources, with the semaphore count initialized to the number of free resources. Threads then atomically increment the count when resources are added and atomically decrement the count when resources are removed.

When the semaphore count becomes zero, indicating that no more resources are present, threads trying to decrement the semaphore block wait until the count becomes greater than zero.

TABLE 4-7 Routines for Semaphores

Operation	Destination Discussion
Initialize a semaphore	“sema_init(3THR)” on page 124
Increment a semaphore	“sem_post(3THR)” on page 126
Block on a semaphore count	“sema_wait(3THR)” on page 126
Decrement a semaphore count	“sema_trywait(3THR)” on page 127
Destroy the semaphore state	“sema_destroy(3THR)” on page 128

Because semaphores need not be acquired and released by the same thread, they can be used for asynchronous event notification (such as in signal handlers). And, because semaphores contain state, they can be used asynchronously without

acquiring a mutex lock as is required by condition variables. However, semaphores are not as efficient as mutex locks.

By default, there is no defined order of unblocking if multiple threads are waiting for a semaphore.

Semaphores must be initialized before use, but they do not have attributes.

Initialize a Semaphore

`sema_init(3THR)`

Prototype:

```
int sem_init(sem_t *sem, int pshared, unsigned int value);
```

```
#include <semaphore.h>
```

```
sem_t sem;  
int pshared;  
int ret;  
int value;
```

```
/* initialize a private semaphore */  
pshared = 0;  
value = 1;  
ret = sem_init(&sem, pshared, value);
```

Use `sema_init(3THR)` to initialize the semaphore variable pointed to by `sem` to `value` amount. If the value of `pshared` is zero, then the semaphore cannot be shared between processes. If the value of `pshared` is nonzero, then the semaphore can be shared between processes. (For Solaris threads, see “`sema_init(3THR)`” on page 219.)

Multiple threads must not initialize the same semaphore.

A semaphore must not be reinitialized while other threads might be using it.

Return Values

`sem_init( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

The value argument exceeds `SEM_VALUE_MAX`.

ENOSPC

A resource required to initialize the semaphore has been exhausted. The limit on semaphores SEM_NSEMS_MAX has been reached.

EPERM

The process lacks the appropriate privileges to initialize the semaphore.

Initializing Semaphores With Intraprocess Scope

When *pshared* is 0, the semaphore can be used by all the threads in this process only.

```
#include <semaphore.h>

sem_t sem;
int ret;
int count = 4;

/* to be used within this process only */
ret = sem_init(&sem, 0, count);
```

Initializing Semaphores With Interprocess Scope

When *pshared* is nonzero, the semaphore can be shared by other processes.

```
#include <semaphore.h>

sem_t sem;
int ret;
int count = 4;

/* to be shared among processes */
ret = sem_init(&sem, 1, count);
```

Named Semaphores

The functions `sem_open(3R)`, `sem_getvalue(3R)`, `sem_close(3R)`, and `sem_unlink(3R)` are available to open, retrieve, close, and remove named semaphores. Using `sem_open( )`, you can create a semaphore that has a name defined in the file system name space.

Named semaphores are like process shared semaphores, except that they are referenced with a pathname rather than a *pshared* value.

For more information about named semaphores, see `sem_open(3R)`, `sem_getvalue(3R)`, `sem_close(3R)`, and `sem_unlink(3R)`.

Increment a Semaphore

`sem_post(3THR)`

Prototype:

```
int sem_post(sem_t *sem);
```

```
#include <semaphore.h>
```

```
sem_t sem;
```

```
int ret;
```

```
ret = sem_post(&sem); /* semaphore is posted */
```

Use `sema_post(3THR)` to atomically increment the semaphore pointed to by *sem*. When any threads are blocked on the semaphore, one of them is unblocked. (For Solaris threads, see “`sema_post(3THR)`” on page 220.)

Return Values

`sem_post( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

EINVAL

sem points to an illegal address.

Block on a Semaphore Count

`sema_wait(3THR)`

Prototype:

```
int sem_wait(sem_t *sem);
```

```
#include <semaphore.h>
```

```
sem_t sem;
```

```
int ret;
```

```
ret = sem_wait(&sem); /* wait for semaphore */
```

Use `sema_wait(3THR)` to block the calling thread until the count in the semaphore pointed to by *sem* becomes greater than zero, then atomically decrement it.

Return Values

`sem_wait( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

sem points to an illegal address.

EINTR

A signal interrupted this function.

Decrement a Semaphore Count

`sema_trywait(3THR)`

Prototype:

```
int sem_trywait(sem_t *sem);
```

```
#include <semaphore.h>
```

```
sem_t sem;  
int ret;
```

```
ret = sem_trywait(&sem); /* try to wait for semaphore*/
```

Use `sema_trywait(3THR)` to try to atomically decrement the count in the semaphore pointed to by *sem* when the count is greater than zero. This function is a nonblocking version of `sem_wait( )`; that is it returns immediately if unsuccessful.

Return Values

`sem_trywait( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

sem points to an illegal address.

EINTR

A signal interrupted this function.

EAGAIN

The semaphore was already locked, so it cannot be immediately locked by the `sem_trywait()` operation.

Destroy the Semaphore State

`sema_destroy(3THR)`

Prototype:

```
int sem_destroy(sem_t *sem);
```

```
#include <semaphore.h>
```

```
sem_t sem;  
int ret;
```

```
ret = sem_destroy(&sem); /* the semaphore is destroyed */
```

Use `sema_destroy(3THR)` to destroy any state associated with the semaphore pointed to by `sem`. The space for storing the semaphore is not freed. (For Solaris threads, see “`sem_destroy(3RT)`” on page 221.)

Return Values

`sem_destroy()` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, the function fails and returns the corresponding value.

EINVAL

`sem` points to an illegal address.

The Producer/Consumer Problem, Using Semaphores

The data structure in Code Example 4–14 is similar to that used for the condition variables example (see Code Example 4–11). Two semaphores represent the number of full and empty buffers and ensure that producers wait until there are empty buffers and that consumers wait until there are full buffers.

CODE EXAMPLE 4–14 The Producer/Consumer Problem With Semaphores

```
typedef struct {  
    char buf[BSIZE];
```

```

    sem_t occupied;
    sem_t empty;
    int nextin;
    int nextout;
    sem_t pmut;
    sem_t cmut;
} buffer_t;

buffer_t buffer;

sem_init(&buffer.occupied, 0, 0);

sem_init(&buffer.empty, 0, BSIZE);
sem_init(&buffer.pmut, 0, 1);
sem_init(&buffer.cmut, 0, 1);
buffer.nextin = buffer.nextout = 0;

```

Another pair of (binary) semaphores plays the same role as mutexes, controlling access to the buffer when there are multiple producers and multiple empty buffer slots, and when there are multiple consumers and multiple full buffer slots. Mutexes would work better here, but would not provide as good an example of semaphore use.

CODE EXAMPLE 4-15 The Producer/Consumer Problem—the Producer

```

void producer(buffer_t *b, char item) {
    sem_wait(&b->empty);

    sem_wait(&b->pmut);

    b->buf[b->nextin] = item;
    b->nextin++;
    b->nextin %= BSIZE;

    sem_post(&b->pmut);

    sem_post(&b->occupied);
}

```

CODE EXAMPLE 4-16 The Producer/Consumer Problem—the Consumer

```

char consumer(buffer_t *b) {
    char item;

    sem_wait(&b->occupied);

    sem_wait(&b->cmut);

    item = b->buf[b->nextout];
    b->nextout++;
    b->nextout %= BSIZE;

    sem_post(&b->cmut);

    sem_post(&b->empty);
}

```

```

    return(item);
}

```

Read-Write Lock Attributes

Read-write locks permit concurrent reads and exclusive writes to a protected shared resource. The read-write lock is a single entity that can be locked in *read* or *write* mode. To modify a resource, a thread must first acquire the exclusive write lock. An exclusive write lock is not permitted until all read locks have been released.

Database access can be synchronized with a read-write lock. Read-write locks support concurrent reads of database records because the read operation does not change the record's information. When the database is to be updated, the write operation must acquire an exclusive write lock.

To change the default read-write lock attributes, you can declare and initialize an attribute object. Often, the read-write lock attributes are set up in one place at the beginning of the application so they can be located quickly and modified easily. The following table lists the functions discussed in this section that manipulate read-write lock attributes.

See “Similar Synchronization Functions—Read-Write Locks” on page 198 for the Solaris threads implementation of read-write locks.

TABLE 4-8 Routines for Read-Write Lock Attributes

Operation	Destination Discussion
Initialize a read-write lock attribute	“pthread_rwlockattr_init(3THR)” on page 131
Destroy a read-write lock attribute	“pthread_rwlockattr_destroy(3THR)” on page 131
Set a read-write lock attribute	“pthread_rwlockattr_setpshared(3THR)” on page 132
Get a read-write lock attribute	“pthread_rwlockattr_getpshared(3THR)” on page 132

Initialize a Read-Write Lock Attribute

`pthread_rwlockattr_init(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlockattr_init(pthread_rwlockattr_t *attr);
```

`pthread_rwlockattr_init(3THR)` initializes a read-write lock attributes object *attr* with the default value for all of the attributes defined by the implementation.

Results are undefined if `pthread_rwlockattr_init` is called specifying an already initialized read-write lock attributes object. After a read-write lock attributes object has been used to initialize one or more read-write locks, any function affecting the attributes object (including destruction) does not affect any previously initialized read-write locks.

Return Values

If successful, `pthread_rwlockattr_init( )` returns zero. Otherwise, an error number is returned to indicate the error.

ENOMEM

Insufficient memory exists to initialize the rwlock attributes object.

Destroy a Read-Write Lock Attribute

`pthread_rwlockattr_destroy(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlockattr_destroy(pthread_rwlockattr_t *attr);
```

`pthread_rwlockattr_destroy(3THR)` destroys a read-write lock attributes object. The effect of subsequent use of the object is undefined until the object is re-initialized by another call to `pthread_rwlockattr_init( )`. An implementation can cause `pthread_rwlockattr_destroy( )` to set the object referenced by *attr* to an invalid value.

Return Values

If successful, `pthread_rwlockattr_destroy( )` returns zero. Otherwise, an error number is returned to indicate the error.

EINVAL

The value specified by *attr* is invalid.

Set a Read-Write Lock Attribute

`pthread_rwlockattr_setpshared(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlockattr_setpshared(pthread_rwlockattr_t *attr, int pshared);
```

`pthread_rwlockattr_setpshared(3THR)` sets the process-shared read-write lock attribute.

PTHREAD_PROCESS_SHARED

Permits a read-write lock to be operated on by any thread that has access to the memory where the read-write lock is allocated, even if the read-write lock is allocated in memory that is shared by multiple processes.

PTHREAD_PROCESS_PRIVATE

The read-write lock will only be operated upon by threads created within the same process as the thread that initialized the read-write lock; if threads of differing processes attempt to operate on such a read-write lock, the behavior is undefined. The default value of the process-shared attribute is **PTHREAD_PROCESS_PRIVATE**.

Return Value

If successful, `pthread_rwlockattr_setpshared( )` returns zero. Otherwise, an error number is returned to indicate the error.

EINVAL

The value specified by *attr* or *pshared* is invalid.

Get a Read-Write Lock Attribute

`pthread_rwlockattr_getpshared(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlockattr_getpshared(const pthread_rwlockattr_t *attr, int *pshared);
```

`pthread_rwlockattr_getpshared(3THR)` gets the process-shared read-write lock attribute.

`pthread_rwlockattr_getpshared( )` obtains the value of the process-shared attribute from the initialized attributes object referenced by *attr*.

Return Value

If successful, `pthread_rwlockattr_getpshared( )` returns zero. Otherwise, an error number is returned to indicate the error.

EINVAL

The value specified by *attr* or *pshared* is invalid.

Using Read-Write Locks

After the attributes for a read-write lock are configured, you initialize the read-write lock itself. The following functions are used to initialize or destroy, lock or unlock, or try to lock a read-write lock. The following table lists the functions discussed in this section that manipulate read-write locks.

TABLE 4-9 Routines that Manipulate Read-Write Locks

Operation	Destination Discussion
Initialize a read-write lock	“pthread_rwlock_init(3THR)” on page 134
Read lock on read-write lock	“pthread_rwlock_rdlock(3THR)” on page 135
Read lock with a nonblocking read-write lock	“pthread_rwlock_tryrdlock(3THR)” on page 136
Write lock on read-write lock	“pthread_rwlock_wrlock(3THR)” on page 136
Write lock with a nonblocking read-write lock	“pthread_rwlock_trywrlock(3THR)” on page 137

TABLE 4-9 Routines that Manipulate Read-Write Locks *(continued)*

Operation	Destination Discussion
Unlock a read-write lock	“pthread_rwlock_unlock(3THR)” on page 138
Destroy a read-write lock	“pthread_rwlock_destroy(3THR)” on page 139

Initialize a Read-Write Lock

pthread_rwlock_init(3THR)

```
#include <pthread.h>
```

```
int pthread_rwlock_init(pthread_rwlock_t *rwlock, const pthread_rwlockattr_t *attr);
```

```
pthread_rwlock_t  rwlock = PTHREAD_RWLOCK_INITIALIZER;
```

Use `pthread_rwlock_init(3THR)` to initialize the read-write lock referenced by *rwlock* with the attributes referenced by *attr*. If *attr* is `NULL`, the default read-write lock attributes are used; the effect is the same as passing the address of a default read-write lock attributes object. Once initialized, the lock can be used any number of times without being re-initialized. On successful initialization, the state of the read-write lock becomes initialized and unlocked. Results are undefined if `pthread_rwlock_init()` is called specifying an already initialized read-write lock. Results are undefined if a read-write lock is used without first being initialized. (For Solaris threads, see “`rwlock_init(3THR)`” on page 198.)

In cases where default read-write lock attributes are appropriate, the macro `PTHREAD_RWLOCK_INITIALIZER` can be used to initialize read-write locks that are statically allocated. The effect is equivalent to dynamic initialization by a call to `pthread_rwlock_init()` with the parameter *attr* specified as `NULL`, except that no error checks are performed.

Return Value

If successful, `pthread_rwlock_init()` returns zero. Otherwise, an error number is returned to indicate the error.

If `pthread_rwlock_init()` fails, *rwlock* is not initialized and the contents of *rwlock* are undefined.

EINVAL

The value specified by *attr* or *rwlock* is invalid.

Read Lock on Read-Write Lock

`pthread_rwlock_rdlock(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlock_rdlock(pthread_rwlock_t *rwlock );
```

`pthread_rwlock_rdlock(3THR)` applies a read lock to the read-write lock referenced by *rwlock*. The calling thread acquires the read lock if a writer does not hold the lock and there are no writers blocked on the lock. It is unspecified whether the calling thread acquires the lock when a writer does not hold the lock and there are writers waiting for the lock. If a writer holds the lock, the calling thread will not acquire the read lock. If the read lock is not acquired, the calling thread blocks (that is, it does not return from the `pthread_rwlock_rdlock( )`) until it can acquire the lock. Results are undefined if the calling thread holds a write lock on *rwlock* at the time the call is made.

Implementations are allowed to favor writers over readers to avoid writer starvation. (For instance, the Solaris threads implementation favors writers over readers. See “`rw_rdlock(3THR)`” on page 200.)

A thread can hold multiple concurrent read locks on *rwlock* (that is, successfully call `pthread_rwlock_rdlock( )` *n* times). If so, the thread must perform matching unlocks (that is, it must call `pthread_rwlock_unlock( )` *n* times).

Results are undefined if `pthread_rwlock_rdlock( )` is called with an uninitialized read-write lock.

If a signal is delivered to a thread waiting for a read-write lock for reading, on return from the signal handler the thread resumes waiting for the read-write lock for reading as if it was not interrupted.

Return Value

If successful, `pthread_rwlock_rdlock( )` returns zero. Otherwise, an error number is returned to indicate the error.

EINVAL

The value specified by *attr* or *rwlock* is invalid.

Read Lock With a Nonblocking Read-Write Lock

`pthread_rwlock_tryrdlock(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlock_tryrdlock(pthread_rwlock_t *rwlock);
```

`pthread_rwlock_tryrdlock(3THR)` applies a read lock as in `pthread_rwlock_rdlock( )` with the exception that the function fails if any thread holds a write lock on *rwlock* or there are writers blocked on *rwlock*. (For Solaris threads, see “`rw_tryrdlock(3THR)`” on page 200.)

Return Value

`pthread_rwlock_tryrdlock( )` returns zero if the lock for reading on the read-write lock object referenced by *rwlock* is acquired. Otherwise an error number is returned to indicate the error.

EBUSY

The read-write lock could not be acquired for reading because a writer holds the lock or was blocked on it.

Write Lock on Read-Write Lock

`pthread_rwlock_wrlock(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlock_wrlock(pthread_rwlock_t *rwlock );
```

`pthread_rwlock_wrlock(3THR)` applies a write lock to the read-write lock referenced by *rwlock*. The calling thread acquires the write lock if no other thread (reader or writer) holds the read-write lock *rwlock*. Otherwise, the thread blocks (that is, does not return from the `pthread_rwlock_wrlock( )` call) until it can acquire the lock. Results are undefined if the calling thread holds the read-write lock (whether a read or write lock) at the time the call is made.

Implementations are allowed to favor writers over readers to avoid writer starvation. (For instance, the Solaris threads implementation favors writers over readers. See “`rw_wrlock(3THR)`” on page 201.)

Results are undefined if `pthread_rwlock_wrlock( )` is called with an uninitialized read-write lock.

If a signal is delivered to a thread waiting for a read-write lock for writing, upon return from the signal handler the thread resumes waiting for the read-write lock for writing as if it was not interrupted.

Return Value

`pthread_rwlock_rwlock( )` returns zero if the lock for writing on the read-write lock object referenced by *rwlock* is acquired. Otherwise an error number is returned to indicate the error.

Write Lock With a Nonblocking Read-Write Lock

`pthread_rwlock_trywrlock(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlock_trywrlock(pthread_rwlock_t *rwlock);
```

`pthread_rwlock_trywrlock(3THR)` applies a write lock like `pthread_rwlock_wrlock( )`, with the exception that the function fails if any thread currently holds *rwlock* (for reading or writing). (For Solaris threads, see “`rw_trywrlock(3THR)`” on page 202.)

Results are undefined if `pthread_rwlock_trywrlock( )` is called with an uninitialized read-write lock.

If a signal is delivered to a thread waiting for a read-write lock for writing, on return from the signal handler the thread resumes waiting for the read-write lock for writing as if it was not interrupted.

Return Value

If successful, `pthread_rwlock_trywrlock( )` returns zero if the lock for writing on the read-write lock object referenced by *rwlock* is acquired. Otherwise, an error number is returned to indicate the error.

EBUSY

The read-write lock could not be acquired for writing because it is already locked for reading or writing.

Unlock a Read-Write Lock

`pthread_rwlock_unlock(3THR)`

`#include <pthread.h>`

`pthread_rwlock_unlock(3THR)` releases a lock held on the read-write lock object referenced by *rwlock*. Results are undefined if the read-write lock *rwlock* is not held by the calling thread. (For Solaris threads, see “`rw_unlock(3THR)`” on page 202.)

If `pthread_rwlock_unlock` is called to release a read lock from the read-write lock object and there are other read locks currently held on this read-write lock object, the read-write lock object remains in the read locked state. If `pthread_rwlock_unlock( )` releases the calling thread’s last read lock on this read-write lock object, then the calling thread is no longer one of the owners of the object. If `pthread_rwlock_unlock( )` releases the last read lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.

If `pthread_rwlock_unlock( )` is called to release a write lock for this read-write lock object, the read-write lock object will be put in the unlocked state with no owners.

If the call to the `pthread_rwlock_unlock( )` results in the read-write lock object becoming unlocked and there are multiple threads waiting to acquire the read-write lock object for writing, the scheduling policy is used to determine which thread acquires the read-write lock object for writing. If there are multiple threads waiting to acquire the read-write lock object for reading, the scheduling policy is used to determine the order in which the waiting threads acquire the read-write lock object for reading. If there are multiple threads blocked on *rwlock* for both read locks and write locks, it is unspecified whether the readers acquire the lock first or whether a writer acquires the lock first.

Results are undefined if `pthread_rwlock_unlock( )` is called with an uninitialized read-write lock.

Return Value

If successful, `pthread_rwlock_unlock( )` returns zero. Otherwise, an error number is returned to indicate the error.

Destroy a Read-Write Lock

`pthread_rwlock_destroy(3THR)`

```
#include <pthread.h>
```

```
int pthread_rwlock_destroy(pthread_rwlock_t *rwlock);
```

```
pthread_rwlock_t  rwlock = PTHREAD_RWLOCK_INITIALIZER;
```

`pthread_rwlock_destroy(3THR)` destroys the read-write lock object referenced by *rwlock* and releases any resources used by the lock. The effect of subsequent use of the lock is undefined until the lock is re-initialized by another call to `pthread_rwlock_init( )`. An implementation can cause `pthread_rwlock_destroy( )` to set the object referenced by *rwlock* to an invalid value. Results are undefined if `pthread_rwlock_destroy( )` is called when any thread holds *rwlock*. Attempting to destroy an uninitialized read-write lock results in undefined behavior. A destroyed read-write lock object can be re-initialized using `pthread_rwlock_init( )`; the results of otherwise referencing the read-write lock object after it has been destroyed are undefined. (For Solaris threads, see “`pthread_rwlock_destroy(3THR)`” on page 203.)

Return Value

If successful, `pthread_rwlock_destroy( )` returns zero. Otherwise, an error number is returned to indicate the error.

EINVAL

The value specified by *attr* or *rwlock* is invalid.

Synchronization Across Process Boundaries

Each of the synchronization primitives can be set up to be used across process boundaries. This is done quite simply by ensuring that the synchronization variable is located in a shared memory segment and by calling the appropriate `init( )` routine, after the primitive has been initialized with its shared attribute set as `interprocess`.

Producer/Consumer Problem Example

Code Example 4-17 shows the producer/consumer problem with the producer and consumer in separate processes. The main routine maps zero-filled memory (that it shares with its child process) into its address space.

A child process is created that runs the consumer. The parent runs the producer.

This example also shows the drivers for the producer and consumer. The `producer_driver()` simply reads characters from `stdin` and calls `producer()`. The `consumer_driver()` gets characters by calling `consumer()` and writes them to `stdout`.

The data structure for Code Example 4-17 is the same as that used for the condition variables example (see Code Example 4-4). Two semaphores represent the number of full and empty buffers and ensure that producers wait until there are empty buffers and that consumers wait until there are full buffers.

CODE EXAMPLE 4-17 Synchronization Across Process Boundaries

```
main() {
    int zfd;
    buffer_t *buffer;
    pthread_mutexattr_t mattr;
    pthread_condattr_t cvattr_less, cvattr_more;

    zfd = open("/dev/zero", O_RDWR);
    buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
        PROT_READ|PROT_WRITE, MAP_SHARED, zfd, 0);
    buffer->occupied = buffer->nextin = buffer->nextout = 0;

    pthread_mutexattr_init(&mattr);
    pthread_mutexattr_setpshared(&mattr,
        PTHREAD_PROCESS_SHARED);

    pthread_mutex_init(&buffer->lock, &mattr);
    pthread_condattr_init(&cvattr_less);
    pthread_condattr_setpshared(&cvattr_less, PTHREAD_PROCESS_SHARED);
    pthread_cond_init(&buffer->less, &cvattr_less);
    pthread_condattr_init(&cvattr_more);
    pthread_condattr_setpshared(&cvattr_more,
        PTHREAD_PROCESS_SHARED);
    pthread_cond_init(&buffer->more, &cvattr_more);

    if (fork() == 0)
        consumer_driver(buffer);
    else
        producer_driver(buffer);
}

void producer_driver(buffer_t *b) {
    int item;

    while (1) {
        item = getchar();
        if (item == EOF) {
            producer(b, '\0');
        }
    }
}
```

```

        break;
    } else
        producer(b, (char)item);
    }
}

void consumer_driver(buffer_t *b) {
    char item;

    while (1) {
        if ((item = consumer(b)) == '\0')
            break;
        putchar(item);
    }
}

```

Interprocess Locking Without the Threads Library

Although not generally recommended, you can use Solaris threads to do interprocess locking without using the threads library. If this is something you want to do, see the instructions in “Using LWPs Between Processes” on page 222.

Comparing Primitives

The most basic synchronization primitive in threads is the mutual exclusion lock. So, it is the most efficient mechanism in both memory use and execution time. The basic use of a mutual exclusion lock is to serialize access to a resource.

The next most efficient primitive in threads is the condition variable. The basic use of a condition variable is to block on a change of state; that is it provides a thread wait facility. Remember that a mutex lock must be acquired before blocking on a condition variable and must be unlocked after returning from `pthread_cond_wait( )`. The mutex lock must also be held across the change of state that occurs before the corresponding call to `pthread_cond_signal( )`.

The semaphore uses more memory than the condition variable. It is easier to use in some circumstances because a semaphore variable functions on state rather than on control. Unlike a lock, a semaphore does not have an owner. Any thread can increment a semaphore that has blocked.

The read-write lock permits concurrent reads and exclusive writes to a protected resource. The read-write lock is a single entity that can be locked in *read* or *write*

mode. To modify a resource, a thread must first acquire the exclusive write lock. An exclusive write lock is not permitted until all read locks have been released.

Programming With the Operating Environment

This chapter describes how multithreading interacts with the Solaris operating environment and how the operating environment has changed to support multithreading.

- “Process Creation—`exec(2)` and `exit(2)` Issues” on page 148
- “Timers, Alarms, and Profiling” on page 148
- “Nonlocal Goto—`setjmp(3C)` and `longjmp(3C)`” on page 150
- “Resource Limits” on page 150
- “LWPs and Scheduling Classes” on page 150
- “Extending Traditional Signals” on page 154
- “I/O Issues” on page 164

Process Creation—Forking Issues

The default handling of the `fork( )` function in the Solaris operating environment is somewhat different from the way `fork( )` is handled in POSIX threads, although the Solaris operating environment does support both mechanisms.

Table 5–1 compares the differences and similarities of Solaris and pthreads `fork( )` handling. When the comparable interface is not available either in POSIX threads or in Solaris threads, the ‘—’ character appears in the table column.

TABLE 5-1 Comparing POSIX and Solaris `fork( )` Handling

	Solaris Operating Environment Interface	POSIX Threads Interface
Fork-one model	<code>fork1(2)</code>	<code>fork(2)</code>
Fork-all model	<code>fork(2)</code>	—
Fork safety	—	<code>pthread_atfork(3T)</code>

The Fork-One Model

As shown in Table 5-1, the behavior of the pthreads `fork(2)` function is the same as that of the Solaris `fork1(2)` function. Both the pthreads `fork(2)` function and the Solaris `fork1(2)` create a new process, duplicating the complete address space in the child, but duplicating only the calling thread in the child process.

This is useful when the child process immediately calls `exec( )`, which is what happens after most calls to `fork( )`. In this case, the child process does not need a duplicate of any thread other than the one that called `fork( )`.

In the child, do not call any library functions after calling `fork()` and before calling `exec( )` because one of the library functions might use a lock that was held in the parent at the time of the `fork( )`. The child process might execute only Async-Signal-Safe operations until one of the `exec( )` handlers is called.

The Fork-One Safety Problem and Solution

In addition to all of the usual concerns such as locking shared data, a library should be well behaved with respect to forking a child process when only one thread is running (the one that called `fork( )`). The problem is that the sole thread in the child process might try to grab a lock that is held by a thread that wasn't duplicated in the child.

This is not a problem most programs are likely to encounter. Most programs call `exec( )` in the child right after the return from `fork( )`. However, if the program wishes to carry out some actions in the child before the call to `exec( )`, or never calls `exec( )`, then the child *could* encounter deadlock scenarios.

Each library writer should provide a safe solution, although not providing a fork-safe library is not a large concern because this condition is rare.

For example, assume that T1 is in the middle of printing something (and so is holding a lock for `printf()`), when T2 forks a new process. In the child process, if the sole thread (T2) calls `printf()`, it promptly deadlocks.

The POSIX `fork()` or Solaris `fork1()` duplicates only the thread that calls it. (Calling the Solaris `fork()` duplicates all threads, so this issue does not come up.)

To prevent deadlock, ensure that no such locks are being held at the time of forking. The most obvious way to do this is to have the forking thread acquire all the locks that could possibly be used by the child. Because you cannot do this for locks like those in `printf()` (because `printf()` is owned by `libc`), you must ensure that `printf()` is not being used at `fork()` time.

To manage the locks in your library:

- Identify all the locks used by the library.
- Identify the locking order for the locks used by the library. (If a strict locking order is not used, then lock acquisition must be managed carefully.)
- Arrange to acquire those locks at fork time. In Solaris threads this must be done manually, obtaining the locks just before calling `fork1()`, and releasing them right after.

In the following example, the list of locks used by the library is $\{L1, \dots, L_n\}$, and the locking order for these locks is also $L1 \dots L_n$.

```
mutex_lock(L1);
mutex_lock(L2);
fork1(...);
mutex_unlock(L1);
mutex_unlock(L2);
```

In pthreads, you can add a call to `pthread_atfork(f1, f2, f3)` in your library's `.init()` section, where `f1()`, `f2()`, `f3()` are defined as follows:

```
f1() /* This is executed just before the process forks. */
{
    mutex_lock(L1); |
    mutex_lock(...); | -- ordered in lock order
    mutex_lock(Ln); |
} V

f2() /* This is executed in the child after the process forks. */
{
    mutex_unlock(L1);
    mutex_unlock(...);
    mutex_unlock(Ln);
}

f3() /* This is executed in the parent after the process forks. */
{
    mutex_unlock(L1);
    mutex_unlock(...);
    mutex_unlock(Ln);
}
```

Another example of deadlock would be a thread in the parent process—other than the one that called Solaris `fork1(2)`—that has locked a mutex. This mutex is copied into the child process in its locked state, but no thread is copied over to unlock the mutex. So, any thread in the child that tries to lock the mutex waits forever.

Virtual Forks—`vfork(2)`

The standard `vfork(2)` function is unsafe in multithreaded programs. `vfork(2)` is like `fork1(2)` in that only the calling thread is copied in the child process. As in nonthreaded implementations, `vfork( )` does not copy the address space for the child process.

Be careful that the thread in the child process does not change memory before it calls `exec(2)`. Remember that `vfork( )` gives the parent address space to the child. The parent gets its address space back after the child calls `exec( )` or exits. It is important that the child not change the state of the parent.

For example, it is dangerous to create new threads between the call to `vfork( )` and the call to `exec( )`. This is an issue only if the fork-one model is used, and only if the child does more than just call `exec( )`. Most libraries are not fork safe, so use `pthread_atfork( )` to implement fork safety.

The Solution—`pthread_atfork(3T)`

Use `pthread_atfork( )` to prevent deadlocks whenever you use the fork-one model.

```
#include <pthread.h>
```

```
int pthread_atfork(void (*prepare) (void), void (*parent) (void),  
void (*child) (void) );
```

The `pthread_atfork( )` function declares `fork( )` handlers that are called before and after `fork( )` in the context of the thread that called `fork( )`.

- The *prepare* handler is called before `fork( )` starts.
- The *parent* handler is called after `fork( )` returns in the parent.
- The *child* handler is called after `fork( )` returns in the child.

Any one of these can be set to `NULL`. The order in which successive calls to `pthread_atfork( )` are made is significant.

For example, a *prepare* handler could acquire all the mutexes needed, and then the *parent* and *child* handlers could release them. This ensures that all the relevant locks are held by the thread that calls the fork function *before* the process is forked, preventing the deadlock in the child.

Using the fork all model avoids the deadlock problem described in “The Fork-One Safety Problem and Solution” on page 144.

Return Values

`pthread_atfork( )` returns a zero when it completes successfully. Any other returned value indicates that an error occurred. If the following condition is detected, `pthread_atfork(3T)` fails and returns the corresponding value.

ENOMEM

Insufficient table space exists to record the fork handler addresses.

The Fork-All Model

The Solaris `fork(2)` function duplicates the address space and all the threads (and LWPs) in the child. This is useful, for example, when the child process never calls `exec(2)` but does use its copy of the parent address space. The fork-all functionality is not available in POSIX threads.

Note that when one thread in a process calls Solaris `fork(2)`, threads that are blocked in an interruptible system call return `EINTR`.

Also, be careful not to create locks that are held by both the parent and child processes. This can happen when locks are allocated in memory that is sharable (that is use `mmap( )` with the `MAP_SHARED` flag). Note that this is not a problem if the fork-one model is used.

Choosing the Right Fork

You determine whether `fork( )` has a “fork-all” or a “fork-one” semantic in your application by linking with the appropriate library. Linking with `-lthread` gives you the “fork-all” semantic for `fork( )`, and linking with `-lpthread` gives the “fork-one” semantic for `fork( )` (see Figure 7-1 for an explanation of compiling options).

Cautions for Any Fork

Be careful when using global state after a call to any `fork( )` function.

For example, when one thread reads a file serially and another thread in the process successfully calls one of the forks, each process then contains a thread that is reading the file. Because the seek pointer for a file descriptor is shared after a `fork( )`, the thread in the parent gets some data while the thread in the child gets the other. This introduces gaps in the sequential read accesses.

Process Creation—`exec(2)` and `exit(2)` Issues

Both the `exec(2)` and `exit(2)` system calls work as they do in single-threaded processes except that they destroy all the threads in the address space. Both calls block until all the execution resources (and so all active threads) are destroyed.

When `exec( )` rebuilds the process, it creates a single lightweight process (LWP) . The process startup code builds the initial thread. As usual, if the initial thread returns, it calls `exit( )` and the process is destroyed.

When all the threads in a process exit, the process exits. A call to any `exec( )` function from a process with more than one thread terminates all threads, and loads and executes the new executable image. No destructor functions are called.

Timers, Alarms, and Profiling

The “End of Life” announcements for per-LWP timers (see `timer_create(3R)`) and per-thread alarms (see `alarm(2)` or `setitimer(2)`) were made in the Solaris 2.5 release. Both features are now supplemented with the per-process variants described in this section.

Originally, each LWP had a unique realtime interval timer and alarm that a thread bound to the LWP could use. The timer or alarm delivered one signal to the thread when the timer or alarm expired.

Each LWP also had a virtual time or profile interval timer that a thread bound to the LWP could use. When the interval timer expired, either `SIGVTALRM` or `SIGPROF`, as appropriate, was sent to the LWP that owned the interval timer.

Per-LWP POSIX Timers

In the Solaris 2.3 and 2.4 releases, the `timer_create(3R)` function returned a timer object with a timer ID meaningful only within the calling LWP and with expiration signals delivered to that LWP. Because of this, the only threads that could use the POSIX timer facility were bound threads.

Even with this restricted use, POSIX timers in the Solaris 2.3 and 2.4 releases for multithreaded applications were unreliable about masking the resulting signals and delivering the associated value from the `sigevent` structure.

Beginning with the Solaris 2.5 release, an application that is compiled defining the macro `_POSIX_PER_PROCESS_TIMERS`, or with a value greater than 199506L for the symbol `_POSIX_C_SOURCE`, can create per-process timers.

Applications compiled with a release before the Solaris 2.5 release, or without the feature test macros, will continue to create per-LWP POSIX timers. In some future release, calls to create per-LWP timers will return per-process timers.

The timer IDs of per-process timers are usable from any LWP, and the expiration signals are generated for the process rather than directed to a specific LWP.

The per-process timers are deleted only by `timer_delete(3R)` or when the process terminates.

Per-Thread Alarms

In the Solaris Operating Environment 2.3 and 2.4 releases, a call to `alarm(2)` or `setitimer(2)` was meaningful only within the calling LWP. Such timers were deleted automatically when the creating LWP terminated. Because of this, the only threads that could use `alarm()` or `setitimer()` were bound threads.

Even with this restricted use, `alarm()` and `setitimer()` timers in Solaris Operating Environment 2.3 and 2.4 multithreaded applications were unreliable about masking the signals from the bound thread that issued these calls. When such masking was not required, then these two system calls worked reliably from bound threads.

With the Solaris Operating Environment 2.5 release, an application linking with `-lpthread` (POSIX) threads will get per-process delivery of `SIGALRM` when calling `alarm()`. The `SIGALRM` generated by `alarm()` is generated for the process rather than directed to a specific LWP. Also, the alarm is reset when the process terminates.

Applications compiled with a release before the Solaris Operating Environment 2.5 release, or not linked with `-lpthread`, will continue to see a per-LWP delivery of signals generated by `alarm()` and `setitimer()`.

In some future release, calls to `alarm()` or to `setitimer()` with the `ITIMER_REAL` flag will cause the resulting `SIGALRM` to be sent to the process. For other flags, `setitimer()` will continue to be per-LWP. Flags other than the `ITIMER_REAL` flag, used by `setitimer()`, will continue to result in the generated signal being delivered to the LWP that issued the call, and so are usable only from bound threads.

Profiling

You can profile each LWP with `profil(2)`, giving each LWP its own buffer, or sharing buffers between LWPs. Profiling data is updated at each clock tick in LWP user time. The profile state is inherited from the creating LWP.

Nonlocal Goto—`setjmp(3C)` and `longjmp(3C)`

The scope of `setjmp( )` and `longjmp( )` is limited to one thread, which is fine most of the time. However, this does mean that a thread that handles a signal can `longjmp( )` only when `setjmp( )` is performed in the same thread.

Resource Limits

Resource limits are set on the entire process and are determined by adding the resource use of all threads in the process. When a soft resource limit is exceeded, the offending thread is sent the appropriate signal. The sum of the resources used in the process is available through `getrusage(3B)`.

LWPs and Scheduling Classes

As mentioned in the “Scheduling” section of the “Chapter 1, Covering Multithreading Basics”, the Solaris pthreads implementation supports only the `SCHED_OTHER` scheduling policy. The others are optional under POSIX.

The POSIX `SCHED_FIFO` and `SCHED_RR` policies can be duplicated or emulated using the standard Solaris mechanisms. These scheduling mechanisms are described in this section.

The Solaris kernel has three classes of scheduling. The highest-priority scheduling class is Realtime (RT). The middle-priority scheduling class is `system`. The `system` class cannot be applied to a user process. The lowest-priority scheduling class is timeshare (TS), which is also the default class.

Scheduling class is maintained for each LWP. When a process is created, the initial LWP inherits the scheduling class and priority of the creating LWP in the parent process. As more LWPs are created to run unbound threads, they also inherit this scheduling class and priority.

All unbound threads in a process have the same scheduling class and priority. Each scheduling class maps the priority of the LWP it is scheduling to an overall dispatching priority according to the configurable priority of the scheduling class.

Bound threads have the scheduling class and priority of their underlying LWPs. Each bound thread in a process can have a unique scheduling class and priority that is visible to the kernel. Bound threads are scheduled with respect to all other LWPs in the system.

Thread priorities regulate access to LWP resources. By default LWPs are in the timesharing class. For compute-bound multithreading, thread priorities are not very useful. For multithreaded applications that do a lot of synchronization using the MT libraries, thread priorities become more meaningful.

The scheduling class is set by `prctl(2)`. How you specify the first two arguments determines whether just the calling LWP or all the LWPs of one or more processes are affected. The third argument of `prctl( )` is the command, which can be one of the following.

- `PC_GETCID`. Get the class ID and class attributes for a specific class.
- `PC_GETCLINFO`. Get the class name and class attributes for a specific class.
- `PC_GETPARMS`. Get the class identifier and the class-specific scheduling parameters of a process, an LWP with a process, or a group of processes.
- `PC_SETPARMS`. Set the class identifier and the class-specific scheduling parameters of a process, an LWP with a process, or a group of processes.

Use `prctl( )` only on bound threads. To affect the priority of unbound threads, use `pthread_setprio(3T)`.

Timeshare Scheduling

Timeshare scheduling distributes the processing resource fairly among the LWPs in this scheduling class. Other parts of the kernel can monopolize the processor for short intervals without degrading response time as seen by the user.

The `prctl(2)` call sets the `nice(2)` level of one or more processes. The `prctl( )` call also affects the `nice( )` level of all the timesharing class LWPs in the process. The `nice( )` level ranges from 0 to +20 normally and from -20 to +20 for processes with superuser privilege. The lower the value, the higher the priority.

The dispatch priority of time shared LWPs is calculated from the instantaneous CPU use rate of the LWP and from its `nice( )` level. The `nice( )` level indicates the relative priority of the LWPs to the timeshare scheduler.

LWPs with a greater `nice( )` value get a smaller, but nonzero, share of the total processing. An LWP that has received a larger amount of processing is given lower priority than one that has received little or no processing.

Realtime Scheduling

The Realtime class (RT) can be applied to a whole process or to one or more LWPs in a process. This requires superuser privilege.

Unlike the `nice(2)` level of the timeshare class, LWPs that are classified Realtime can be assigned priorities either individually or jointly. A `prioctl(2)` call affects the attributes of all the Realtime LWPs in the process.

The scheduler always dispatches the highest-priority Realtime LWP. It preempts a lower-priority LWP when a higher-priority LWP becomes runnable. A preempted LWP is placed at the head of its level queue.

A Realtime LWP retains control of a processor until it is preempted, it suspends, or its Realtime priority is changed. LWPs in the RT class have absolute priority over processes in the TS class.

A new LWP inherits the scheduling class of the parent process or LWP. An RT class LWP inherits the parent's time slice, whether finite or infinite.

An LWP with a finite time slice runs until it terminates, blocks (for example, to wait for an I/O event), is preempted by a higher-priority runnable Realtime process, or the time slice expires.

An LWP with an infinite time slice ceases execution only when it terminates, blocks, or is preempted.

LWP Scheduling and Thread Binding

The threads library automatically adjusts the number of LWPs in the pool used to run unbound threads. Its objectives are:

- To prevent the program from being blocked by a lack of unblocked LWPs.
For example, if there are more runnable unbound threads than LWPs and all the active threads block in the kernel in indefinite waits (such as while reading a tty), the process cannot progress until a waiting thread returns.
- To make efficient use of LWPs.
For example, if the library creates one LWP for each thread, many LWPs will usually be idle and the operating environment is overloaded by the resource requirements of the unused LWPs.

Keep in mind that LWPs are time sliced, not threads. This means that when there is only one LWP, there is no time slicing within the process—threads run on the LWP until they block (through interthread synchronization), are preempted, or terminate.

You can assign priorities to threads with `pthread_setprio(3T)`; lower-priority unbound threads are assigned to LWPs only when no higher-priority unbound threads are available. Bound threads, of course, do not compete for LWPs because

they have their own. Note that the thread priority that is set with `pthread_setprio( )` regulates threads access to LWPs, not to CPUs.

Bind threads to your LWPs to get precise control over whatever is being scheduled. This control is not possible when many unbound threads compete for an LWP.

In particular, a lower-priority unbound thread could be on a higher-priority LWP and running on a CPU, while a higher-priority unbound thread assigned to a lower-priority LWP is not running. In this sense, thread priorities are just a hint about access to CPUs.

Realtime threads are useful for getting a quick response to external stimuli. Consider a thread used for mouse tracking that must respond instantly to mouse clicks. By binding the thread to an LWP, you guarantee that there is an LWP available when it is needed. By assigning the LWP to the Realtime scheduling class, you ensure that the LWP is scheduled quickly in response to mouse clicks.

SIGWAITING—Creating LWPs for Waiting Threads

The library usually ensures that there are enough LWPs in its pool for a program to proceed.

When all the LWPs in the process are blocked in indefinite waits (such as blocked reading from a tty or network), the operating environment sends the new signal, `SIGWAITING`, to the process. This signal is handled by the threads library. When the process contains a thread that is waiting to run, a new LWP is created and the appropriate waiting thread is assigned to it for execution.

The `SIGWAITING` mechanism does not ensure that an additional LWP is created when one or more threads are compute bound and another thread becomes runnable. A compute-bound thread can prevent multiple runnable threads from being started because of a shortage of LWPs.

This can be prevented by calling `thr_setconcurrency(3THR)`. While using `thr_setconcurrency( )` with POSIX threads is not POSIX compliant, its use is recommended to avoid LWP shortages for unbound threads in some computationally intensive situations. (The only way to be *completely* POSIX compliant *and* avoid LWP shortages is to create only `PTHREAD_SCOPE_SYSTEM` bound threads.)

See “Thread Concurrency (Solaris Threads Only)” on page 235 for more information about using the `thr_setconcurrency(3T)` function.

In Solaris threads, you can also use `THR_NEW_LWP` in calls to `thr_create(3THR)` to create another LWP.

Aging LWPs

When the number of active threads is reduced, some of the LWPs in the pool are no longer needed. When there are more LWPs than active threads, the threads library destroys the unneeded LWPs. The library ages LWPs—they are deleted when they are unused for a “long” time; the default is five minutes.

Extending Traditional Signals

The traditional UNIX signal model is extended to threads in a fairly natural way. The key characteristics are that the signal disposition is process-wide, but the signal mask is per-thread. The process-wide disposition of signals is established using the traditional mechanisms (`signal(2)`, `sigaction(2)`, and so on).

When a signal handler is marked `SIG_DFL` or `SIG_IGN`, the action on receipt of the signal (exit, core dump, stop, continue, or ignore) is performed on the entire receiving process, affecting all threads in the process. For these signals that don’t have handlers, the issue of which thread picks the signal is unimportant, because the action on receipt of the signal is carried out on the whole process. See `signal(5)` for basic information about signals.

Each thread has its own signal mask. This lets a thread block some signals while it uses memory or another state that is also used by a signal handler. All threads in a process share the set of signal handlers set up by `sigaction(2)` and its variants.

A thread in one process cannot send a signal to a specific thread in another process. A signal sent by `kill(2)` or `sigsend(2)` to a process is handled by any one of the receptive threads in the process.

Unbound threads cannot use alternate signal stacks. A bound thread can use an alternate stack because the state is associated with the execution resource. An alternate stack must be enabled for the signal through `sigaction(2)`, and declared and enabled through `signalstack(2)`.

An application can have per-thread signal handlers based on the per-process signal handlers. One way is for the process-wide signal handler to use the identifier of the thread handling the signal as an index into a table of per-thread handlers. Note that there is no thread zero.

Signals are divided into two categories: traps and exceptions (synchronously generated signals) and interrupts (asynchronously generated signals).

As in traditional UNIX, if a signal is pending, additional occurrences of that signal have no additional effect—a pending signal is represented by a bit, not by a counter. In other words, signal delivery is idempotent.

As is the case with single-threaded processes, when a thread receives a signal while blocked in a system call, the thread might return early, either with the `EINTR` error code, or, in the case of I/O calls, with fewer bytes transferred than requested.

Of particular importance to multithreaded programs is the effect of signals on `pthread_cond_wait(3THR)`. This call usually returns in response to a `pthread_cond_signal(3THR)` or a `pthread_cond_broadcast(3THR)`, but, if the waiting thread receives a traditional UNIX signal, it returns with the error code `EINTR`. See “Interrupted Waits on Condition Variables (Solaris Threads Only)” on page 162 for more information.

Synchronous Signals

Traps (such as `SIGILL`, `SIGFPE`, `SIGSEGV`) result from something a thread does to itself, such as dividing by zero or explicitly sending itself a signal. A trap is handled only by the thread that caused it. Several threads in a process can generate and handle the same type of trap simultaneously.

Extending the idea of signals to individual threads is easy for synchronous signals—the signal is dealt with by the thread that caused the problem.

However, if the thread has not chosen to deal with the problem, such as by establishing a signal handler with `sigaction(2)`, the handler is invoked on the thread that receives the synchronous signal.

Because such a synchronous signal usually means that something is seriously wrong with the whole process, and not just with a thread, terminating the process is often a good choice.

Asynchronous Signals

Interrupts (such as `SIGINT` and `SIGIO`) are asynchronous with any thread and result from some action outside the process. They might be signals sent explicitly by other threads, or they might represent external actions such as a user typing `Control-C`. Dealing with asynchronous signals is more complicated than dealing with synchronous signals.

An interrupt can be handled by any thread whose signal mask allows it. When more than one thread is able to receive the interrupt, only one is chosen.

When multiple occurrences of the same signal are sent to a process, then each occurrence can be handled by a separate thread, as long as threads are available that do not have it masked. When all threads have the signal masked, then the signal is marked *pending* and the first thread to unmask the signal handles it.

Continuation Semantics

Continuation semantics are the traditional way to deal with signals. The idea is that when a signal handler returns, control resumes where it was at the time of the interruption. This is well suited for asynchronous signals in single-threaded processes, as shown in Code Example 5–1.

This is also used as the exception-handling mechanism in some programming languages, such as PL/1.

CODE EXAMPLE 5–1 Continuation Semantics

```
unsigned int nestcount;

unsigned int A(int i, int j) {
    nestcount++;

    if (i==0)
        return(j+1)
    else if (j==0)
        return(A(i-1, 1));
    else
        return(A(i-1, A(i, j-1)));
}

void sig(int i) {
    printf("nestcount = %d\n", nestcount);
}

main() {
    sigset(SIGINT, sig);
    A(4,4);
}
```

Operations on Signals

pthread_sigmask(3THR)

`pthread_sigmask(3THR)` does for a thread what `sigprocmask(2)` does for a process—it sets the thread's signal mask. When a new thread is created, its initial mask is inherited from its creator.

The call to `sigprocmask( )` in a multithreaded process is equivalent to a call to `pthread_sigsetmask( )`. See the `sigprocmask(2)` page for more information.

pthread_kill(3THR)

`pthread_kill(3THR)` is the thread analog of `kill(2)`—it sends a signal to a specific thread. This, of course, is different from sending a signal to a process. When a signal is sent to a process, the signal can be handled by any thread in the process. A signal sent by `pthread_kill( )` can be handled only by the specified thread.

Note that you can use `pthread_kill()` to send signals only to threads in the current process. This is because the thread identifier (type `thread_t`) is local in scope—it is not possible to name a thread in any process but your own.

Note also that the action taken (handler, `SIG_DFL`, `SIG_IGN`) on receipt of a signal by the target thread is global, as usual. This means, for example, that if you send `SIGXXX` to a thread, and the `SIGXXX` signal disposition for the process is to kill the process, then the whole process is killed when the target thread receives the signal.

`sigwait(2)`

For multithreaded programs, `sigwait(2)` is the preferred interface to use, because it deals so well with asynchronously generated signals.

`sigwait()` causes the calling thread to wait until any signal identified by its set argument is delivered to the thread. While the thread is waiting, signals identified by the set argument are unmasked, but the original mask is restored when the call returns.

All signals identified by the set argument must be blocked on all threads, including the calling thread; otherwise, `sigwait()` might not work correctly.

Use `sigwait()` to separate threads from asynchronous signals. You can create one thread that is listening for asynchronous signals while your other threads are created to block any asynchronous signals that might be set to this process.

New `sigwait()` Implementations

Two versions of `sigwait()` are available beginning with the Solaris Operating Environment 2.5 release: the new Solaris Operating Environment 2.5 version, and the POSIX.1c version. New applications and libraries should use the POSIX standard interface, as the Solaris Operating Environment version might not be available in future releases.

Note - The new Solaris Operating Environment 2.5 `sigwait()` does not override the signal's ignore disposition. Applications relying on the older `sigwait(2)` behavior can break unless you install a dummy signal handler to change the disposition from `SIG_IGN` to having a handler, so calls to `sigwait()` for this signal catch it.

The syntax for the two versions of `sigwait()` is shown below.

```
#include <signal.h>

/* the Solaris 2.5 version */
int sigwait(sigset_t *set);

/* the POSIX.1c version */
int sigwait(const sigset_t *set, int *sig);
```

When the signal is delivered, the POSIX.1c `sigwait()` clears the pending signal and places the signal number in `sig`. Many threads can call `sigwait()` at the same time, but only one thread returns for each signal that is received.

With `sigwait()` you can treat asynchronous signals synchronously—a thread that deals with such signals simply calls `sigwait()` and returns as soon as a signal arrives. By ensuring that all threads (including the caller of `sigwait()`) have such signals masked, you can be sure that signals are handled only by the intended handler and that they are handled safely.

By always masking all signals in all threads, and just calling `sigwait()` as necessary, your application will be much safer for threads that depend on signals.

Usually, you use `sigwait()` to create one or more threads that wait for signals. Because `sigwait()` can retrieve even masked signals, be sure to block the signals of interest in all other threads so they are not accidentally delivered.

When the signals arrive, a thread returns from `sigwait()`, handles the signal, and waits for more signals. The signal-handling thread is not restricted to using Async-Signal-Safe functions and can synchronize with other threads in the usual way. (The Async-Signal-Safe category is defined in “MT Interface Safety Levels” on page 170.)

Note - `sigwait()` should *never* be used with synchronous signals.

sigtimedwait(2)

`sigtimedwait(2)` is similar to `sigwait(2)` except that it fails and returns an error when a signal is not received in the indicated amount of time.

Thread-Directed Signals

The UNIX signal mechanism is extended with the idea of thread-directed signals. These are just like ordinary asynchronous signals, except that they are sent to a particular thread instead of to a process.

Waiting for asynchronous signals in a separate thread can be safer and easier than installing a signal handler and processing the signals there.

A better way to deal with asynchronous signals is to treat them synchronously. By calling `sigwait(2)`, discussed in “`sigwait(2)`” on page 157, a thread can wait until a signal occurs.

CODE EXAMPLE 5-2 Asynchronous Signals and `sigwait(2)`

```
main() {
    sigset_t set;
    void runA(void);
```

```

int sig;

sigemptyset(&set);
sigaddset(&set, SIGINT);
pthread_sigsetmask(SIG_BLOCK, &set, NULL);
pthread_create(NULL, 0, runA, NULL, PTHREAD_DETACHED, NULL);

while (1) {
    sigwait(&set, &sig);
    printf("nestcount = %d\n", nestcount);
    printf("received signal %d\n", sig);
}

}

void runA() {
    A(4,4);
    exit(0);
}

```

This example modifies the code of Code Example 5–1: the main routine masks the SIGINT signal, creates a child thread that calls the function A of the previous example, and then issues `sigwait()` to handle the SIGINT signal.

Note that the signal is masked in the compute thread because the compute thread inherits its signal mask from the main thread. The main thread is protected from SIGINT while, and only while, it is not blocked inside of `sigwait()`.

Also, note that there is never any danger of having system calls interrupted when you use `sigwait()`.

Completion Semantics

Another way to deal with signals is with completion semantics.

Use completion semantics when a signal indicates that something so catastrophic has happened that there is no reason to continue executing the current code block. The signal handler runs instead of the remainder of the block that had the problem. In other words, the signal handler completes the block.

In Code Example 5–3, the block in question is the body of the `then` part of the `if` statement. The call to `setjmp(3C)` saves the current register state of the program in *jbuf* and returns 0, thereby executing the block.

CODE EXAMPLE 5–3 Completion Semantics

```

sigjmp_buf jbuf;
void mult_divide(void) {
    int a, b, c, d;
    void problem();

    sigset(SIGFPE, problem);
    while (1) {
        if (sigsetjmp(&jbuf) == 0) {

```

```

        printf("Three numbers, please:\n");
        scanf("%d %d %d", &a, &b, &c);
        d = a*b/c;
        printf("%d*%d/%d = %d\n", a, b, c, d);
    }
}

void problem(int sig) {
    printf("Couldn't deal with them, try again\n");
    siglongjmp(&jbuf, 1);
}

```

If a SIGFPE (a floating-point exception) occurs, the signal handler is invoked.

The signal handler calls `siglongjmp(3C)`, which restores the register state saved in `jbuf`, causing the program to return from `sigsetjmp( )` again (among the registers saved are the program counter and the stack pointer).

This time, however, `sigsetjmp(3C)` returns the second argument of `siglongjmp( )`, which is 1. Notice that the block is skipped over, only to be executed during the next iteration of the `while` loop.

Note that you can use `sigsetjmp(3C)` and `siglongjmp(3C)` in multithreaded programs, but be careful that a thread never does a `siglongjmp( )` using the results of another thread's `sigsetjmp( )`.

Also, `sigsetjmp( )` and `siglongjmp( )` save and restore the signal mask, but `setjmp(3C)` and `longjmp(3C)` do not.

It is best to use `sigsetjmp( )` and `siglongjmp( )` when you work with signal handlers.

Completion semantics are often used to deal with exceptions. In particular, the Sun Ada™ programming language uses this model.

Note - Remember, `sigwait(2)` should *never* be used with synchronous signals.

Signal Handlers and Async-Signal Safety

A concept similar to thread safety is Async-Signal safety. Async-Signal-Safe operations are guaranteed not to interfere with operations that are being interrupted.

The problem of Async-Signal safety arises when the actions of a signal handler can interfere with the operation that is being interrupted.

For example, suppose a program is in the middle of a call to `printf(3S)` and a signal occurs whose handler itself calls `printf( )`. In this case, the output of the two `printf( )` statements would be intertwined. To avoid this, the handler should not call `printf( )` itself when `printf( )` might be interrupted by a signal.

This problem cannot be solved by using synchronization primitives because any attempted synchronization between the signal handler and the operation being synchronized would produce immediate deadlock.

Suppose that `printf()` is to protect itself by using a mutex. Now suppose that a thread that is in a call to `printf()`, and so holds the lock on the mutex, is interrupted by a signal.

If the handler (being called by the thread that is still inside of `printf()`) itself calls `printf()`, the thread that holds the lock on the mutex will attempt to take it again, resulting in an instant deadlock.

To avoid interference between the handler and the operation, either ensure that the situation never arises (perhaps by masking off signals at critical moments) or invoke only Async-Signal-Safe operations from inside signal handlers.

Because setting a thread's mask is an inexpensive user-level operation, you can inexpensively make functions or sections of code fit in the Async-Signal-Safe category.

The only routines that POSIX guarantees to be Async-Signal-Safe are listed in Table 5-2. Any signal handler can safely call in to one of these functions.

TABLE 5-2 Async-Signal-Safe Functions

<code>_exit()</code>	<code>fstat()</code>	<code>read()</code>	<code>sysconf()</code>
<code>access()</code>	<code>getegid()</code>	<code>rename()</code>	<code>tcdrain()</code>
<code>alarm()</code>	<code>geteuid()</code>	<code>rmdir()</code>	<code>tcflow()</code>
<code>cfgetispeed()</code>	<code>getgid()</code>	<code>setgid()</code>	<code>tcflush()</code>
<code>cfgetospeed()</code>	<code>getgroups()</code>	<code>setpgid()</code>	<code>tcgetattr()</code>
<code>cfsetispeed()</code>	<code>getpgrp()</code>	<code>setsid()</code>	<code>tcgetpgrp()</code>
<code>cfsetospeed()</code>	<code>getpid()</code>	<code>setuid()</code>	<code>tcsendbreak()</code>
<code>chdir()</code>	<code>getppid()</code>	<code>sigaction()</code>	<code>tcsetattr()</code>
<code>chmod()</code>	<code>getuid()</code>	<code>sigaddset()</code>	<code>tcsetpgrp()</code>
<code>chown()</code>	<code>kill()</code>	<code>sigdelset()</code>	<code>time()</code>

TABLE 5-2 Async-Signal-Safe Functions *(continued)*

<code>close()</code>	<code>link()</code>	<code>sigemptyset()</code>	<code>times()</code>
<code>creat()</code>	<code>lseek()</code>	<code>sigfillset()</code>	<code>umask()</code>
<code>dup2()</code>	<code>mkdir()</code>	<code>sigismember()</code>	<code>uname()</code>
<code>dup()</code>	<code>mkfifo()</code>	<code>sigpending()</code>	<code>unlink()</code>
<code>execle()</code>	<code>open()</code>	<code>sigprocmask()</code>	<code>utime()</code>
<code>execve()</code>	<code>pathconf()</code>	<code>sigsuspend()</code>	<code>wait()</code>
<code>fcntl()</code>	<code>pause()</code>	<code>sleep()</code>	<code>waitpid()</code>
<code>fork()</code>	<code>pipe()</code>	<code>stat()</code>	<code>write()</code>

Interrupted Waits on Condition Variables (Solaris Threads Only)

When a signal is delivered to a thread while the thread is waiting on a condition variable, the old convention (assuming that the process is not terminated) is that interrupted calls return `EINTR`.

The ideal new condition would be that when `cond_wait(3T)` and `cond_timedwait(3T)` return, the lock has been retaken on the mutex.

This is what is done in Solaris threads: when a thread is blocked in `cond_wait()` or `cond_timedwait()` and an unmasked, caught signal is delivered to the thread, the handler is invoked and the call to `cond_wait()` or `cond_timedwait()` returns `EINTR` with the mutex locked.

This implies that the mutex is locked in the signal handler because the handler might have to clean up after the thread. While this is true in the Solaris Operating Environment 2.5 release, it might change in the future, so do not rely on this behavior.

Note - In POSIX threads, `pthread_cond_wait(3T)` returns from signals, but this is not an error, `pthread_cond_wait( )` returns zero as a spurious wake-up.

Handler cleanup is illustrated by Code Example 5-4.

CODE EXAMPLE 5-4 Condition Variables and Interrupted Waits

```
int sig_catcher() {
    sigset_t set;
    void hdlr();

    mutex_lock(&mut);

    sigemptyset(&set);
    sigaddset(&set, SIGINT);
    sigsetmask(SIG_UNBLOCK, &set, 0);

    if (cond_wait(&cond, &mut) == EINTR) {
        /* signal occurred and lock is held */
        cleanup();
        mutex_unlock(&mut);
        return(0);
    }
    normal_processing();
    mutex_unlock(&mut);
    return(1);
}

void hdlr() {
    /* lock is held in the handler */
    ...
}
```

Assume that the `SIGINT` signal is blocked in all threads on entry to `sig_catcher( )` and that `hdlr( )` has been established (with a call to `sigaction(2)`) as the handler for the `SIGINT` signal. When an unmasked and caught instance of the `SIGINT` signal is delivered to the thread while it is in `cond_wait( )`, the thread first reacquires the lock on the mutex, then calls `hdlr( )`, and then returns `EINTR` from `cond_wait( )`.

Note that whether `SA_RESTART` has been specified as a flag to `sigaction( )` has no effect here; `cond_wait(3T)` is not a system call and is not automatically restarted. When a caught signal occurs while a thread is blocked in `cond_wait( )`, the call always returns `EINTR`. Again, the application should not rely on an interrupted `cond_wait( )` reacquiring the mutex, because this behavior could change in the future.

I/O Issues

One of the attractions of multithreaded programming is I/O performance. The traditional UNIX API gave you little assistance in this area—you either used the facilities of the file system or bypassed the file system entirely.

This section shows how to use threads to get more flexibility through I/O concurrency and multibuffering. This section also discusses the differences and similarities between the approaches of synchronous I/O (with threads) and asynchronous I/O (with and without threads).

I/O as a Remote Procedure Call

In the traditional UNIX model, I/O appears to be synchronous, as if you were placing a remote procedure call to the I/O device. Once the call returns, then the I/O has completed (or at least it appears to have completed—a write request, for example, might merely result in the transfer of the data to a buffer in the operating environment).

The advantage of this model is that it is easy to understand because, as a programmer you are very familiar with the concept of procedure calls.

An alternative approach not found in traditional UNIX systems is the asynchronous model, in which an I/O request merely starts an operation. The program must somehow discover when the operation completes.

This approach is not as simple as the synchronous model, but it has the advantage of allowing concurrent I/O and processing in traditional, single-threaded UNIX processes.

Tamed Asynchrony

You can get most of the benefits of asynchronous I/O by using synchronous I/O in a multithreaded program. Where, with asynchronous I/O, you would issue a request and check later to determine when it completes, you can instead have a separate thread perform the I/O synchronously. The main thread can then check (perhaps by calling `pthread_join(3T)`) for the completion of the operation at some later time.

Asynchronous I/O

In most situations there is no need for asynchronous I/O, since its effects can be achieved with the use of threads, with each thread doing synchronous I/O. However, in a few situations, threads cannot achieve what asynchronous I/O can.

The most straightforward example is writing to a tape drive to make the tape drive stream. Streaming prevents the tape drive from stopping while it is being written to and moves the tape forward at high speed while supplying a constant stream of data that is written to tape.

To do this, the tape driver in the kernel must issue a queued write request when the tape driver responds to an interrupt that indicates that the previous tape-write operation has completed.

Threads cannot guarantee that asynchronous writes will be ordered because the order in which threads execute is indeterminate. Specifying the order of a write to a tape, for example, is not possible.

Asynchronous I/O Operations

```
#include <sys/asynch.h>

int aioread(int files, char *bufp, int bufs, off_t offset,
            int whence, aio_result_t *resultp);

int aiowrite(int files, const char *bufp, int bufs,
             off_t offset, int whence, aio_result_t *resultp);

aio_result_t *aiowait(const struct timeval *timeout);

int aiocancel(aio_result_t *resultp);
```

`aioread(3)` and `aiowrite(3)` are similar in form to `pread(2)` and `pwrite(2)`, except for the addition of the last argument. Calls to `aioread( )` and `aiowrite( )` result in the initiation (or queueing) of an I/O operation.

The call returns without blocking, and the status of the call is returned in the structure pointed to by *resultp*. This is an item of type `aio_result_t` that contains the following:

```
int aio_return;
int aio_errno;
```

When a call fails immediately, the failure code can be found in `aio_errno`. Otherwise, this field contains `AIO_INPROGRESS`, meaning that the operation has been successfully queued.

You can wait for an outstanding asynchronous I/O operation to complete by calling `aiowait(3)`. This returns a pointer to the `aio_result_t` structure supplied with the original `aioread(3)` or `aiowrite(3)` call.

This time `aio_result_t` contains whatever `read(2)` or `write(2)` would have returned if one of them had been called instead of the asynchronous version. If the `read( )` or `write( )` is successful, `aio_return` contains the number of bytes that were read or written; if it was not successful, `aio_return` is -1, and `aio_errno` contains the error code.

`aiowait( )` takes a *timeout* argument, which indicates how long the caller is willing to wait. As usual, a NULL pointer here means that the caller is willing to wait indefinitely, and a pointer to a structure containing a zero value means that the caller is unwilling to wait at all.

You might start an asynchronous I/O operation, do some work, then call `aiowait( )` to wait for the request to complete. Or you can use `SIGIO` to be notified, asynchronously, when the operation completes.

Finally, a pending asynchronous I/O operation can be cancelled by calling `aiocancel( )`. This routine is called with the address of the result area as an argument. This result area identifies which operation is being cancelled.

Shared I/O and New I/O System Calls

When multiple threads are performing I/O operations at the same time with the same file descriptor, you might discover that the traditional UNIX I/O interface is not thread safe. The problem occurs with nonsequential I/O. This uses the `lseek(2)` system call to set the file offset, which is then used in the next `read(2)` or `write(2)` call to indicate where in the file the operation should start. When two or more threads are issuing `lseek( )` to the same file descriptor, a conflict results.

To avoid this conflict, use the `pread(2)` and `pwrite(2)` system calls.

```
#include <sys/types.h>
#include <unistd.h>

ssize_t pread(int fildes, void *buf, size_t nbyte, off_t offset);

ssize_t pwrite(int fildes, void *buf, size_t nbyte,
               off_t offset);
```

These behave just like `read(2)` and `write(2)` except that they take an additional argument, the file offset. With this argument, you specify the offset without using `lseek(2)`, so multiple threads can use these routines safely for I/O on the same file descriptor.

Alternatives to `getc(3S)` and `putc(3S)`

An additional problem occurs with standard I/O. Programmers are accustomed to routines such as `getc(3S)` and `putc(3S)` being very quick—they are implemented as macros. Because of this, they can be used within the inner loop of a program with no concerns about efficiency.

However, when they are made thread safe they suddenly become more expensive—they now require (at least) two internal subroutine calls, to lock and unlock a mutex.

To get around this problem, alternative versions of these routines are supplied, `getc_unlocked(3S)` and `putc_unlocked(3S)`.

These do not acquire locks on a mutex and so are as quick as the original, nonthread-safe versions of `getc(3S)` and `putc(3S)`.

However, to use them in a thread-safe way, you must explicitly lock and release the mutexes that protect the standard I/O streams, using `flockfile(3S)` and `funlockfile(3S)`. The calls to these latter routines are placed outside the loop, and the calls to `getc_unlocked( )` or `putc_unlocked( )` are placed inside the loop.

Safe and Unsafe Interfaces

This chapter defines MT-safety levels for functions and libraries.

- “Thread Safety” on page 169
- “MT Interface Safety Levels” on page 170
- “Async-Signal-Safe Functions” on page 173
- “MT Safety Levels for Libraries” on page 173

Thread Safety

Thread safety is the avoidance of data races—situations in which data are set to either correct or incorrect values, depending upon the order in which multiple threads access and modify the data.

When no sharing is intended, give each thread a private copy of the data. When sharing is important, provide explicit synchronization to make certain that the program behaves in a deterministic manner.

A procedure is thread safe when it is logically correct when executed simultaneously by several threads. At a practical level, it is convenient to recognize three levels of safety.

- Unsafe
- Thread safe - Serializable
- Thread safe - MT-Safe

An unsafe procedure can be made thread safe and serializable by surrounding it with statements to lock and unlock a mutex. Code Example 6–1 shows three simplified implementations of `fputs( )`, initially thread unsafe.

Next is a serializable version of this routine with a single mutex protecting the procedure from concurrent execution problems. Actually, this is stronger synchronization than is usually necessary. When two threads are sending output to different files using `fputs()`, one need not wait for the other—the threads need synchronization only when they are sharing an output file.

The last version is MT-safe. It uses one lock for each file, allowing two threads to print to different files at the same time. So, a routine is MT-safe when it is thread safe and its execution does not negatively affect performance.

CODE EXAMPLE 6-1 Degrees of Thread Safety

```
/* not thread-safe */
fputs(const char *s, FILE *stream) {
    char *p;
    for (p=s; *p; p++)
        putc((int)*p, stream);
}

/* serializable */
fputs(const char *s, FILE *stream) {
    static mutex_t mut;
    char *p;
    mutex_lock(&mut);
    for (p=s; *p; p++)
        putc((int)*p, stream);

    mutex_unlock(&mut);
}

/* MT-Safe */
mutex_t m[NFILE];
fputs(const char *s, FILE *stream) {
    static mutex_t mut;
    char *p;
    mutex_lock(&m[fileno(stream)]);
    for (p=s; *p; p++)
        putc((int)*p, stream);
    mutex_unlock(&m[fileno(stream)]);
}
```

MT Interface Safety Levels

The threads man pages, `man(3t)`, use the safety level categories listed in Table 6-1 to describe how well an interface supports threads (these categories are explained more fully in the `Intro(3)` reference manual page).

TABLE 6-1 Interface Safety Levels

Category	Description
Safe	This code can be called from a multithreaded application
Safe with exceptions	See the NOTES sections of these pages for a description of the exceptions.
Unsafe	This interface is not safe to use with multithreaded applications unless the application arranges for only one thread at a time to execute within the library.
MT-Safe	This interface is fully prepared for multithreaded access in that it is both <i>safe</i> and it supports some concurrency.
MT-Safe with exceptions	See the NOTES sections of these pages in the man(3t): Library Routines for a list of the exceptions.
Async-Signal-Safe	This routine can safely be called from a signal handler. A thread that is executing an Async-Signal-Safe routine does not deadlock with itself when it is interrupted by a signal.
Fork1-Safe	This interface releases locks it has held whenever the Solaris fork1(2) or the POSIX fork(2) is called.

See the table in Appendix C, MT Safety Levels: Library Interfaces, for the safety levels of interfaces from the man(3): Library Routines. Check the man page to be sure of the level.

Some functions have purposely not been made safe for the following reasons.

- Making the interface MT-Safe would have negatively affected the performance of single-threaded applications.
- The library has an Unsafe interface. For example, a function might return a pointer to a buffer in the stack. You can use reentrant counterparts for some of these functions. The reentrant function name is the original function name with “_r” appended.



Caution - There is no way to be certain that a function with a name not ending in “_r” is MT-Safe other than by checking its reference manual page. Use of a function identified as not MT-Safe must be protected by a synchronizing device or by restriction to the initial thread.

Reentrant Functions for Unsafe Interfaces

For most functions with Unsafe interfaces, an MT-Safe version of the routine exists. The name of the new MT-Safe routine is always the name of the old Unsafe routine with “_r” appended. The Table 6-2 “_r” routines are supplied in the Solaris environment.

TABLE 6-2 Reentrant Functions

asctime_r(3c)	gethostbyname_r(3n)	getservbyname_r(3n)
ctermid_r(3s)	gethostent_r(3n)	getservbyport_r(3n)
ctime_r(3c)	getlogin_r(3c)	getservent_r(3n)
fgetgrent_r(3c)	getnetbyaddr_r(3n)	getspent_r(3c)
fgetpwent_r(3c)	getnetbyname_r(3n)	getspnam_r(3c)
fgetspent_r(3c)	getnetent_r(3n)	gmtime_r(3c)
gamma_r(3m)	getnetgrent_r(3n)	lgamma_r(3m)
getaaclassent_r(3)	getprotobyname_r(3n)	localtime_r(3c)
getaaclassnam_r(3)	getprotobynumber_r(3n)	nis_sperror_r(3n)
getauevent_r(3)	getprotoent_r(3n)	rand_r(3c)
getauevnam_r(3)	getpwent_r(3c)	readdir_r(3c)
getauevnum_r(3)	getpwnam_r(3c)	strtok_r(3c)
getgrent_r(3c)	getpwuid_r(3c)	tmpnam_r(3s)
getgrgid_r(3c)	getrpcbyname_r(3n)	ttynam_r(3c)

TABLE 6-2 Reentrant Functions *(continued)*

getgrnam_r(3c)	getrpcbynumber_r(3n)
gethostbyaddr_r(3n)	getrpccent_r(3n)

Async-Signal-Safe Functions

Functions that can safely be called from signal handlers are *Async-Signal-Safe*. The POSIX standard defines and lists Async-Signal-Safe functions (IEEE Std 1003.1-1990, 3.3.1.3 (3)(f), page 55). In addition to the POSIX Async-Signal-Safe functions, these three functions from the Solaris threads library are also Async-Signal-Safe.

- `sema_post(3T)`
- `thr_sigsetmask(3T)`, similar to `pthread_sigmask(3T)`
- `thr_kill(3T)`, similar to `pthread_kill(3T)`

MT Safety Levels for Libraries

All routines that can potentially be called by a thread from a multithreaded program should be MT-Safe.

This means that two or more activations of a routine must be able to *correctly* execute concurrently. So, every library interface that a multithreaded program uses must be MT-Safe.

Not all libraries are now MT-Safe. The commonly used libraries that are MT-Safe are listed in Table 6-3. Additional libraries will eventually be modified to be MT-Safe.

TABLE 6-3 Some MT-Safe Libraries

Library	Comments
<code>lib/libc</code>	Interfaces that are not safe have thread-safe interfaces of the form <code>*_r</code> (often with different semantics)
<code>lib/libdl_stubs</code>	To support static switch compiling

TABLE 6-3 Some MT-Safe Libraries *(continued)*

Library	Comments
lib/libintl	Internationalization library
lib/libm	Math library compliant with System V Interface Definition, Edition 3, X/Open and ANSI C
lib/libmalloc	Space-efficient memory allocation library; see malloc(3X)
lib/libmapmalloc	Alternative mmap(2)-based memory allocation library; see mapmalloc(3X)
lib/libnsl	The TLI interface, XDR, RPC clients and servers, netdir, netselect and getXXbyYY interfaces are not safe, but have thread-safe interfaces of the form getXXbyYY_r
lib/libresolv	Thread-specific errno support
lib/libsocket	Socket library for making network connections
lib/libw	Wide character and wide string functions for supporting multibyte locales
lib/straddr	Network name-to-address translation library
lib/libX11	X11 Windows library routines
lib/libC	C++ runtime shared objects

Unsafe Libraries

Routines in libraries that are not guaranteed to be MT-Safe can safely be called by multithreaded programs only when such calls are single threaded.

Compiling and Debugging

This chapter describes how to compile and debug multithreaded programs.

- “Compiling a Multithreaded Application” on page 175
- “Debugging a Multithreaded Program” on page 181

Compiling a Multithreaded Application

Many options are available for header files, define flags, and linking.

Preparing for Compilation

The following items are required to compile and link a multithreaded program. Except for the C compiler, all should come with your Solaris operating environment.

- A standard C compiler
- Include files:
 - `<thread.h>` and `<pthread.h>`
 - `<errno.h>`, `<limits.h>`, `<signal.h>`, `<unistd.h>`
- The regular Solaris linker, `ld(1)`
- The Solaris threads library (`libthread`), the POSIX threads library (`libpthread`), and possibly the POSIX realtime library (`librt`) for semaphores
- MT-safe libraries (`libc`, `libm`, `libw`, `libintl`, `libnsl`, `libsocket`, `libmalloc`, `libmapmalloc`, and so on)

Choosing Solaris or POSIX Semantics

Certain functions, including the ones listed in Table 7–1, have different semantics in the POSIX 1003.1c standard than in the Solaris 2.4 Operating Environment release, which was based on an earlier POSIX draft. Function definitions are chosen at compile time. See the man Pages(3): Library Routines for a description of the differences in expected parameters and return values.

TABLE 7–1 Functions With POSIX/Solaris Semantic Differences

sigwait(2)	
ctime_r(3C)	asctime_r(3C)
ftrylockfile(3S) - new	getlogin_r(3C)
getgrnam_r(3C)	getgrgid_r(3C)
getpwnam_r(3C)	getpwuid_r(3C)
readdir_r(3C)	ttyname_r(3C)

The Solaris `fork(2)` function duplicates all threads (*fork-all* behavior), while the POSIX `fork(2)` function duplicates only the calling thread (*fork-one* behavior), as does the Solaris `fork1( )` function.

The handling of an `alarm(2)` is also different: a Solaris alarm goes to the thread's LWP, while a POSIX alarm goes to the whole process (see “Per-Thread Alarms” on page 149).

Including `<thread.h>` or `<pthread.h>`

The include file `<thread.h>`, used with the `-lthread` library, compiles code that is upward compatible with earlier releases of the Solaris Operating Environment. This library contains both interfaces—those with Solaris semantics and those with POSIX semantics. To call `thr_setconcurrency(3T)` with POSIX threads, your program needs to include `<thread.h>`.

The include file `<pthread.h>`, used with the `-lpthread` library, compiles code that is conformant with the multithreading interfaces defined by the POSIX 1003.1c standard. For complete POSIX compliance, the define flag `_POSIX_C_SOURCE` should be set to a (long) value ≥ 199506 :

```
cc [flags] file... -D_POSIX_C_SOURCE=N (where N 199506L)
```

You can mix Solaris threads and POSIX threads in the same application, by including both `<thread.h>` and `<pthread.h>`, and linking with either the `-lthread` or `-lpthread` library.

In mixed use, Solaris semantics prevail when compiling with `-D_REENTRANT` and linking with `-lthread`, whereas POSIX semantics prevail when compiling with `-D_POSIX_C_SOURCE` and linking with `-lpthread`.

Defining `_REENTRANT` or `_POSIX_C_SOURCE`

For POSIX behavior, compile applications with the `-D_POSIX_C_SOURCE` flag set $\geq 199506L$. For Solaris behavior, compile multithreaded programs with the `-D_REENTRANT` flag. This applies to every module of an application.

For mixed applications (for example, Solaris threads with POSIX semantics), compile with the `-D_REENTRANT` and `-D_POSIX_PTHREAD_SEMANTICS` flags.

To compile a single-threaded application, define neither the `-D_REENTRANT` nor the `-D_POSIX_C_SOURCE` flag. When these flags are not present, all the old definitions for `errno`, `stdio`, and so on, remain in effect.

Note - Compile single-threaded applications, not linked with either of the thread libraries (`libthread.so.1` or `libpthread.so.1`), without the `-D_REENTRANT` flag. This eliminates performance degradation incurred when macros, such as `putc(3s)`, are converted into reentrant function calls.

To summarize, POSIX applications that define `-D_POSIX_C_SOURCE` get the POSIX 1003.1c semantics for the routines listed in Table 7-1. Applications that define only `-D_REENTRANT` get the Solaris semantics for these routines. Solaris applications that define `-D_POSIX_PTHREAD_SEMANTICS` get the POSIX semantics for these routines, but can still use the Solaris threads interface.

Applications that define both `-D_POSIX_C_SOURCE` and `-D_REENTRANT` get the POSIX semantics.

Linking With `libthread` or `libpthread`

For POSIX threads behavior, load the `libpthread` library. For Solaris threads behavior, load the `libthread` library. Some POSIX programmers might want to link with `-lthread` to preserve the Solaris distinction between `fork()` and `fork1()`. All that `-lpthread` really does is to make `fork()` behave the same way as the Solaris `fork1()` call, and change the behavior of `alarm(2)`.

To use `libthread`, specify `-lthread` before `-lc` on the `ld` command line, or last on the `CC` command line.

To use `libpthread`, specify `-lpthread` before `-lc` on the `ld` command line, or last on the `CC` command line.

Do not link a nonthreaded program with `-lthread` or `-lpthread`. Doing so establishes multithreading mechanisms at link time that are initiated at runtime. These slow down a single-threaded application, waste system resources, and produce misleading results when you debug your code.

Figure 7-1 summarizes the compile options.




Figure 7-1 Compilation Flowchart

In mixed usage, you need to include both `thread.h` and `pthread.h`.

All calls to `libthread` and `libpthread` are no-ops if the application does not link `-lthread` or `-lpthread`. The runtime library `libc` has many predefined `libthread` and `libpthread` stubs that are null procedures. True procedures are interposed by `libthread` or `libpthread` when the application links both `libc` and the thread library.

The behavior of the C library is undefined if a program is constructed with an `ld` command line that includes the following *incorrect* fragment:

```
.o's ... -lc -lthread ... (this is incorrect)
or
.o's ... -lc -lpthread ... (this is incorrect)
```

Note - For C++ programs that use threads, use the `-mt` option, rather than `-lthread`, to compile and link your application. The `-mt` option links with `libthread` and ensures proper library linking order. Using `-lthread` might cause your program to core dump.

Linking With `-lposix4` for POSIX Semaphores

The Solaris semaphore routines, `sem_*(3T)`, are contained in the `libthread` library. By contrast, you link with the `-lposix4` library to get the standard `sem_*(3R)` POSIX 1003.1c semaphore routines described in “Semaphores” on page 122.

Link Old With New

Table 7-2 shows that multithreaded object modules should be linked with old object modules only with great caution.

TABLE 7-2 Compiling With and Without the `_REENTRANT` Flag

The File Type	Compiled	Reference	And Return
Old object files (nonthreaded) and new object files	<i>Without</i> the <code>_REENTRANT</code> or <code>_POSIX_C_SOURCE</code> flag	Static storage	The traditional <i>errno</i>
New object files	<i>With</i> the <code>_REENTRANT</code> or <code>_POSIX_C_SOURCE</code> flag	<code>__errno</code> , the new binary entry point	The address of the thread's definition of <i>errno</i>
Programs using TLI in <code>libnsl</code> ¹	<i>With</i> the <code>_REENTRANT</code> or <code>_POSIX_C_SOURCE</code> flag (required)	<code>__t_errno</code> , a new entry point	The address of the thread's definition of <i>t_errno</i> .

1. Include `tiuser.h` to get the TLI global error variable.

Linking With the Alternate One-level libthread Library

The standard Solaris threads implementation is built upon a two-level threading model in which user-level threads are multiplexed over possibly fewer lightweight processes, or LWPs. An LWP is the fundamental unit of execution that is dispatched to a processor by the operating system. A mechanism is provided in the standard implementation for associating threads one-to-one with LWPs (the `THR_BOUND` and `PTHREAD_SCOPE_SYSTEM` flags) for writing applications with one-level semantics.

The Solaris 8 Operating Environment provides an alternate threads implementation of a one-level model in which user-level threads are associated one-to-one with LWPs. This implementation is simpler than the standard implementation and may be beneficial to some multithreaded applications. It provides exactly the same interfaces, both for POSIX threads and Solaris threads, as the standard implementation.

To link with the alternate implementation, use the following run-path `-R` option when linking the program.

For POSIX threads use:

```
cc -mt ... -lpthread ... -R /usr/lib/lwp      (32-bit)
cc -mt ... -lpthread ... -R /usr/lib/lwp/64   (64-bit)
```

For Solaris threads use:

```
cc -mt ... -R /usr/lib/lwp      (32-bit)
cc -mt ... -R /usr/lib/lwp/64   (64-bit)
```

For multithreaded programs that have been previously linked with the standard threads library, the environment variable `LD_LIBRARY_PATH` and `LD_LIBRARY_PATH_64` can be set as follows to bind the program at run time to the alternate threads library:

```
LD_LIBRARY_PATH=/usr/lib/lwp
LD_LIBRARY_PATH_64=/usr/lib/lwp:/usr/lib/lwp/64
```

Note that if an `LD_LIBRARY_PATH` environment variable is in effect for a secure process, then only the trusted directories specified by this variable will be used to augment the runtime linker's search rules.

When using the alternate one-level threads implementation, the library might create more LWPs than the standard implementation using unbound threads. LWPs consume operating system memory in contrast to threads, which consume only user-level memory. Thus, a multithreaded application linked with this library that creates thousands of threads would create an equal number of LWPs and might run the system out of resources required to support the application.

Debugging a Multithreaded Program

Common Oversights

The following list points out some of the more frequent oversights that can cause bugs in multithreaded programs.

- Passing a pointer to the caller's stack as an argument to a new thread.
- Accessing global memory (shared changeable state) without the protection of a synchronization mechanism.
- Creating deadlocks caused by two threads trying to acquire rights to the same pair of global resources in alternate order (so that one thread controls the first resource and the other controls the second resource and neither can proceed until the other gives up).
- Trying to reacquire a lock already held (recursive deadlock).
- Creating a hidden gap in synchronization protection. This is caused when a code segment protected by a synchronization mechanism contains a call to a function that frees and then reacquires the synchronization mechanism before it returns to the caller. The result is that it appears to the caller that the global data has been protected when it actually has not.
- Mixing UNIX signals with threads—it is better to use the `sigwait(2)` model for handling asynchronous signals.
- Using `setjmp(3B)` and `longjmp(3B)`, and then long-jumping away without releasing the mutex locks.
- Failing to reevaluate the conditions after returning from a call to `*_cond_wait(3T)` or `*_cond_timedwait(3T)`.
- Forgetting that default threads are created `PTHREAD_CREATE_JOINABLE` and must be reclaimed with `pthread_join(3T)`; note, `pthread_exit(3T)` does not free up its storage space.

- Making deeply nested, recursive calls and using large automatic arrays can cause problems because multithreaded programs have a more limited stack size than single-threaded programs.
- Specifying an inadequate stack size, or using non-default stacks.

And, note that multithreaded programs (especially those containing bugs) often behave differently in two successive runs, given identical inputs, because of differences in the thread scheduling order.

In general, multithreading bugs are statistical instead of deterministic. Tracing is usually a more effective method of finding order of execution problems than is breakpoint-based debugging.

Tracing and Debugging With the TNF Utilities

Use the TNF utilities (included as part of the Solaris system) to trace, debug, and gather performance analysis information from your applications and libraries. The TNF utilities integrate trace information from the kernel and from multiple user processes and threads, and so are especially useful for multithreaded code.

With the TNF utilities, you can easily trace and debug multithreaded programs. See the TNF utilities chapter in the *Programming Utilities Guide* for detailed information on using `prex(1)`, `tnfdump(1)`, and other TNF utilities.

Using `truss(1)`

See `truss(1)` for information on tracing system calls and signals.

Using `adb(1)`

When you bind all threads in a multithreaded program, a thread and an LWP are synonymous. Then you can access each thread with the following `adb` commands that support multithreaded programming.

TABLE 7-3 MT `adb` Commands

<code>pid:A</code>	Attaches to process # <i>pid</i> . This stops the process and all its LWPs.
<code>:R</code>	Detaches from process. This resumes the process and all its LWPs.
<code>\$L</code>	Lists all active LWPs in the (stopped) process.

TABLE 7-3 MT adb Commands *(continued)*

n:l	Switches focus to LWP # <i>n</i> .
\$l	Shows the LWP currently focused.
num:i	Ignores signal number <i>num</i> .

These commands to set conditional breakpoints are often useful.

TABLE 7-4 Setting adb Breakpoints

[label], [count]:b [expression]	Breakpoint is detected when <i>expression</i> equals zero
foo,ffff:b <g7-0xabcdef	Stop at <i>foo</i> when <i>g7</i> = the hex value 0xABCDEF

Using dbx

With the dbx utility you can debug and execute source programs written in C++, ANSI C, and FORTRAN. dbx accepts the same commands as the Debugger, but uses a standard terminal (TTY) interface. Both dbx and the Debugger support debugging multithreaded programs. For a full overview of dbx and Debugger features see the dbx(1) reference manual page and the *Using Sun Workshop* user's guide.

All the dbx options listed in Table 7-5 can support multithreaded applications.

TABLE 7-5 dbx Options for MT Programs

Option	Meaning
cont at line [sig signo id]	Continues execution at <i>line</i> with signal <i>signo</i> . The <i>id</i> , if present, specifies which thread or LWP to continue. The default value is <i>all</i> .
lwp	Displays current LWP. Switches to given LWP [lwpid].
lwps	Lists all LWPs in the current process.

TABLE 7-5 dbx Options for MT Programs *(continued)*

Option	Meaning
next ... tid	Steps the given thread. When a function call is skipped, all LWPs are implicitly resumed for the duration of that function call. Nonactive threads cannot be stepped.
next ... lid	Steps the given LWP. Does not implicitly resume all LWPs when skipping a function. The LWP on which the given thread is active. Does not implicitly resume all LWP when skipping a function.
step... tid	Steps the given thread. When a function call is skipped, all LWPs are implicitly resumed for the duration of that function call. Nonactive threads cannot be stepped.
step... lid	Steps the given LWP. Does not implicitly resume all LWPs when skipping a function.
stepi... lid	The given LWP.
stepi... tid	The LWP on which the given thread is active.
thread	Displays current thread. Switches to thread <i>tid</i> . In all the following variations, an optional tid implies the current thread.
thread -info [tid]	Prints everything known about the given thread.
thread -locks [tid]	Prints all locks held by the given thread.
thread -suspend [tid]	Puts the given thread into suspended state.
thread -continue [tid]	Unsusends the given thread.
thread -hide [tid]	<i>Hides</i> the given (or current) thread. It will not appear in the generic threads listing.
thread -unhide [tid]	<i>Unhides</i> the given (or current) thread.
allthread-unhide	<i>Unhides</i> all threads.
threads	Prints the list of all known threads.

TABLE 7-5 dbx Options for MT Programs *(continued)*

Option	Meaning
threads-all	Prints threads that are not usually printed (zombies).
all filterthreads-mode	Controls whether threads prints all threads or filters them by default.
auto manualthreads-mode	Enables automatic updating of the thread listing.
threads-mode	Echoes the current modes. Any of the previous forms can be followed by a thread or LWP ID to get the traceback for the specified entity.

Programming With Solaris Threads

This chapter compares the Application Programming Interface (API) for Solaris and POSIX threads, and explains the Solaris features that are not found in POSIX threads.

- “Comparing APIs for Solaris Threads and POSIX Threads” on page 187
- “Unique Solaris Threads Functions” on page 194
- “Similar Synchronization Functions—Read-Write Locks” on page 198
- “Similar Solaris Threads Functions” on page 204
- “Similar Synchronization Functions—Mutual Exclusion Locks” on page 213
- “Similar Synchronization Functions—Condition Variables” on page 216
- “Similar Synchronization Functions—Semaphores” on page 219
- “Special Issues for fork() and Solaris Threads” on page 224

Comparing APIs for Solaris Threads and POSIX Threads

The Solaris threads API and the pthreads API are two solutions to the same problem: building parallelism into application software. Although each API is complete in itself, you can safely mix Solaris threads functions and pthread functions in the same program.

The two APIs do not match exactly, however. Solaris threads supports functions that are not found in pthreads, and pthreads includes functions that are not supported in the Solaris interface. For those functions that *do* match, the associated arguments might not, although the information content is effectively the same.

By combining the two APIs, you can use features not found in one to enhance the other. Similarly, you can run applications using Solaris threads, exclusively, with applications using pthreads, exclusively, on the same system.

Major API Differences

Solaris threads and pthreads are very similar in both API action and syntax. The major differences are listed in Table 8–1.

TABLE 8–1 Unique Solaris Threads and pthreads Features

Solaris Threads (libthread)	POSIX Threads (libpthread)
thr_ prefix for threads function names; sema_ prefix for semaphore function names	pthread_ prefix for pthreads function names; sem_ prefix for semaphore function names
Readers/Writer locks	Attribute objects (these replace many Solaris arguments or flags with pointers to pthreads attribute objects)
Ability to create “daemon” threads	Cancellation semantics
Suspending and continuing a thread	Scheduling policies
Setting concurrency (requesting a new LWP); determining concurrency level	

Function Comparison Table

The following table compares Solaris threads functions with pthreads functions. Note that even when Solaris threads and pthreads functions appear to be essentially the same, the arguments to the functions can differ.

When a comparable interface is not available either in pthreads or Solaris threads, a hyphen ‘-’ appears in the column. Entries in the pthreads column that are followed by “POSIX 1003.4” or “POSIX.4” are part of the POSIX Realtime standard specification and are not part of pthreads.

TABLE 8-2 Solaris Threads and POSIX pthreads Comparison

Solaris Threads (libthread)	pthreads (libpthread)
<code>thr_create()</code>	<code>pthread_create()</code>
<code>thr_exit()</code>	<code>pthread_exit()</code>
<code>thr_join()</code>	<code>pthread_join()</code>
<code>thr_yield()</code>	<code>sched_yield()</code> POSIX.4
<code>thr_self()</code>	<code>pthread_self()</code>
<code>thr_kill()</code>	<code>pthread_kill()</code>
<code>thr_sigsetmask()</code>	<code>pthread_sigmask()</code>
<code>thr_setprio()</code>	<code>pthread_setschedparam()</code>
<code>thr_getprio()</code>	<code>pthread_getschedparam()</code>
<code>thr_setconcurrency()</code>	<code>pthread_setconcurrency()</code>
<code>thr_getconcurrency()</code>	<code>pthread_getconcurrency()</code>
<code>thr_suspend()</code>	-
<code>thr_continue()</code>	-
<code>thr_keycreate()</code>	<code>pthread_key_create()</code>
-	<code>pthread_key_delete()</code>
<code>thr_setspecific()</code>	<code>pthread_setspecific()</code>
<code>thr_getspecific()</code>	<code>pthread_getspecific()</code>

TABLE 8-2 Solaris Threads and POSIX pthreads Comparison *(continued)*

Solaris Threads (libthread)	pthreads (libpthread)
-	pthread_once()
-	pthread_equal()
-	pthread_cancel()
-	pthread_testcancel()
-	pthread_cleanup_push()
-	pthread_cleanup_pop()
-	pthread_setcanceltype()
-	pthread_setcancelstate()
mutex_lock()	pthread_mutex_lock()
mutex_unlock()	pthread_mutex_unlock()
mutex_trylock()	pthread_mutex_trylock()
mutex_init()	pthread_mutex_init()
mutex_destroy()	pthread_mutex_destroy()
cond_wait()	pthread_cond_wait()
cond_timedwait()	pthread_cond_timedwait()
cond_signal()	pthread_cond_signal()
cond_broadcast()	pthread_cond_broadcast()

TABLE 8-2 Solaris Threads and POSIX pthreads Comparison *(continued)*

Solaris Threads (libthread)	pthreads (libpthread)
<code>cond_init()</code>	<code>pthread_cond_init()</code>
<code>cond_destroy()</code>	<code>pthread_cond_destroy()</code>
<code>rwlock_init()</code>	<code>pthread_rwlock_init()</code>
<code>rwlock_destroy()</code>	<code>pthread_rwlock_destroy()</code>
<code>rw_rdlock()</code>	<code>pthread_rwlock_rdlock()</code>
<code>rw_wrlock()</code>	<code>pthread_rwlock_wrlock()</code>
<code>rw_unlock()</code>	<code>pthread_rwlock_unlock()</code>
<code>rw_tryrdlock()</code>	<code>pthread_rwlock_tryrdlock()</code>
<code>rw_trywrlock()</code>	<code>pthread_rwlock_trywrlock()</code>
-	<code>pthread_rwlockattr_init()</code>
-	<code>pthread_rwlockattr_destroy()</code>
-	<code>pthread_rwlockattr_getpshared()</code>
-	<code>pthread_rwlockattr_setpshared()</code>
<code>sema_init()</code>	<code>sem_init()</code> POSIX 1003.4
<code>sema_destroy()</code>	<code>sem_destroy()</code> POSIX 1003.4
<code>sema_wait()</code>	<code>sem_wait()</code> POSIX 1003.4
<code>sema_post()</code>	<code>sem_post()</code> POSIX 1003.4

TABLE 8-2 Solaris Threads and POSIX pthreads Comparison *(continued)*

Solaris Threads (libthread)	pthreads (libpthread)
<code>sema_trywait()</code>	<code>sem_trywait()</code> POSIX 1003.4
<code>fork1()</code>	<code>fork()</code>
-	<code>pthread_atfork()</code>
<code>fork()</code> (multiple thread copy)	-
-	<code>pthread_mutexattr_init()</code>
-	<code>pthread_mutexattr_destroy()</code>
<code>type()</code> argument in <code>cond_init()</code>	<code>pthread_mutexattr_setpshared()</code>
-	<code>pthread_mutexattr_getpshared()</code>
-	<code>pthread_mutex_attr_settype()</code>
-	<code>pthread_mutex_attr_gettype()</code>
-	<code>pthread_condattr_init()</code>
-	<code>pthread_condattr_destroy()</code>
<code>type()</code> argument in <code>cond_init()</code>	<code>pthread_condattr_setpshared()</code>
-	<code>pthread_condattr_getpshared()</code>
-	<code>pthread_attr_init()</code>
-	<code>pthread_attr_destroy()</code>
<code>THR_BOUND</code> flag in <code>thr_create()</code>	<code>pthread_attr_setscope()</code>
-	<code>pthread_attr_getscope()</code>

TABLE 8-2 Solaris Threads and POSIX pthreads Comparison *(continued)*

Solaris Threads (libthread)	pthreads (libpthread)
-	pthread_attr_setguardsize()
-	pthread_attr_getguardsize()
stack_size() argument in thr_create()	pthread_attr_setstacksize()
-	pthread_attr_getstacksize()
stack_addr() argument in thr_create()	pthread_attr_setstackaddr()
-	pthread_attr_getstackaddr()
THR_DETACH flag in thr_create()	pthread_attr_setdetachstate()
-	pthread_attr_getdetachstate()
-	pthread_attr_setschedparam()
-	pthread_attr_getschedparam()
-	pthread_attr_setinheritsched()
-	pthread_attr_getinheritsched()
-	pthread_attr_setsschedpolicy()
-	pthread_attr_getschedpolicy()

To use the Solaris threads functions described in this chapter, you must link with the Solaris threads library `-lthread`).

Where functionality is virtually the same for both Solaris threads and for pthreads, (even though the function names or arguments might differ), only a brief example

consisting of the correct include file and the function prototype is presented. Where return values are not given for the Solaris threads functions, see the appropriate pages in `man(3): Library Routines` for the function return values.

For more information on Solaris related functions, see the related `pthread`s documentation for the similarly named function.

Where Solaris threads functions offer capabilities that are not available in `pthread`s, a full description of the functions is provided.

Unique Solaris Threads Functions

- “Suspend Thread Execution” on page 194
- “Continue a Suspended Thread” on page 195
- “Set Thread Concurrency Level” on page 196
- “Get Thread Concurrency” on page 197

Suspend Thread Execution

`thr_suspend(3THR)`

`thr_suspend(3THR)` immediately suspends the execution of the thread specified by `target_thread`. On successful return from `thr_suspend( )`, the suspended thread is no longer executing.

Once a thread is suspended, subsequent calls to `thr_suspend( )` have no effect. Signals cannot awaken the suspended thread; they remain pending until the thread resumes execution.

```
#include <thread.h>
```

```
int thr_suspend(thread_t tid);
```

In the following synopsis, `pthread_t tid` as defined in `pthread`s is the same as `thread_t tid` in Solaris threads. `tid` values can be used interchangeably either by assignment or through the use of casts.

```
thread_t tid; /* tid from thr_create() */
```

```
/* pthreads equivalent of Solaris tid from thread created */  
/* with pthread_create() */  
pthread_t ptid;
```

```
int ret;
```

```
ret = thr_suspend(tid);

/* using pthreads ID variable with a cast */
ret = thr_suspend((thread_t) ptid);
```

Return Values

`thr_suspend( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, `thr_suspend( )` fails and returns the corresponding value.

ESRCH

tid cannot be found in the current process.

Continue a Suspended Thread

thr_continue(3THR)

`thr_continue(3THR)` resumes the execution of a suspended thread. Once a suspended thread is continued, subsequent calls to `thr_continue( )` have no effect.

```
#include <thread.h>
```

```
int thr_continue(thread_t tid);
```

A suspended thread will not be awakened by a signal. The signal stays pending until the execution of the thread is resumed by `thr_continue( )`.

`pthread_t tid` as defined in `pthreads` is the same as `thread_t tid` in Solaris threads. *tid* values can be used interchangeably either by assignment or through the use of casts.

```
thread_t tid; /* tid from thr_create() */

/* pthreads equivalent of Solaris tid from thread created */
/* with pthread_create() */
pthread_t ptid;

int ret;

ret = thr_continue(tid);

/* using pthreads ID variable with a cast */
ret = thr_continue((thread_t) ptid)
```

Return Values

`thr_continue( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When the following condition occurs, `thr_continue( )` fails and returns the corresponding value.

ESRCH

tid cannot be found in the current process.

Set Thread Concurrency Level

By default, Solaris threads attempt to adjust the system execution resources (LWPs) used to run unbound threads to match the real number of active threads. While the Solaris threads package cannot make perfect decisions, it at least ensures that the process continues to make progress.

When you have some idea of the number of unbound threads that should be simultaneously active (executing code or system calls), tell the library through `thr_setconcurrency( )`. To get the number of threads being used, use `thr_getconcurrency( )`.

`thr_setconcurrency(3THR)`

`thr_setconcurrency(3THR)` provides a hint to the system about the required level of concurrency in the application. The system ensures that a sufficient number of threads are active so that the process continues to make progress.

```
#include <thread.h>
```

```
int new_level;  
int ret;
```

```
ret = thr_setconcurrency(new_level);
```

Unbound threads in a process might or might not be required to be simultaneously active. To conserve system resources, the threads system ensures by default that enough threads are active for the process to make progress, and that the process will not deadlock through a lack of concurrency.

Because this might not produce the most effective level of concurrency, `thr_setconcurrency( )` permits the application to give the threads system a hint, specified by *new_level*, for the desired level of concurrency.

The actual number of simultaneously active threads can be larger or smaller than *new_level*.

Note that an application with multiple compute-bound threads can fail to schedule all the runnable threads if `thr_setconcurrency( )` has not been called to adjust the level of execution resources.

You can also affect the value for the desired concurrency level by setting the `THR_NEW_LWP` flag in `thr_create( )`. This effectively increments the current level by one.

Return Values

Returns a zero when it completes successfully. Any other returned value indicates that an error occurred. When any of the following conditions is detected, `thr_setconcurrency( )` fails and returns the corresponding value.

EAGAIN

The specified concurrency level would cause a system resource to be exceeded.

EINVAL

The value for *new_level* is negative.

Get Thread Concurrency

`thr_getconcurrency(3THR)`

Use `thr_getconcurrency(3THR)` to get the current value of the concurrency level previously set by `thr_setconcurrency( )`. Note that the actual number of simultaneously active threads can be larger or smaller than this number.

```
#include <thread.h>
```

```
int thr_getconcurrency(void)
```

Return Value

`thr_getconcurrency( )` always returns the current value for the desired concurrency level.

Similar Synchronization Functions—Read-Write Locks

Read-write locks allow simultaneous read access by many threads while restricting write access to only one thread at a time.

- “Initialize a Read-Write Lock” on page 198
- “Acquire a Read Lock” on page 200
- “Try to Acquire a Read Lock” on page 200
- “Acquire a Write Lock” on page 201
- “Try to Acquire a Write Lock” on page 202
- “Unlock a Read-Write Lock” on page 202
- “Destroy Read-Write Lock State” on page 203

When any thread holds the lock for reading, other threads can also acquire the lock for reading but must wait to acquire the lock for writing. If one thread holds the lock for writing, or is waiting to acquire the lock for writing, other threads must wait to acquire the lock for either reading or writing.

Read-write locks are slower than mutexes, but can improve performance when they protect data that are not frequently written but that are read by many concurrent threads.

Use read-write locks to synchronize threads in this process and other processes by allocating them in memory that is writable and shared among the cooperating processes (see `mmap(2)`) and by initializing them for this behavior.

By default, the acquisition order is not defined when multiple threads are waiting for a read-write lock. However, to avoid writer starvation, the Solaris threads package tends to favor writers over readers.

Read-write locks must be initialized before use.

Initialize a Read-Write Lock

`rwlock_init(3THR)`

```
#include <synch.h> (or #include <thread.h>)

int rwlock_init(rwlock_t *rwlp, int type, void * arg);
```

Use `rwlock_init(3THR)` to initialize the read-write lock pointed to by `rwlp` and to set the lock state to unlocked. `type` can be one of the following (note that `arg` is currently ignored). (For POSIX threads, see “`pthread_rwlock_init(3THR)`” on page 134.)

- **USYNC_PROCESS** The read-write lock can be used to synchronize threads in this process and other processes. `arg` is ignored.
- **USYNC_THREAD** The read-write lock can be used to synchronize threads in this process, only. `arg` is ignored.

Multiple threads must not initialize the same read-write lock simultaneously. Read-write locks can also be initialized by allocation in zeroed memory, in which case a type of **USYNC_THREAD** is assumed. A read-write lock must not be reinitialized while other threads might be using it.

Initializing Read-Write Locks With Intraprocess Scope

```
#include <thread.h>

rwlock_t rwlp;
int ret;

/* to be used within this process only */
ret = rwlock_init(&rwlp, USYNC_THREAD, 0);
```

Initializing Read-Write Locks With Interprocess Scope

```
#include <thread.h>

rwlock_t rwlp;
int ret;

/* to be used among all processes */
ret = rwlock_init(&rwlp, USYNC_PROCESS, 0);
```

Return Values

`rwlock_init( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

Invalid argument.

EFAULT

`rwlp` or `arg` points to an illegal address.

Acquire a Read Lock

`rw_rdlock(3THR)`

```
#include <synch.h> (or #include <thread.h>)
```

```
int rw_rdlock(rwlock_t *rwlp);
```

Use `rw_rdlock(3THR)` to acquire a read lock on the read-write lock pointed to by `rwlp`. When the read-write lock is already locked for writing, the calling thread blocks until the write lock is released. Otherwise, the read lock is acquired. (For POSIX threads, see “`pthread_rwlock_rdlock(3THR)`” on page 135.)

Return Values

`rw_rdlock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

Invalid argument.

EFAULT

`rwlp` points to an illegal address.

Try to Acquire a Read Lock

`rw_tryrdlock(3THR)`

```
#include <synch.h> (or #include <thread.h>)
```

```
int rw_tryrdlock(rwlock_t *rwlp);
```

Use `rw_tryrdlock(3THR)` to attempt to acquire a read lock on the read-write lock pointed to by `rwlp`. When the read-write lock is already locked for writing, it returns an error. Otherwise, the read lock is acquired. (For POSIX threads, see “`pthread_rwlock_tryrdlock(3THR)`” on page 136.)

Return Values

`rw_tryrdlock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

Invalid argument.

EFAULT

rwlp points to an illegal address.

EBUSY

The read-write lock pointed to by *rwlp* was already locked.

Acquire a Write Lock

`rw_wrlock(3THR)`

`#include <synch.h>` (or `#include <thread.h>`)

`int rw_wrlock(rwlock_t *rwlp);`

Use `rw_wrlock(3THR)` to acquire a write lock on the read-write lock pointed to by *rwlp*. When the read-write lock is already locked for reading or writing, the calling thread blocks until all the read locks and write locks are released. Only one thread at a time can hold a write lock on a read-write lock. (For POSIX threads, see “`pthread_rwlock_wrlock(3THR)`” on page 136.)

Return Values

`rw_wrlock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

Invalid argument.

EFAULT

rwlp points to an illegal address.

Try to Acquire a Write Lock

`rw_trywrlock(3THR)`

`#include <synch.h> (or #include <thread.h>)`

`int rw_trywrlock(rwlock_t *rwlp);`

Use `rw_trywrlock(3THR)` to attempt to acquire a write lock on the read-write lock pointed to by `rwlp`. When the read-write lock is already locked for reading or writing, it returns an error. (For POSIX threads, see “`pthread_rwlock_trywrlock(3THR)`” on page 137.)

Return Values

`rw_trywrlock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

Invalid argument.

EFAULT

`rwlp` points to an illegal address.

EBUSY

The read-write lock pointed to by `rwlp` was already locked.

Unlock a Read-Write Lock

`rw_unlock(3THR)`

`#include <synch.h> (or #include <thread.h>)`

`int rw_unlock(rwlock_t *rwlp);`

Use `rw_unlock(3THR)` to unlock a read-write lock pointed to by `rwlp`. The read-write lock must be locked and the calling thread must hold the lock either for reading or writing. When any other threads are waiting for the read-write lock to become available, one of them is unblocked. (For POSIX threads, see “`pthread_rwlock_unlock(3THR)`” on page 138.)

Return Values

`rw_unlock( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

Invalid argument.

EFAULT

rwlp points to an illegal address.

Destroy Read-Write Lock State

`rwlock_destroy(3THR)`

```
#include <synch.h> (or #include <thread.h>)
```

```
int rwlock_destroy(rwlock_t *rwlp);
```

Use `rwlock_destroy(3THR)` to destroy any state associated with the read-write lock pointed to by *rwlp*. The space for storing the read-write lock is not freed. (For POSIX threads, see “`pthread_rwlock_destroy(3THR)`” on page 139.)

Return Values

`rwlock_destroy( )` returns zero after completing successfully. Any other returned value indicates that an error occurred. When any of the following conditions occurs, the function fails and returns the corresponding value.

EINVAL

Invalid argument.

EFAULT

rwlp points to an illegal address.

Read-Write Lock Example

Code Example 8–1 uses a bank account to demonstrate read-write locks. While the program could allow multiple threads to have concurrent read-only access to the account balance, only a single writer is allowed. Note that the `get_balance()` function

needs the lock to ensure that the addition of the checking and saving balances occurs atomically.

CODE EXAMPLE 8-1 Read-Write Bank Account

```
rwlock_t account_lock;
float checking_balance = 100.0;
float saving_balance = 100.0;
...
rwlock_init(&account_lock, 0, NULL);
...

float
get_balance() {
    float bal;

    rw_rdlock(&account_lock);
    bal = checking_balance + saving_balance;
    rw_unlock(&account_lock);
    return(bal);
}

void
transfer_checking_to_savings(float amount) {
    rw_wrlock(&account_lock);
    checking_balance = checking_balance - amount;
    saving_balance = saving_balance + amount;
    rw_unlock(&account_lock);
}
```

Similar Solaris Threads Functions

Operation	Destination Discussion
Create a thread	"thr_create(3THR)" on page 205
Get the minimal stack size	"thr_min_stack(3THR)" on page 208
Get the thread identifier	"thr_self(3THR)" on page 209
Yield thread execution	"thr_yield(3THR)" on page 209
Send a signal to a thread	"thr_kill(3THR)" on page 209

Operation	Destination Discussion
Access the signal mask of the calling thread	“thr_sigsetmask(3THR)” on page 209
Terminate a thread	“thr_exit(3THR)” on page 210
Wait for thread termination	“thr_join(3THR)” on page 210
Create a thread-specific data key	“thr_keycreate(3THR)” on page 211
Set the thread-specific data key	“thr_setspecific(3THR)” on page 211
Get the thread-specific data key	“thr_getspecific(3THR)” on page 212
Set the thread priority	“thr_setprio(3THR)” on page 212
Get the thread priority	“thr_getprio(3THR)” on page 213

Create a Thread

The `thr_create(3THR)` routine is one of the most elaborate of all the Solaris threads library routines.

`thr_create(3THR)`

Use `thr_create(3THR)` to add a new thread of control to the current process. (For POSIX threads, see “`pthread_create(3THR)`” on page 28.)

Note that the new thread does not inherit pending signals, but it does inherit priority and signal masks.

```
#include <thread.h>
```

```
int thr_create(void *stack_base, size_t stack_size,
               void *(*start_routine) (void *), void *arg, long flags,
               thread_t *new_thread);
```

```
size_t thr_min_stack(void);
```

stack_base—Contains the address for the stack that the new thread uses. If *stack_base* is NULL then `thr_create()` allocates a stack for the new thread with at least *stack_size* bytes.

stack_size—Contains the size, in number of bytes, for the stack that the new thread uses. If *stack_size* is zero, a default size is used. In most cases, a zero value works best. If *stack_size* is not zero, it must be greater than the value returned by `thr_min_stack()`.

There is no general need to allocate stack space for threads. The threads library allocates 1 Mbyte of virtual memory for each thread's stack with no swap space reserved. (The library uses the `-MAP_NORESERVE` option of `mmap(2)` to make the allocations.)

start_routine—Contains the function with which the new thread begins execution. When `start_routine()` returns, the thread exits with the exit status set to the value returned by *start_routine* (see “`thr_exit(3THR)`” on page 210).

arg—Can be anything that is described by `void`, which is typically any 4-byte value. Anything larger must be passed indirectly by having the argument point to it.

Note that you can supply only one argument. To get your procedure to take multiple arguments, encode them as one (such as by putting them in a structure).

flags—() Specifies attributes for the created thread. In most cases a zero value works best.

The value in *flags* is constructed from the bitwise inclusive OR of the following:

- **THR_SUSPENDED**—Suspends the new thread and does not execute *start_routine* until the thread is started by `thr_continue()`. Use this to operate on the thread (such as changing its priority) before you run it. The termination of a detached thread is ignored.
- **THR_DETACHED**—Detaches the new thread so that its thread ID and other resources can be reused as soon as the thread terminates. Set this when you do not want to wait for the thread to terminate.

Note - When there is no explicit synchronization to prevent it, an unsuspended, detached thread can die and have its thread ID reassigned to another new thread before its creator returns from `thr_create()`.

- **THR_BOUND**—Permanently binds the new thread to an LWP (the new thread is a *bound thread*).
- **THR_NEW_LWP**—Increases the concurrency level for unbound threads by one. The effect is similar to incrementing concurrency by one with `thr_setconcurrency(3T)`, although **THR_NEW_LWP** does not affect the level set through the `thr_setconcurrency()` function. Typically, **THR_NEW_LWP** adds a new LWP to the pool of LWPs running unbound threads.

- When you specify both `THR_BOUND` and `THR_NEW_LWP`, two LWPs are typically created—one for the bound thread and another for the pool of LWPs running unbound threads.
- `THR_DAEMON`—Marks the new thread as a daemon. The process exits when all nondaemon threads exit. Daemon threads do not affect the process exit status and are ignored when counting the number of thread exits.

A process can exit either by calling `exit()` or by having every thread in the process that was not created with the `THR_DAEMON` flag call `thr_exit(3T)`. An application, or a library it calls, can create one or more threads that should be ignored (not counted) in the decision of whether to exit. The `THR_DAEMON` flag identifies threads that are not counted in the process exit criterion.

new_thread—Points to a location (when *new_thread* is not `NULL`) where the ID of the new thread is stored when `thr_create()` is successful. The caller is responsible for supplying the storage this argument points to. The ID is valid only within the calling process.

If you are not interested in this identifier, supply a zero value to *new_thread*.

Return Values

Returns a zero and exits when it completes successfully. Any other returned value indicates that an error occurred. When any of the following conditions is detected, `thr_create()` fails and returns the corresponding value.

`EAGAIN`

A system limit is exceeded, such as when too many LWPs have been created.

`ENOMEM`

Not enough memory was available to create the new thread.

`EINVAL`

stack_base is not `NULL` and *stack_size* is less than the value returned by `thr_min_stack()`.

Stack Behavior

Stack behavior in Solaris threads is generally the same as that in pthreads. For more information about stack setup and operation, see “About Stacks” on page 71.

You can get the absolute minimum on stack size by calling `thr_min_stack()`, which returns the amount of stack space required for a thread that executes a null procedure. Useful threads need more than this, so be very careful when reducing the stack size.

You can specify a custom stack in two ways. The first is to supply a `NULL` for the stack location, thereby asking the runtime library to allocate the space for the stack, but to supply the desired size in the `stacksize` parameter to `thr_create()`.

The other approach is to take overall aspects of stack management and supply a pointer to the stack to `thr_create()`. This means that you are responsible not only for stack allocation but also for stack deallocation—when the thread terminates, you must arrange for the disposal of its stack.

When you allocate your own stack, be sure to append a red zone to its end by calling `mprotect(2)`.

Get the Minimal Stack Size

`thr_min_stack(3THR)`

Use `thr_min_stack(3THR)` to get the minimum stack size for a thread.

```
#include <thread.h>
```

```
size_t thr_min_stack(void);
```

`thr_min_stack()` returns the amount of space needed to execute a null thread (a null thread is a thread that is created to execute a null procedure).

A thread that does more than execute a null procedure should allocate a stack size greater than the size of `thr_min_stack()`.

When a thread is created with a user-supplied stack, the user must reserve enough space to run the thread. In a dynamically linked execution environment, it is difficult to know what the thread minimal stack requirements are.

Most users should not create threads with user-supplied stacks. User-supplied stacks exist only to support applications that require complete control over their execution environments.

Instead, users should let the threads library manage stack allocation. The threads library provides default stacks that should meet the requirements of any created thread.

Get the Thread Identifier

thr_self(3THR)

Use `thr_self(3THR)` to get the ID of the calling thread. (For POSIX threads, see “`pthread_self(3THR)`” on page 38.)

```
#include <thread.h>

thread_t thr_self(void);
```

Yield Thread Execution

thr_yield(3THR)

`thr_yield(3THR)` causes the current thread to yield its execution in favor of another thread with the same or greater priority; otherwise it has no effect. There is no guarantee that a thread calling `thr_yield( )` will do so.

```
#include <thread.h>

void thr_yield(void);
```

Send a Signal to a Thread

thr_kill(3THR)

`thr_kill(3THR)` sends a signal to a thread. (For POSIX threads, see “`pthread_kill(3THR)`” on page 42.)

```
#include <thread.h>
#include <signal.h>

int thr_kill(thread_t target_thread, int sig);
```

Access the Signal Mask of the Calling Thread

thr_sigsetmask(3THR)

Use `thr_sigsetmask(3THR)` to change or examine the signal mask of the calling thread.

```
#include <thread.h>
#include <signal.h>
```

```
int thr_sigsetmask(int how, const sigset_t *set, sigset_t *oset);
```

Terminate a Thread

thr_exit(3THR)

Use `thr_exit(3THR)` to terminate a thread. (For POSIX threads, see “`pthread_exit(3THR)`” on page 44.)

```
#include <thread.h>
```

```
void thr_exit(void *status);
```

Wait for Thread Termination

thr_join(3THR)

Use `thr_join(3THR)` to wait for a thread to terminate. (For POSIX threads, see “`pthread_join(3THR)`” on page 29.)

```
#include <thread.h>
```

```
int thr_join(thread_t tid, thread_t *departedid, void **status);
```

Join specific

```
#include <thread.h>
```

```
thread_t tid;  
thread_t departedid;  
int ret;  
int status;
```

```
/* waiting to join thread "tid" with status */  
ret = thr_join(tid, &departedid, (void**)&status);
```

```
/* waiting to join thread "tid" without status */  
ret = thr_join(tid, &departedid, NULL);
```

```
/* waiting to join thread "tid" without return id and status */  
ret = thr_join(tid, NULL, NULL);
```

When the *tid* is `(thread_t)0`, then `thread_join( )` waits for any undetached thread in the process to terminate. In other words, when no thread identifier is specified, any undetached thread that exits causes `thread_join( )` to return.

Join any

```
#include <thread.h>

thread_t tid;
thread_t departedid;
int ret;
int status;

/* waiting to join thread "tid" with status */
ret = thr_join(NULL, &departedid, (void **)&status);
```

By indicating NULL as thread id in the Solaris `thr_join( )`, a join will take place when any non detached thread in the process exits. The *departedid* will indicate the thread ID of exiting thread.

Create a Thread-Specific Data Key

Except for the function names and arguments, thread specific data is the same for Solaris as it is for POSIX. The synopses for the Solaris functions are given in this section.

thr_keycreate(3THR)

`thr_keycreate(3THR)` allocates a key that is used to identify thread-specific data in a process. (For POSIX threads, see “`pthread_key_create(3THR)`” on page 33.)

```
#include <thread.h>

int thr_keycreate(thread_key_t *keyp,
    void (*destructor) (void *value));
```

Set the Thread-Specific Data Key

thr_setspecific(3THR)

`thr_setspecific(3THR)` binds *value* to the thread-specific data key, *key*, for the calling thread. (For POSIX threads, see “`pthread_setspecific(3THR)`” on page 35.)

```
#include <thread.h>

int thr_setspecific(thread_key_t key, void *value);
```

Get the Thread-Specific Data Key

thr_getspecific(3THR)

`thr_getspecific(3THR)` stores the current value bound to *key* for the calling thread into the location pointed to by *valuep*. (For POSIX threads, see “`pthread_getspecific(3THR)`” on page 35.)

```
#include <thread.h>
```

```
int thr_getspecific(thread_key_t key, void **valuep);
```

Set the Thread Priority

In Solaris threads, if a thread is to be created with a priority other than that of its parent’s, it is created in SUSPEND mode. While suspended, the thread’s priority is modified using the `thr_setprio(3T)` function call; then it is continued.

An unbound thread is usually scheduled only with respect to other threads in the process using simple priority levels with no adjustments and no kernel involvement. Its system priority is usually uniform and is inherited from the creating process.

thr_setprio(3THR)

`thr_setprio(3THR)` changes the priority of the thread, specified by *tid*, within the current process to the priority specified by *newprio*. (For POSIX threads, see “`pthread_setschedparam(3THR)`” on page 40.)

```
#include <thread.h>
```

```
int thr_setprio(thread_t tid, int newprio)
```

By default, threads are scheduled based on fixed priorities that range from zero, the least significant, to the largest integer. The *tid* will preempt lower-priority threads, and will yield to higher-priority threads.

```
thread_t tid;
```

```
int ret;
```

```
int newprio = 20;
```

```
/* suspended thread creation */
```

```
ret = thr_create(NULL, NULL, func, arg, THR_SUSPEND, &tid);
```

```
/* set the new priority of suspended child thread */
```

```
ret = thr_setprio(tid, newprio);
```

```
/* suspended child thread starts executing with new priority */
```

```
ret = thr_continue(tid);
```

Get the Thread Priority

thr_getprio(3THR)

Use `thr_getprio(3THR)` to get the current priority for the thread. Each thread inherits a priority from its creator. `thr_getprio( )` stores the current priority, *tid*, in the location pointed to by *newprio*. (For POSIX threads, see “`pthread_getschedparam(3THR)`” on page 41.)

```
#include <thread.h>
```

```
int thr_getprio(thread_t tid, int *newprio)
```

Similar Synchronization Functions—Mutual Exclusion Locks

- “Initialize a Mutex” on page 213
- “Destroy a Mutex” on page 215
- “Acquire a Mutex” on page 215
- “Release a Mutex” on page 215
- “Try to Acquire a Mutex” on page 216

Initialize a Mutex

mutex_init(3THR)

```
#include <synch.h> (or  
#include <thread.h>)
```

```
int mutex_init(mutex_t *mp, int type, void *arg);
```

Use `mutex_init(3THR)` to initialize the mutex pointed to by *mp*. The *type* can be one of the following (note that *arg* is currently ignored). (For POSIX threads, see “Initialize a Mutex” on page 94.)

- `USYNC_PROCESS` The mutex can be used to synchronize threads in this and other processes.

- `USYNC_PROCESS_ROBUST` The mutex can be used to *robustly* synchronize threads in this and other processes.
- `USYNC_THREAD` The mutex can be used to synchronize threads in this process only.

When a process dies while holding a `USYNC_PROCESS` lock, subsequent requestors of that lock hang. This is a problem for systems which share locks with client processes because the client processes can be abnormally killed. To avoid the problem of hanging on a lock held by a dead process, use `USYNC_PROCESS_ROBUST` to lock the mutex. `USYNC_PROCESS_ROBUST` adds two capabilities:

- In the case of process death, all owned locks held by that process are unlocked.
- The next requestor for any of the locks owned by the dead process receives the lock, but with an error return indicating that the previous owner died while holding the lock..

Mutexes can also be initialized by allocation in zeroed memory, in which case a *type* of `USYNC_THREAD` is assumed.

Multiple threads must not initialize the same mutex simultaneously. A mutex lock must not be reinitialized while other threads might be using it.

Mutexes With Intraprocess Scope

```
#include <thread.h>

mutex_t mp;
int ret;

/* to be used within this process only */
ret = mutex_init(&mp, USYNC_THREAD, 0);
```

Mutexes With Interprocess Scope

```
#include <thread.h>

mutex_t mp;
int ret;

/* to be used among all processes */
ret = mutex_init(&mp, USYNC_PROCESS, 0);
```

Mutexes With Interprocess Scope-Robust

```
#include <thread.h>

mutex_t mp;
int ret;

/* to be used among all processes */
ret = mutex_init(&mp, USYNC_PROCESS_ROBUST, 0);
```

Destroy a Mutex

`mutex_destroy(3THR)`

```
#include <thread.h>
```

```
int mutex_destroy (mutex_t *mp);
```

Use `mutex_destroy(3THR)` to destroy any state associated with the mutex pointed to by *mp*. Note that the space for storing the mutex is not freed. (For POSIX threads, see “`pthread_mutex_destroy(3THR)`” on page 101.)

Acquire a Mutex

`mutex_lock(3THR)`

```
#include <thread.h>
```

```
int mutex_lock(mutex_t *mp);
```

Use `mutex_lock(3THR)` to lock the mutex pointed to by *mp*. When the mutex is already locked, the calling thread blocks until the mutex becomes available (blocked threads wait on a prioritized queue). (For POSIX threads, see “`pthread_mutex_lock(3THR)`” on page 97.)

Release a Mutex

`mutex_unlock(3THR)`

```
#include <thread.h>
```

```
int mutex_unlock(mutex_t *mp);
```

Use `mutex_unlock(3THR)` to unlock the mutex pointed to by *mp*. The mutex must be locked and the calling thread must be the one that last locked the mutex (the owner). (For POSIX threads, see “`pthread_mutex_unlock(3THR)`” on page 98.)

Try to Acquire a Mutex

`mutex_trylock(3THR)`

```
#include <thread.h>
```

```
int mutex_trylock(mutex_t *mp);
```

Use `mutex_trylock(3THR)` to attempt to lock the mutex pointed to by *mp*. This function is a nonblocking version of `mutex_lock( )`. (For POSIX threads, see “`pthread_mutex_trylock(3THR)`” on page 99.)

Similar Synchronization Functions—Condition Variables

- “Initialize a Condition Variable” on page 216
- “Destroy a Condition Variable” on page 217
- “Wait for a Condition” on page 217
- “Wait for an Absolute Time” on page 218
- “Signal One Condition Variable” on page 218
- “Signal All Condition Variables” on page 219

Initialize a Condition Variable

`cond_init(3THR)`

```
#include <thread.h>
```

```
int cond_init(cond_t *cv, int type, int arg);
```

Use `cond_init(3THR)` to initialize the condition variable pointed to by *cv*. The *type* can be one of the following (note that *arg* is currently ignored). (For POSIX threads, see “`pthread_condattr_init(3THR)`” on page 107.)

- `USYNC_PROCESS` The condition variable can be used to synchronize threads in this and other processes. *arg* is ignored.
- `USYNC_THREAD` The condition variable can be used to synchronize threads in this process only. *arg* is ignored.

Condition variables can also be initialized by allocation in zeroed memory, in which case a type of USYNC_THREAD is assumed.

Multiple threads must not initialize the same condition variable simultaneously. A condition variable must not be reinitialized while other threads might be using it.

Condition Variables With Intraprocess Scope

```
#include <thread.h>

cond_t cv;
int ret;

/* to be used within this process only */
ret = cond_init(cv, USYNC_THREAD, 0);
```

Condition Variables With Interprocess Scope

```
#include <thread.h>

cond_t cv;
int ret;

/* to be used among all processes */
ret = cond_init(&cv, USYNC_PROCESS, 0);
```

Destroy a Condition Variable

`cond_destroy(3THR)`

```
#include <thread.h>

int cond_destroy(cond_t *cv);
```

Use `cond_destroy(3THR)` to destroy state associated with the condition variable pointed to by `cv`. The space for storing the condition variable is not freed. (For POSIX threads, see “`pthread_condattr_destroy(3THR)`” on page 108.)

Wait for a Condition

`cond_wait(3THR)`

```
#include <thread.h>

int cond_wait(cond_t *cv, mutex_t *mp);
```

Use `cond_wait(3THR)` to atomically release the mutex pointed to by *mp* and to cause the calling thread to block on the condition variable pointed to by *cv*. The blocked thread can be awakened by `cond_signal()`, `cond_broadcast()`, or when interrupted by delivery of a signal or a `fork()`. (For POSIX threads, see “`pthread_cond_wait(3THR)`” on page 112.)

Wait for an Absolute Time

`cond_timedwait(3THR)`

```
#include <thread.h>
```

```
int cond_timedwait(cond_t *cv, mutex_t *mp, timestruct_t abstime)
```

Use `cond_timedwait(3THR)` as you would use `cond_wait()`, except that `cond_timedwait()` does not block past the time of day specified by *abstime*. (For POSIX threads, see “`pthread_cond_timedwait(3THR)`” on page 115.)

`cond_timedwait()` always returns with the mutex locked and owned by the calling thread even when returning an error.

The `cond_timedwait()` function blocks until the condition is signaled or until the time of day specified by the last argument has passed. The timeout is specified as the time of day so the condition can be retested efficiently without recomputing the time-out value.

Signal One Condition Variable

`cond_signal(3THR)`

```
#include <thread.h>
```

```
int cond_signal(cond_t *cv);
```

Use `cond_signal(3THR)` to unblock one thread that is blocked on the condition variable pointed to by *cv*. Call this function under protection of the same mutex used with the condition variable being signaled. Otherwise, the condition could be signaled between its test and `cond_wait()`, causing an infinite wait.

Signal All Condition Variables

`cond_broadcast(3THR)`

```
#include <thread.h>
```

```
int cond_broadcast(cond_t *cv);
```

Use `cond_broadcast(3THR)` to unblock all threads that are blocked on the condition variable pointed to by `cv`. When no threads are blocked on the condition variable then `cond_broadcast( )` has no effect.

Similar Synchronization Functions—Semaphores

Semaphore operations are the same in both the Solaris Operating Environment and the POSIX environment. The function name changed from `sema_` in the Solaris Operating Environment to `sem_` in pthreads.

- “Initialize a Semaphore” on page 219
- “Increment a Semaphore” on page 220
- “Block on a Semaphore Count” on page 220
- “Decrement a Semaphore Count” on page 221
- “Destroy the Semaphore State” on page 221

Initialize a Semaphore

`sema_init(3THR)`

```
#include <thread.h>
```

```
int sema_init(sema_t *sp, unsigned int count, int type,  
              void *arg);
```

Use `sema_init(3THR)` to initialize the semaphore variable pointed to by `sp` by `count` amount. `type` can be one of the following (note that `arg` is currently ignored).

USYNC_PROCESS The semaphore can be used to synchronize threads in this process and other processes. Only one process should initialize the semaphore. `arg` is ignored.

USYNC_THREAD The semaphore can be used to synchronize threads in this process, only. *arg* is ignored.

Multiple threads must not initialize the same semaphore simultaneously. A semaphore must not be reinitialized while other threads might be using it.

Semaphores With Intraprocess Scope

```
#include <thread.h>

sema_t sp;
int ret;
int count;
count = 4;

/* to be used within this process only */
ret = sema_init(&sp, count, USYNC_THREAD, 0);
```

Semaphores With Interprocess Scope

```
#include <thread.h>

sema_t sp;
int ret;
int count;
count = 4;

/* to be used among all the processes */
ret = sema_init (&sp, count, USYNC_PROCESS, 0);
```

Increment a Semaphore

`sema_post(3THR)`

```
#include <thread.h>

int sema_post(sema_t *sp);
```

Use `sema_post(3THR)` to atomically increment the semaphore pointed to by *sp*. When any threads are blocked on the semaphore, one is unblocked.

Block on a Semaphore Count

`sema_wait(3THR)`

```
#include <thread.h>
```

```
int sema_wait(sema_t *sp);
```

Use `sema_wait(3THR)` to block the calling thread until the count in the semaphore pointed to by `sp` becomes greater than zero, then atomically decrement it.

Decrement a Semaphore Count

```
sema_trywait(3THR)
```

```
#include <thread.h>
```

```
int sema_trywait(sema_t *sp);
```

Use `sema_trywait(3THR)` to atomically decrement the count in the semaphore pointed to by `sp` when the count is greater than zero. This function is a nonblocking version of `sema_wait( )`.

Destroy the Semaphore State

```
sem_destroy(3RT)
```

```
#include <thread.h>
```

```
int sem_destroy(sema_t *sp);
```

Use `sem_destroy(3RT)` to destroy any state associated with the semaphore pointed to by `sp`. The space for storing the semaphore is not freed.

Synchronization Across Process Boundaries

Each of the synchronization primitives can be set up to be used across process boundaries. This is done quite simply by ensuring that the synchronization variable is located in a shared memory segment and by calling the appropriate `init` routine with `type` set to `USYNC_PROCESS`.

If this has been done, then the operations on the synchronization variables work just as they do when `type` is `USYNC_THREAD`.

```

mutex_init(&m, USYNC_PROCESS, 0);

rwlock_init(&rw, USYNC_PROCESS, 0);

cond_init(&cv, USYNC_PROCESS, 0);

sema_init(&s, count, USYNC_PROCESS, 0);

```

Using LWPs Between Processes

Using locks and condition variables between processes does not require using the threads library. The recommended approach is to use the threads library interfaces, but when this is not desirable, then the `_lwp_mutex_*` and `_lwp_cond_*` interfaces can be used as follows:

1. Allocate the locks and condition variables as usual in shared memory (either with `shmop(2)` or `mmap(2)`).
2. Then initialize the newly allocated objects appropriately with the `USYNC_PROCESS` type. Because no interface is available to perform the initialization (`_lwp_mutex_init(2)` and `_lwp_cond_init(2)` do not exist), the objects can be initialized using statically allocated and initialized dummy objects.

For example, to initialize `lockp`:

```

lwp_mutex_t *lwp_lockp;
lwp_mutex_t dummy_shared_mutex = SHARED_MUTEX;
/* SHARED_MUTEX is defined in /usr/include/synch.h */
...
...
lwp_lockp = alloc_shared_lock();
*lwp_lockp = dummy_shared_mutex;

```

Similarly, for condition variables:

```

lwp_cond_t *lwp_condp;
lwp_cond_t dummy_shared_cv = SHARED_CV;
/* SHARED_CV is defined in /usr/include/synch.h */
...
...
lwp_condp = alloc_shared_cv();
*lwp_condp = dummy_shared_cv;

```

Producer/Consumer Problem Example

Code Example 8-2 shows the producer/consumer problem with the producer and consumer in separate processes. The main routine maps zero-filled memory (that it shares with its child process) into its address space. Note that `mutex_init()` and `cond_init()` must be called because the type of the synchronization variables is `USYNC_PROCESS`.

A child process is created that runs the consumer. The parent runs the producer.

This example also shows the drivers for the producer and consumer. The `producer_driver()` simply reads characters from `stdin` and calls `producer()`. The `consumer_driver()` gets characters by calling `consumer()` and writes them to `stdout`.

The data structure for Code Example 8-2 is the same as that used for the solution with condition variables (see “Nested Locking With a Singly Linked List” on page 104).

CODE EXAMPLE 8-2 The Producer/Consumer Problem, Using `USYNC_PROCESS`

```
main() {
    int zfd;
    buffer_t *buffer;

    zfd = open('/dev/zero', O_RDWR);
    buffer = (buffer_t *)mmap(NULL, sizeof(buffer_t),
        PROT_READ|PROT_WRITE, MAP_SHARED, zfd, 0);
    buffer->occupied = buffer->nextin = buffer->nextout = 0;

    mutex_init(&buffer->lock, USYNC_PROCESS, 0);
    cond_init(&buffer->less, USYNC_PROCESS, 0);
    cond_init(&buffer->more, USYNC_PROCESS, 0);
    if (fork() == 0)
        consumer_driver(buffer);
    else
        producer_driver(buffer);
}

void producer_driver(buffer_t *b) {
    int item;

    while (1) {
        item = getchar();
        if (item == EOF) {
            producer(b, '\0');
            break;
        } else
            producer(b, (char)item);
    }
}

void consumer_driver(buffer_t *b) {
    char item;

    while (1) {
        if ((item = consumer(b)) == '\0')
            break;
        putchar(item);
    }
}
```

A child process is created to run the consumer; the parent runs the producer.

Special Issues for `fork()` and Solaris Threads

Solaris threads and POSIX threads define the behavior of `fork()` differently. See “Process Creation—`exec(2)` and `exit(2)` Issues” on page 148 for a thorough discussion of `fork()` issues.

Solaris `libthread` supports both `fork()` and `fork1()`. The `fork()` call has “fork-all” semantics—it duplicates everything in the process, including threads and LWPs, creating a true clone of the parent. The `fork1()` call creates a clone that has only one thread; the process state and address space are duplicated, but only the calling thread is cloned.

POSIX `libpthread` supports only `fork()`, which has the same semantics as `fork1()` in Solaris threads.

Whether `fork()` has “fork-all” semantics or “fork-one” semantics is dependent on which library is used. Linking with `-lthread` assigns “fork-all” semantics to `fork()`, while linking with `-lpthread` assigns “fork-one” semantics to `fork()`.

See “Linking With `libthread` or `libpthread`” on page 177 for more details.

Programming Guidelines

This chapter gives some pointers on programming with threads. Most pointers apply to both Solaris and POSIX threads, but where functionality differs, it is noted. Changing from single-threaded thinking to multithreaded thinking is emphasized in this chapter.

- “Rethinking Global Variables” on page 225
- “Providing for Static Local Variables” on page 226
- “Synchronizing Threads” on page 227
- “Avoiding Deadlock” on page 230
- “Following Some Basic Guidelines” on page 232
- “Creating and Using Threads” on page 232
- “Working With Multiprocessors” on page 236
- “Summary” on page 241

Rethinking Global Variables

Historically, most code has been designed for single-threaded programs. This is especially true for most of the library routines called from C programs. The following implicit assumptions were made for single-threaded code:

- When you write into a global variable and then, a moment later, read from it, what you read is exactly what you just wrote.
- This is also true for nonglobal, static storage.
- You do not need synchronization because there is nothing to synchronize with.

The next few examples discuss some of the problems that arise in multithreaded programs because of these assumptions, and how you can deal with them.

Traditional, single-threaded C and UNIX have a convention for handling errors detected in system calls. System calls can return anything as a functional value (for example, `write( )` returns the number of bytes that were transferred). However, the value `-1` is reserved to indicate that something went wrong. So, when a system call returns `-1`, you know that it failed.

CODE EXAMPLE 9-1 Global Variables and *errno*

```
extern int errno;
...
if (write(file_desc, buffer, size) == -1) {
    /* the system call failed */
    fprintf(stderr, "something went wrong, "
               "error code = %d\n", errno);
    exit(1);
}
...
```

Rather than return the actual error code (which could be confused with normal return values), the error code is placed into the global variable `errno`. When the system call fails, you can look in `errno` to find out what went wrong.

Now consider what happens in a multithreaded environment when two threads fail at about the same time, but with different errors. Both expect to find their error codes in `errno`, but one copy of `errno` cannot hold both values. This global variable approach simply does not work for multithreaded programs.

Threads solves this problem through a conceptually new storage class—thread-specific data. This storage is similar to global storage in that it can be accessed from any procedure in which a thread might be running. However, it is private to the thread—when two threads refer to the thread-specific data location of the same name, they are referring to two different areas of storage.

So, when using threads, each reference to `errno` is thread specific because each thread has a private copy of `errno`. This is achieved in this implementation by making `errno` a macro that expands to a function call.

Providing for Static Local Variables

Code Example 9-2 shows a problem similar to the `errno` problem, but involving static storage instead of global storage. The function `gethostbyname(3N)` is called with the computer name as its argument. The return value is a pointer to a structure containing the required information for contacting the computer through network communications.

CODE EXAMPLE 9-2 The gethostbyname() Problem

```
struct hostent *gethostbyname(char *name) {
    static struct hostent result;
    /* Lookup name in hosts database */
    /* Put answer in result */
    return(&result);
}
```

Returning a pointer to a local variable is generally not a good idea, although it works in this case because the variable is static. However, when two threads call this variable at once with different computer names, the use of static storage conflicts.

Thread-specific data could be used as a replacement for static storage, as in the `errno` problem, but this involves dynamic allocation of storage and adds to the expense of the call.

A better way to handle this kind of problem is to make the caller of `gethostbyname( )` supply the storage for the result of the call. This is done by having the caller supply an additional argument, an output argument, to the routine. This requires a new interface to `gethostbyname( )`.

This technique is used in threads to fix many of these problems. In most cases, the name of the new interface is the old name with “_r” appended, as in `gethostbyname_r(3N)`.

Synchronizing Threads

The threads in an application must cooperate and synchronize when sharing the data and the resources of the process.

A problem arises when multiple threads call something that manipulates an object. In a single-threaded world, synchronizing access to such objects is not a problem, but as Code Example 9-3 illustrates, this is a concern with multithreaded code. (Note that the `printf(3S)` function is safe to call for a multithreaded program; this example illustrates what could happen if `printf( )` were not safe.)

CODE EXAMPLE 9-3 The printf() Problem

```
/* thread 1: */
printf("go to statement reached");
```

```
/* thread 2: */
printf("hello world");
```

```
printed on display:
go to hello
```

Single-Threaded Strategy

One strategy is to have a single, application-wide mutex lock that is acquired whenever any thread in the application is running and is released before it must block. Since only one thread can be accessing shared data at any one time, each thread has a consistent view of memory.

Because this is effectively a single-threaded program, very little is gained by this strategy.

Reentrance

A better approach is to take advantage of the principles of modularity and data encapsulation. A reentrant function is one that behaves correctly if it is called simultaneously by several threads. Writing a reentrant function is a matter of understanding just what *behaves correctly* means for this particular function.

Functions that are callable by several threads must be made reentrant. This might require changes to the function interface or to the implementation.

Functions that access global state, like memory or files, have reentrance problems. These functions need to protect their use of global state with the appropriate synchronization mechanisms provided by threads.

The two basic strategies for making functions in modules reentrant are code locking and data locking.

Code Locking

Code locking is done at the function call level and guarantees that a function executes entirely under the protection of a lock. The assumption is that all access to data is done through functions. Functions that share data should execute under the same lock.

Some parallel programming languages provide a construct called a monitor that implicitly does code locking for functions that are defined within the scope of the monitor. A monitor can also be implemented by a mutex lock.

Functions under the protection of the same mutex lock or within the same monitor are guaranteed to execute atomically with respect to each other.

Data Locking

Data locking guarantees that access to a collection of data is maintained consistently. For data locking, the concept of locking code is still there, but code locking is around references to shared (global) data, only. For a mutual exclusion locking protocol, only one thread can be in the critical section for each collection of data.

Alternatively, in a multiple readers, single writer protocol, several readers can be allowed for each collection of data or one writer. Multiple threads can execute in a single module when they operate on different data collections and do not conflict on a single collection for the multiple readers, single writer protocol. So, data locking typically allows more concurrency than does code locking. (Note that Solaris threads has “Readers/Writer Lock” functionality built in.)

What strategy should you use when using locks (whether implemented with mutexes, condition variables, or semaphores) in a program? Should you try to achieve maximum parallelism by locking only when necessary and unlocking as soon as possible (fine-grained locking)? Or should you hold locks for long periods to minimize the overhead of taking and releasing them (coarse-grained locking)?

The granularity of the lock depends on the amount of data it protects. A very coarse-grained lock might be a single lock to protect all data. Dividing how the data is protected by the appropriate number of locks is very important. Too fine a grain of locking can degrade performance. The overhead associated with acquiring and releasing locks can become significant when there are too many locks.

The common wisdom is to start with a coarse-grained approach, identify bottlenecks, and add finer-grained locking where necessary to alleviate the bottlenecks. This is reasonably sound advice, but use your own judgment about finding the balance between maximizing parallelism and minimizing lock overhead.

Invariants

For both code locking and data locking, *invariants* are important to control locking complexity. An invariant is a condition or relation that is always true.

The definition is modified somewhat for concurrent execution: an invariant is a condition or relation that is true when the associated lock is being set. Once the lock is set, the invariant can be false. However, the code holding the lock must reestablish the invariant before releasing the lock.

An invariant can also be a condition or relation that is true when a lock is being set. Condition variables can be thought of as having an invariant that is the condition.

CODE EXAMPLE 9-4 Testing the Invariant With assert(3X)

```
mutex_lock(&lock);
while((condition)==FALSE)
    cond_wait(&cv,&lock);
assert((condition)==TRUE);
.
.
.
mutex_unlock(&lock);
```

The `assert( )` statement is testing the invariant. The `cond_wait( )` function does not preserve the invariant, which is why the invariant must be reevaluated when the thread returns.

Another example is a module that manages a doubly linked list of elements. For each item on the list a good invariant is the forward pointer of the previous item on the list that should also point to the same thing as the backward pointer of the forward item.

Assume this module uses code-based locking and therefore is protected by a single global mutex lock. When an item is deleted or added the mutex lock is acquired, the correct manipulation of the pointers is made, and the mutex lock is released. Obviously, at some point in the manipulation of the pointers the invariant is false, but the invariant is reestablished before the mutex lock is released.

Avoiding Deadlock

Deadlock is a permanent blocking of a set of threads that are competing for a set of resources. Just because some thread can make progress does not mean that there is not a deadlock somewhere else.

The most common error causing deadlock is *self deadlock* or *recursive deadlock*: a thread tries to acquire a lock it is already holding. Recursive deadlock is very easy to program by mistake.

For example, if a code monitor has every module function grabbing the mutex lock for the duration of the call, then any call between the functions within the module protected by the mutex lock immediately deadlocks. If a function calls some code outside the module which, through some circuitous path, calls back into any method protected by the same mutex lock, then it will deadlock too.

The solution for this kind of deadlock is to avoid calling functions outside the module when you don't know whether they will call back into the module without reestablishing invariants and dropping all module locks before making the call. Of course, after the call completes and the locks are reacquired, the state must be verified to be sure the intended operation is still valid.

An example of another kind of deadlock is when two threads, thread 1 and thread 2, each acquires a mutex lock, A and B, respectively. Suppose that thread 1 tries to acquire mutex lock B and thread 2 tries to acquire mutex lock A. Thread 1 cannot proceed and it is blocked waiting for mutex lock B. Thread 2 cannot proceed and it is blocked waiting for mutex lock A. Nothing can change, so this is a permanent blocking of the threads, and a deadlock.

This kind of deadlock is avoided by establishing an order in which locks are acquired (a *lock hierarchy*). When all threads always acquire locks in the specified order, this deadlock is avoided.

Adhering to a strict order of lock acquisition is not always optimal. When thread 2 has many assumptions about the state of the module while holding mutex lock B, giving up mutex lock B to acquire mutex lock A and then reacquiring mutex lock B

in order would cause it to discard its assumptions and reevaluate the state of the module.

The blocking synchronization primitives usually have variants that attempt to get a lock and fail if they cannot, such as `mutex_trylock()`. This allows threads to violate the lock hierarchy when there is no contention. When there is contention, the held locks must usually be discarded and the locks reacquired in order.

Deadlocks Related to Scheduling

Because there is no guaranteed order in which locks are acquired, a problem in threaded programs is that a particular thread never acquires a lock, even though it seems that it should.

This usually happens when the thread that holds the lock releases it, lets a small amount of time pass, and then reacquires it. Because the lock was released, it might seem that the other thread should acquire the lock. But, because nothing blocks the thread holding the lock, it continues to run from the time it releases the lock until it reacquires the lock, and so no other thread is run.

You can usually solve this type of problem by calling `thr_yield(3T)` just before the call to reacquire the lock. This allows other threads to run and to acquire the lock.

Because the time-slice requirements of applications are so variable, the threads library does not impose any. Use calls to `thr_yield()` to make threads share time as you require.

Locking Guidelines

Here are some simple guidelines for locking.

- Try not to hold locks across long operations like I/O where performance can be adversely affected.
- Don't hold locks when calling a function that is outside the module and that might reenter the module.
- In general, start with a coarse-grained approach, identify bottlenecks, and add finer-grained locking where necessary to alleviate the bottlenecks. Most locks are held for short amounts of time and contention is rare, so fix only those locks that have measured contention.
- When using multiple locks, avoid deadlocks by making sure that all threads acquire the locks in the same order.

Following Some Basic Guidelines

- Know what you are importing and whether it is safe.

A threaded program cannot arbitrarily enter nonthreaded code.
- Threaded code can safely refer to unsafe code only from the initial thread.

This ensures that the static storage associated with the initial thread is used only by that thread.
- Sun-supplied libraries are assumed to be *unsafe* unless explicitly documented as *safe*.

If a reference manual entry does not state explicitly that an interface is *MT-Safe*, you should assume that the interface is *unsafe*.
- Use compilation flags to manage binary incompatible source changes. (See Chapter 7, Compiling and Debugging, for complete instructions.)
 - `-D_REENTRANT` enables multithreading with the `-lthread` library.
 - `-D_POSIX_C_SOURCE` with `-lthread` gives POSIX threads behavior.
 - `-D_POSIX_PTHREADS_SEMANTICS` with `-lthread` gives both Solaris threads and pthreads interfaces with a preference given to the POSIX interfaces when the two interfaces conflict.
- When making a library safe for multithreaded use, do not thread global process operations.

Do not change global operations (or actions with global side effects) to behave in a threaded manner. For example, if file I/O is changed to per-thread operation, threads cannot cooperate in accessing files.

For thread-specific behavior, or *thread cognizant* behavior, use thread facilities. For example, when the termination of `main( )` should terminate only the thread that is exiting `main( )`, the end of `main( )` should be:

```
thr_exit();  
/*NOTREACHED*/
```

Creating and Using Threads

The threads packages will cache the threads data structure, stacks, and LWPs so that the repetitive creation of unbound threads can be inexpensive.

Unbound thread creation has considerable overhead when compared to process creation or even to bound thread creation. In fact, the overhead is similar to unbound thread synchronization when you include the context switches to stop one thread and start another.

So, creating and destroying threads as they are required is usually better than attempting to manage a pool of threads that wait for independent work.

A good example of this is an RPC server that creates a thread for each request and destroys it when the reply is delivered, instead of trying to maintain a pool of threads to service requests.

While thread creation has less overhead compared to that of process creation, it is not efficient when compared to the cost of a few instructions. Create threads for processing that lasts at least a couple of thousand machine instructions.

Lightweight Processes

Figure 9–1 illustrates the relationship between LWPs and the user and kernel levels.




Figure 9–1 Multithreading Levels and Relationships

The user-level threads library, with help from the programmer and the operating environment, ensures that the number of LWPs available is adequate for the currently active user-level threads. However, there is no one-to-one mapping between user threads and LWPs, and user-level threads can freely migrate from one LWP to another.

With Solaris threads, a programmer can tell the threads library how many threads should be “running” at the same time.

For example, if the programmer says that up to three threads should run at the same time, then at least three LWPs should be available. If there are three available processors, the threads run in parallel. If there is only one processor, then the operating environment multiplexes the three LWPs on that one processor. If all the LWPs block, the threads library adds another LWP to the pool.

When a user thread blocks due to synchronization, its LWP transfers to another runnable thread. This transfer is done with a coroutine linkage and not with a system call.

The operating environment decides which LWP should run on which processor and when. It has no knowledge about what user threads are or how many are active in each process.

The kernel schedules LWPs onto CPU resources according to their scheduling classes and priorities. The threads library schedules threads on the process pool of LWPs in much the same way.

Each LWP is independently dispatched by the kernel, performs independent system calls, incurs independent page faults, and runs in parallel on a multiprocessor system.

An LWP has some capabilities that are not exported directly to threads, such as a special scheduling class.

Unbound Threads

The library invokes LWPs as needed and assigns them to execute runnable threads. The LWP assumes the state of the thread and executes its instructions. If the thread becomes blocked on a synchronization mechanism, or if another thread should be run, the thread state is saved in process memory and the threads library assigns another thread to the LWP to run.

Bound Threads

Sometimes having more threads than LWPs, as can happen with unbound threads, is a disadvantage.

For example, a parallel array computation divides the rows of its arrays among different threads. If there is one LWP for each processor, but multiple threads for each LWP, each processor spends time switching between threads. In this case, it is better to have one thread for each LWP, divide the rows among a smaller number of threads, and reduce the number of thread switches.

A mixture of threads that are permanently bound to LWPs and unbound threads is also appropriate for some applications.

An example of this is a realtime application that has some threads with system-wide priority and realtime scheduling, and other threads that attend to background computations. Another example is a window system with unbound threads for most operations and a mouse serviced by a high-priority, bound, realtime thread.

When a user-level thread issues a system call, the LWP running the thread calls into the kernel and remains attached to the thread at least until the system call completes.

Bound threads require more overhead than unbound threads. Because bound threads can change the attributes of the underlying LWP, the LWPs are not cached when the bound threads exit. Instead, the operating environment provides a new LWP when a bound thread is created and destroys it when the bound thread exits.

Use bound threads only when a thread needs resources that are available only through the underlying LWP, such as a virtual time interval timer or an alternate stack, or when the thread must be visible to the kernel to be scheduled with respect to all other active threads in the system, as in realtime scheduling.

Use unbound threads even when you expect all threads to be active simultaneously. This enables Solaris threads to cache LWP and thread resources efficiently so that thread creation and destruction are fast. Use `thr_setconcurrency(3T)` to tell Solaris threads how many threads you expect to be simultaneously active.

Thread Concurrency (Solaris Threads Only)

By default, Solaris threads attempts to adjust the system execution resources (LWPs) used to run unbound threads to match the real number of active threads. While the Solaris threads package cannot make perfect decisions, it at least ensures that the process continues to make progress.

When you have some idea of the number of unbound threads that should be simultaneously active (executing code or system calls), tell the library through `thr_setconcurrency(3THR)`.

For example:

- A database server that has a thread for each user should tell Solaris threads the expected number of simultaneously active users.
- A window server that has one thread for each client should tell Solaris threads the expected number of simultaneously active clients.
- A file copy program that has one reader thread and one writer thread should tell Solaris threads that the desired concurrency level is two.

Alternatively, the concurrency level can be incremented by one through the `THR_NEW_LWP` flag as each thread is created.

Include unbound threads blocked on interprocess (`USYNC_PROCESS`) synchronization variables as active when you compute thread concurrency. Exclude

bound threads—they do not require concurrency support from Solaris threads because they are equivalent to LWPs.

Efficiency

A new thread is created with `thr_create(3T)` in less time than an existing thread can be restarted. This means that it is more efficient to create a new thread when one is needed and have it call `thr_exit(3T)` when it has completed its task than it would be to stockpile an idle thread and restart it.

Thread Creation Guidelines

Here are some simple guidelines for using threads.

- Use threads for independent activities that must do a meaningful amount of work.
- Use Solaris threads to take advantage of CPU concurrency.
- Use bound threads only when absolutely necessary, that is, when some facility of the underlying LWP is required.

Working With Multiprocessors

Multithreading lets you take advantage of multiprocessors, primarily through parallelism and scalability. Programmers should be aware of the differences between the memory models of a multiprocessor and a uniprocessor.

Memory consistency is directly interrelated to the processor interrogating memory. For uniprocessors, memory is obviously consistent because there is only one processor viewing memory.

To improve multiprocessor performance, memory consistency is relaxed. You cannot always assume that changes made to memory by one processor are immediately reflected in the other processors' views of that memory.

You can avoid this complexity by using synchronization variables when you use shared or global variables.

Barrier synchronization is sometimes an efficient way to control parallelism on multiprocessors. An example of barriers can be found in Appendix B, Solaris Threads Example: `barrier.c`.

Another multiprocessor issue is efficient synchronization when threads must wait until all have reached a common point in their execution.

Note - The issues discussed here are not important when the threads synchronization primitives are always used to access shared memory locations.

The Underlying Architecture

When threads synchronize access to shared storage locations using the threads synchronization routines, the effect of running a program on a shared-memory multiprocessor is identical to the effect of running the program on a uniprocessor.

However, in many situations a programmer might be tempted to take advantage of the multiprocessor and use “tricks” to avoid the synchronization routines. As Code Example 9–5 and Code Example 9–6 show, such tricks can be dangerous.

Understanding the memory models supported by common multiprocessor architectures helps to understand the dangers.

The major multiprocessor components are:

- The processors themselves
- Store buffers, which connect the processors to their caches
- *Caches*, which hold the contents of recently accessed or modified storage locations
- memory, which is the primary storage (and is shared by all processors).

In the simple traditional model, the multiprocessor behaves as if the processors are connected directly to memory: when one processor stores into a location and another immediately loads from the same location, the second processor loads what was stored by the first.

Caches can be used to speed the average memory access, and the desired semantics can be achieved when the caches are kept consistent with one another.

A problem with this simple approach is that the processor must often be delayed to make certain that the desired semantics are achieved. Many modern multiprocessors use various techniques to prevent such delays, which, unfortunately, change the semantics of the memory model.

Two of these techniques and their effects are explained in the next two examples.

“Shared-Memory” Multiprocessors

Consider the purported solution to the producer/consumer problem shown in Code Example 9–5.

Although this program works on current SPARC-based multiprocessors, it assumes that all multiprocessors have strongly ordered memory. This program is therefore not portable.

CODE EXAMPLE 9-5 The Producer/Consumer Problem—Shared Memory Multiprocessors

```
char buffer[BSIZE];
unsigned int in = 0;
unsigned int out = 0;

void producer(char item) {
    do
        /* nothing */
    while
        (in - out == BSIZE);

    buffer[in%BSIZE] = item;
    in++;
}

char consumer(void) {
    char item;
    do
        /* nothing */
    while
        (in - out == 0);
    item = buffer[out%BSIZE];
    out++;
}
```

When this program has exactly one producer and exactly one consumer and is run on a shared-memory multiprocessor, it appears to be correct. The difference between *in* and *out* is the number of items in the buffer.

The producer waits (by repeatedly computing this difference) until there is room for a new item, and the consumer waits until there is an item in the buffer.

For memory that is *strongly ordered* (for instance, a modification to memory on one processor is immediately available to the other processors), this solution is correct (it is correct even taking into account that *in* and *out* will eventually overflow, as long as *B*SIZE is less than the largest integer that can be represented in a word).

Shared-memory multiprocessors do not necessarily have strongly ordered memory. A change to memory by one processor is not necessarily available immediately to the other processors. When two changes to different memory locations are made by one processor, the other processors do not necessarily detect the changes in the order in which they were made because changes to memory do not happen immediately.

First the changes are stored in *store buffers* that are not visible to the cache.

The processor checks these store buffers to ensure that a program has a consistent view, but because store buffers are not visible to other processors, a write by one processor does not become visible until it is written to cache.

The synchronization primitives (see Chapter 4, Programming With Synchronization Objects) use special instructions that flush the store buffers to cache. So, using locks around your shared data ensures memory consistency.

When memory ordering is very relaxed, Code Example 9-5 has a problem because the consumer might see that *in* has been incremented by the producer before it sees the change to the corresponding buffer slot.

This is called *weak ordering* because stores made by one processor can appear to happen out of order by another processor (memory, however, is always consistent from the same processor). To fix this, the code should use mutexes to flush the cache.

The trend is toward relaxing memory order. Because of this, programmers are becoming increasingly careful to use locks around all global or shared data.

As demonstrated by Code Example 9–5 and Code Example 9–6, locking is essential.

Peterson’s Algorithm

The code in Code Example 9–6 is an implementation of Peterson’s Algorithm, which handles mutual exclusion between two threads. This code tries to guarantee that there is never more than one thread in the critical section and that, when a thread calls `mut_excl( )`, it enters the critical section sometime “soon.”

An assumption here is that a thread exits fairly quickly after entering the critical section.

CODE EXAMPLE 9–6 Mutual Exclusion for Two Threads?

```
void mut_excl(int me /* 0 or 1 */) {
    static int loser;
    static int interested[2] = {0, 0};
    int other; /* local variable */

    other = 1 - me;
    interested[me] = 1;
    loser = me;
    while (loser == me && interested[other])
        ;

    /* critical section */
    interested[me] = 0;
}
```

This algorithm works some of the time when it is assumed that the multiprocessor has strongly ordered memory.

Some multiprocessors, including some SPARC-based multiprocessors, have store buffers. When a thread issues a store instruction, the data is put into a store buffer. The buffer contents are eventually sent to the cache, but not necessarily right away. (Note that the caches on each of the processors maintain a consistent view of memory, but modified data does not reach the cache right away.)

When multiple memory locations are stored into, the changes reach the cache (and memory) in the correct order, but possibly after a delay. SPARC-based multiprocessors with this property are said to have total store order (TSO).

When one processor stores into location A and then loads from location B, and another processor stores into location B and loads from location A, the expectation is that either the first processor fetches the newly modified value in location B or the second processor fetches the newly modified value in location A, or both. However, the case in which both processors load the old values simply cannot happen.

Moreover, with the delays caused by load and store buffers, the “impossible case” can happen.

What could happen with Peterson’s algorithm is that two threads running on separate processors each stores into its own slot of the particular array and then loads from the other slot. They both read the old values (0), assume that the other party is not present, and both enter the critical section. (Note that this is the sort of problem that might not occur when you test a program, but only much later.)

This problem is avoided when you use the threads synchronization primitives, whose implementations issue special instructions to force the writing of the store buffers to the cache.

Parallelizing a Loop on a Shared-Memory Parallel Computer

In many applications, and especially numerical applications, while part of the algorithm can be parallelized, other parts are inherently sequential, as shown in the following:

Thread ₁	Thread ₂ through Thread _n
<pre>while(many_iterations) { sequential_computation --- Barrier --- parallel_computation }</pre>	<pre>while(many_iterations) { --- Barrier --- parallel_computation }</pre>

For example, you might produce a set of matrixes with a strictly linear computation, then perform operations on the matrixes using a parallel algorithm, then use the results of these operations to produce another set of matrixes, then operate on them in parallel, and so on.

The nature of the parallel algorithms for such a computation is that little synchronization is required during the computation, but synchronization of all the threads employed is required to ensure that the sequential computation is finished before the parallel computation begins.

The barrier forces all the threads that are doing the parallel computation to wait until all threads involved have reached the barrier. When they’ve reached the barrier, they are released and begin computing together.

Summary

This guide has covered a wide variety of important threads programming issues. Look in Appendix A, Sample Application—Multithreaded `grep`, for a `pthread` program example that uses many of the features and styles that have been discussed. Look in Appendix B, Solaris Threads Example, for a program example that uses Solaris threads.

Further Reading

For more in-depth information about multithreading, see the following book:

- *Programming with Threads* by Steve Kleiman, Devang Shah, and Bart Smaalders (Prentice-Hall, published in 1995)

Sample Application—Multithreaded grep

Description of `tgrep`

The `tgrep` sample program is a multithreaded version of `find(1)` combined with `grep(1)`. `tgrep` supports all but the `-w` (word search) options of the normal `grep`, and a few exclusively available options.

By default, the `tgrep` searches are like the following command:

```
find . -exec grep [ options ] pattern {} \;
```

For large directory hierarchies, `tgrep` gets results more quickly than the `find` command, depending on the number of processors available. On uniprocessor machines it is about twice as fast, and on four processor machines it is about four times as fast.

The `-e` option changes the way `tgrep` interprets the pattern string. Ordinarily (without the `-e` option) `tgrep` uses a literal string match. With the `-e` option, `tgrep` uses an MT-Safe public domain version of a regular expression handler. The regular expression method is slower.

The `-B` option tells `tgrep` to use the value of the environment variable called `TGLIMIT` to limit the number of threads it will use during a search. This option has no affect if `TGLIMIT` is not set. Because `tgrep` can use a lot of system resources, this is a way to run it politely on a timesharing system.

Getting Online Source Code

Source for `tgrep` is included on the Catalyst Developer's CD. Contact your sales representative to find out how you can get a copy.

Only the multithreaded `main.c` module appears here. Other modules, including those for regular expression handling, plus documentation and Makefiles, are available on the Catalyst Developer's CD.

CODE EXAMPLE A-1 Source Code for `tgrep` Program

```
/* Copyright (c) 1993, 1994 Ron Winacott */
/* This program may be used, copied, modified, and redistributed freely */
/* for ANY purpose, so long as this notice remains intact. */

#define _REENTRANT

#include <stdio.h>
#include <string.h>
#include <stdlib.h>
#include <unistd.h>
#include <assert.h>
#include <errno.h>
#include <ctype.h>
#include <sys/types.h>
#include <time.h>
#include <sys/stat.h>
#include <dirent.h>

#include "version.h"

#include <fcntl.h>
#include <sys/uio.h>
#include <pthread.h>
#include <sched.h>

#ifdef MARK
#include <prof.h> /* to turn on MARK(), use -DMARK to compile (see man prof5) */
#endif

#include "pmatch.h"

#define PATH_MAX          1024 /* max # of characters in a path name */
#define HOLD_FDS          6   /* stdin,out,err and a buffer */
#define UNLIMITED        99999 /* The default tglimit */
#define MAXREGEXP         10   /* max number of -e options */

#define FB_BLOCK          0x00001
#define FC_COUNT          0x00002
#define FH_HOLDNAME      0x00004
#define FI_IGNCASE       0x00008
#define FL_NAMEONLY      0x00010
#define FN_NUMBER        0x00020
#define FS_NOERROR       0x00040
#define FV_REVERSE       0x00080
```

```

#define FW_WORD            0x00100
#define FR_RECUR           0x00200
#define FU_UNSORT          0x00400
#define FX_STDIN           0x00800
#define TG_BATCH           0x01000
#define TG_FILEPAT         0x02000
#define FE_REGEX           0x04000
#define FS_STATS           0x08000
#define FC_LINE            0x10000
#define TG_PROGRESS        0x20000

#define FILET              1
#define DIRT               2

typedef struct work_st {
    char        *path;
    int         tp;
    struct work_st *next;
} work_t;

typedef struct out_st {
    char        *line;
    int         line_count;
    long        byte_count;
    struct out_st *next;
} out_t;

#define ALPHASIZ           128
typedef struct bm_pattern { /* Boyer - Moore pattern */
    short       p_m;        /* length of pattern string */
    short       p_r[ALPHASIZ]; /* "r" vector */
    short       *p_R;        /* "R" vector */
    char        *p_pat;      /* pattern string */
} BM_PATTERN;

/* bmpmatch.c */
extern BM_PATTERN *bm_makepat(char *p);
extern char *bm_pmatch(BM_PATTERN *pat, register char *s);
extern void bm_freepat(BM_PATTERN *pattern);
BM_PATTERN      *bm_pat; /* the global target read only after main */

/* pmatch.c */
extern char *pmatch(register PATTERN *pattern, register char *string, int *len);
extern PATTERN *makepat(char *string, char *metas);
extern void freepat(register PATTERN *pat);
extern void printpat(PATTERN *pat);
PATTERN      *pm_pat[MAXREGEXP]; /* global targets read only for pmatch */

#include "proto.h" /* function prototypes of main.c */

/* local functions to POSIX only */
void pthread_setconcurrency_np(int con);
int pthread_getconcurrency_np(void);
void pthread_yield_np(void);

pthread_attr_t detached_attr;
pthread_mutex_t output_print_lk;
pthread_mutex_t global_count_lk;

int      global_count = 0;

```

```

work_t          *work_q = NULL;
pthread_cond_t  work_q_cv;
pthread_mutex_t work_q_lk;
pthread_mutex_t debug_lock;

#include "debug.h" /* must be included AFTER the
                  mutex_t debug_lock line */

work_t          *search_q = NULL;
pthread_mutex_t  search_q_lk;
pthread_cond_t  search_q_cv;
int             search_pool_cnt = 0; /* the count in the pool now */
int             search_thr_limit = 0; /* the max in the pool */

work_t          *cascade_q = NULL;
pthread_mutex_t  cascade_q_lk;
pthread_cond_t  cascade_q_cv;
int             cascade_pool_cnt = 0;
int             cascade_thr_limit = 0;

int             running = 0;
pthread_mutex_t  running_lk;

pthread_mutex_t  stat_lk;
time_t          st_start = 0;
int             st_dir_search = 0;
int             st_file_search = 0;
int             st_line_search = 0;
int             st_cascade = 0;
int             st_cascade_pool = 0;
int             st_cascade_destroy = 0;
int             st_search = 0;
int             st_pool = 0;
int             st_maxrun = 0;
int             st_worknull = 0;
int             st_workfds = 0;
int             st_worklimit = 0;
int             st_destroy = 0;

int             all_done = 0;
int             work_cnt = 0;
int             current_open_files = 0;
int             tglimit = UNLIMITED; /* if -B limit the number of
                                     threads */

int             progress_offset = 1;
int             progress = 0; /* protected by the print_lock ! */
unsigned int     flags = 0;
int             regexp_cnt = 0;
char            *string[MAXREGEXP];
int             debug = 0;
int             use_pmatch = 0;
char            file_pat[255]; /* file patten match */
PATTERN         *pm_file_pat; /* compiled file target string (pmatch()) */

/*
 * Main: This is where the fun starts
 */
int
main(int argc, char **argv)

```

```

{
    int          c,out_thr_flags;
    long         max_open_files = 01, ncpus = 01;
    extern int   optind;
    extern char  *optarg;
    int         prio = 0;
    struct stat  sbuf;
    pthread_t    tid,dtid;
    void         *status;
    char         *e = NULL, *d = NULL; /* for debug flags */
    int         debug_file = 0;
    struct sigaction sigact;
    sigset_t     set,oset;
    int         err = 0, i = 0, pm_file_len = 0;
    work_t       *work;
    int         restart_cnt = 10;

    /* NO OTHER THREADS ARE RUNNING */
    flags = FR_RECUR; /* the default */

    while ((c = getopt(argc, argv, "d:e:bchilnsvwruf:p:BCSZzHP:")) != EOF) {
        switch (c) {
#ifdef DEBUG
            case 'd':
                debug = atoi(optarg);
                if (debug == 0)
                    debug_usage();

                d = optarg;
                fprintf(stderr,"tgrep: Debug on at level(s) ");
                while (*d) {
                    for (i=0; i<9; i++)
                        if (debug_set[i].level == *d) {
                            debug_levels |= debug_set[i].flag;
                            fprintf(stderr,"%c ",debug_set[i].level);
                            break;
                        }
                    d++;
                }
                fprintf(stderr,"\n");
                break;
            case 'f': debug_file = atoi(optarg); break;
#endif
            case 'B':
                flags |= TG_BATCH;
#ifdef __lock_lint
            /* locklint complains here, but there are no other threads */
                if ((e = getenv("TGLIMIT")) != NULL) {
                    tglimit = atoi(e);
                }
            else {
                if (!(flags & FS_NOERROR)) /* order dependent! */
                    fprintf(stderr,"env TGLIMIT not set, overriding -B\n");
                flags &= ~TG_BATCH;
            }
#endif
                break;
            case 'p':
                flags |= TG_FILEPAT;

```

```

        strcpy(file_pat,optarg);
        pm_file_pat = makepat(file_pat,NULL);
        break;
    case 'P':
        flags |= TG_PROGRESS;
        progress_offset = atoi(optarg);
        break;
    case 'S': flags |= FS_STATS;      break;
    case 'b': flags |= FB_BLOCK;      break;
    case 'c': flags |= FC_COUNT;      break;
    case 'h': flags |= FH_HOLDNAME;   break;
    case 'i': flags |= FI_IGNCASE;     break;
    case 'l': flags |= FL_NAMEONLY;    break;
    case 'n': flags |= FN_NUMBER;      break;
    case 's': flags |= FS_NOERROR;     break;
    case 'v': flags |= FV_REVERSE;     break;
    case 'w': flags |= FW_WORD;        break;
    case 'r': flags &= ~FR_RECUR;      break;
    case 'C': flags |= FC_LINE;        break;
    case 'e':
        if (regexp_cnt == MAXREGEXP) {
            fprintf(stderr,"Max number of regexp's (%d) exceeded!\n",
                MAXREGEXP);
            exit(1);
        }
        flags |= FE_REGEXP;
        if ((string[regexp_cnt] =(char *)malloc(strlen(optarg)+1))==NULL){
            fprintf(stderr,"tgrep: No space for search string(s)\n");
            exit(1);
        }
        memset(string[regexp_cnt],0,strlen(optarg)+1);
        strcpy(string[regexp_cnt],optarg);
        regexp_cnt++;
        break;
    case 'z':
    case 'Z': regexp_usage();
        break;
    case 'H':
    case '?':
    default : usage();
        }
    }
    if (flags & FS_STATS)
        st_start = time(NULL);

    if (!(flags & FE_REGEXP)) {
        if (argc - optind < 1) {
            fprintf(stderr,"tgrep: Must supply a search string(s) "
                "and file list or directory\n");
            usage();
        }
        if ((string[0]=(char *)malloc(strlen(argv[optind])+1))==NULL){
            fprintf(stderr,"tgrep: No space for search string(s)\n");
            exit(1);
        }
        memset(string[0],0,strlen(argv[optind])+1);
        strcpy(string[0],argv[optind]);
        regexp_cnt=1;
        optind++;
    }
}

```

```

if (flags & FI_IGNORECASE)
    for (i=0; i<regexp_cnt; i++)
        uncase(string[i]);

if (flags & FE_REGEX) {
    for (i=0; i<regexp_cnt; i++)
        pm_pat[i] = makepat(string[i],NULL);
    use_pmatch = 1;
}
else {
    bm_pat = bm_makepat(string[0]); /* only one allowed */
}

flags |= FX_STDIN;

max_open_files = sysconf(_SC_OPEN_MAX);
ncpus = sysconf(_SC_NPROCESSORS_ONLN);
if ((max_open_files - HOLD_FDS - debug_file) < 1) {
    fprintf(stderr,"tgrep: You MUST have at least ONE fd "
        "that can be used, check limit (>10)\n");
    exit(1);
}
search_thr_limit = max_open_files - HOLD_FDS - debug_file;
cascade_thr_limit = search_thr_limit / 2;
/* the number of files that can be open */
current_open_files = search_thr_limit;

pthread_attr_init(&detached_attr);
pthread_attr_setdetachstate(&detached_attr,
    PTHREAD_CREATE_DETACHED);

pthread_mutex_init(&global_count_lk,NULL);
pthread_mutex_init(&output_print_lk,NULL);
pthread_mutex_init(&work_q_lk,NULL);
pthread_mutex_init(&running_lk,NULL);
pthread_cond_init(&work_q_cv,NULL);
pthread_mutex_init(&search_q_lk,NULL);
pthread_cond_init(&search_q_cv,NULL);
pthread_mutex_init(&cascade_q_lk,NULL);
pthread_cond_init(&cascade_q_cv,NULL);

if ((argc == optind) && ((flags & TG_FILEPAT) || (flags & FR_RECUR))) {
    add_work(".",DIRT);
    flags = (flags & ~FX_STDIN);
}
for ( ; optind < argc; optind++) {
    restart_cnt = 10;
    flags = (flags & ~FX_STDIN);
    STAT_AGAIN:
    if (stat(argv[optind], &sbuf)) {
        if (errno == EINTR) { /* try again !, restart */
            if (--restart_cnt)
                goto STAT_AGAIN;
        }
        if (!(flags & FS_NOERROR))
            fprintf(stderr,"tgrep: Can't stat file/dir %s, %s\n",
                argv[optind], strerror(errno));
        continue;
    }
}

```

```

    }
    switch (sbuf.st_mode & S_IFMT) {
    case S_IFREG :
        if (flags & TG_FILEPAT) {
            if (pmatch(pm_file_pat, argv[optind], &pm_file_len))
                DP(DLEVEL1, "File pat match %s\n", argv[optind]));
            add_work(argv[optind], FILET);
        }
        else {
            add_work(argv[optind], FILET);
        }
        break;
    case S_IFDIR :
        if (flags & FR_RECUR) {
            add_work(argv[optind], DIRT);
        }
        else {
            if (!(flags & FS_NOERROR))
                fprintf(stderr, "tgrep: Can't search directory %s, "
                    "-r option is on. Directory ignored.\n",
                    argv[optind]);
        }
        break;
    }
}

pthread_setconcurrency_np(3);

if (flags & FX_STDIN) {
    fprintf(stderr, "tgrep: stdin option is not coded at this time\n");
    exit(0);
    /* XXX Need to fix this SOON */
    search_thr(NULL);
    if (flags & FC_COUNT) {
        pthread_mutex_lock(&global_count_lk);
        printf("%d\n", global_count);
        pthread_mutex_unlock(&global_count_lk);
    }
    if (flags & FS_STATS)
        prnt_stats();
    exit(0);
}

pthread_mutex_lock(&work_q_lk);
if (!work_q) {
    if (!(flags & FS_NOERROR))
        fprintf(stderr, "tgrep: No files to search.\n");
    exit(0);
}
pthread_mutex_unlock(&work_q_lk);

DP(DLEVEL1, "Starting to loop through the work_q for work\n");

/* OTHER THREADS ARE RUNNING */
while (1) {
    pthread_mutex_lock(&work_q_lk);
    while ((work_q == NULL || current_open_files == 0 || tglimit <= 0) &&
        all_done == 0) {
        if (flags & FS_STATS) {
            pthread_mutex_lock(&stat_lk);
            if (work_q == NULL)

```

```

        st_worknull++;
        if (current_open_files == 0)
            st_workfds++;
        if (tglimit <= 0)
            st_worklimit++;
        pthread_mutex_unlock(&stat_lk);
    }
    pthread_cond_wait(&work_q_cv, &work_q_lk);
}
if (all_done != 0) {
    pthread_mutex_unlock(&work_q_lk);
    DP(DLEVEL1, ("All_done was set to TRUE\n"));
    goto OUT;
}
work = work_q;
work_q = work->next; /* maybe NULL */
work->next = NULL;
current_open_files--;
pthread_mutex_unlock(&work_q_lk);

tid = 0;
switch (work->tp) {
case DIRT:
    pthread_mutex_lock(&cascade_q_lk);
    if (cascade_pool_cnt) {
        if (flags & FS_STATS) {
            pthread_mutex_lock(&stat_lk);
            st_cascade_pool++;
            pthread_mutex_unlock(&stat_lk);
        }
        work->next = cascade_q;
        cascade_q = work;
        pthread_cond_signal(&cascade_q_cv);
        pthread_mutex_unlock(&cascade_q_lk);
        DP(DLEVEL2, ("Sent work to cascade pool thread\n"));
    }
    else {
        pthread_mutex_unlock(&cascade_q_lk);
        err = pthread_create(&tid, &detached_attr, cascade, (void *)work);
        DP(DLEVEL2, ("Sent work to new cascade thread\n"));
        if (flags & FS_STATS) {
            pthread_mutex_lock(&stat_lk);
            st_cascade++;
            pthread_mutex_unlock(&stat_lk);
        }
    }
    break;
case FILET:
    pthread_mutex_lock(&search_q_lk);
    if (search_pool_cnt) {
        if (flags & FS_STATS) {
            pthread_mutex_lock(&stat_lk);
            st_pool++;
            pthread_mutex_unlock(&stat_lk);
        }
        work->next = search_q; /* could be null */
        search_q = work;
        pthread_cond_signal(&search_q_cv);
        pthread_mutex_unlock(&search_q_lk);
        DP(DLEVEL2, ("Sent work to search pool thread\n"));
    }
}

```

```

    }
    else {
        pthread_mutex_unlock(&search_q_lk);
        err = pthread_create(&tid,&detached_attr,
                           search_thr,(void *)work);
        pthread_setconcurrency_np(pthread_getconcurrency_np()+1);
        DP(DLEVEL2,"Sent work to new search thread\n");
        if (flags & FS_STATS) {
            pthread_mutex_lock(&stat_lk);
            st_search++;
            pthread_mutex_unlock(&stat_lk);
        }
    }
    break;
default:
    fprintf(stderr,"tgrep: Internal error, work_t->tp not valid\n");
    exit(1);
}
if (err) { /* NEED TO FIX THIS CODE. Exiting is just wrong */
    fprintf(stderr,"Could not create new thread!\n");
    exit(1);
}
}

OUT:
if (flags & TG_PROGRESS) {
    if (progress)
        fprintf(stderr,".\n");
    else
        fprintf(stderr,"\n");
}
/* we are done, print the stuff. All other threads are parked */
if (flags & FC_COUNT) {
    pthread_mutex_lock(&global_count_lk);
    printf("%d\n",global_count);
    pthread_mutex_unlock(&global_count_lk);
}
if (flags & FS_STATS)
    prnt_stats();
return(0); /* should have a return from main */
}

/*
 * Add_Work: Called from the main thread, and cascade threads to add file
 * and directory names to the work Q.
 */
int
add_work(char *path,int tp)
{
    work_t      *wt,*ww,*wp;

    if ((wt = (work_t *)malloc(sizeof(work_t))) == NULL)
        goto ERROR;
    if ((wt->path = (char *)malloc(strlen(path)+1)) == NULL)
        goto ERROR;

    strcpy(wt->path,path);
    wt->tp = tp;
    wt->next = NULL;
    if (flags & FS_STATS) {

```

```

        pthread_mutex_lock(&stat_lk);
        if (wt->tp == DIRT)
            st_dir_search++;
        else
            st_file_search++;
        pthread_mutex_unlock(&stat_lk);
    }
    pthread_mutex_lock(&work_q_lk);
    work_cnt++;
    wt->next = work_q;
    work_q = wt;
    pthread_cond_signal(&work_q_cv);
    pthread_mutex_unlock(&work_q_lk);
    return(0);
ERROR:
    if (!(flags & FS_NOERROR))
        fprintf(stderr, "tgrep: Could not add %s to work queue. Ignored\n",
            path);
    return(-1);
}

/*
 * Search thread: Started by the main thread when a file name is found
 * on the work Q to be searched. If all the needed resources are ready
 * a new search thread will be created.
 */
void *
search_thr(void *arg) /* work_t *arg */
{
    FILE          *fin;
    char          fin_buf[(BUFSIZ*4)]; /* 4 Kbytes */
    work_t        *wt, std;
    int           line_count;
    char          rline[128];
    char          cline[128];
    char          *line;
    register char *p,*pp;
    int           pm_len;
    int           len = 0;
    long          byte_count;
    long          next_line;
    int           show_line; /* for the -v option */
    register int  slen,plen,i;
    out_t         *out = NULL; /* this threads output list */

    pthread_yield_np();
    wt = (work_t *)arg; /* first pass, wt is passed to use. */

    /* len = strlen(string); */ /* only set on first pass */

    while (1) { /* reuse the search threads */
        /* init all back to zero */
        line_count = 0;
        byte_count = 0;
        next_line = 0;
        show_line = 0;

        pthread_mutex_lock(&running_lk);
        running++;
        pthread_mutex_unlock(&running_lk);
    }

```

```

pthread_mutex_lock(&work_q_lk);
tglimit--;
pthread_mutex_unlock(&work_q_lk);
DP(DLEVEL5, "searching file (STDIO) %s\n", wt->path));

if ((fin = fopen(wt->path, "r")) == NULL) {
    if (!(flags & FS_NOERROR)) {
        fprintf(stderr, "tgrep: %s. File \"%s\" not searched.\n",
            strerror(errno), wt->path);
    }
    goto ERROR;
}
setvbuf(fin, fin_buf, _IOFBF, (BUFSIZ*4)); /* XXX */
DP(DLEVEL5, "Search thread has opened file %s\n", wt->path));
while ((fgets(rline, 127, fin)) != NULL) {
    if (flags & FS_STATS) {
        pthread_mutex_lock(&stat_lk);
        st_line_search++;
        pthread_mutex_unlock(&stat_lk);
    }
    slen = strlen(rline);
    next_line += slen;
    line_count++;
    if (rline[slen-1] == '\n')
        rline[slen-1] = '\0';
    /*
    ** If the uncase flag is set, copy the read in line (rline)
    ** To the uncase line (cline) Set the line pointer to point at
    ** cline.
    ** If the case flag is NOT set, then point line at rline.
    ** line is what is compared, rline is what is printed on a
    ** match.
    */
    if (flags & FI_IGNCASE) {
        strcpy(cline, rline);
        uncase(cline);
        line = cline;
    }
    else {
        line = rline;
    }
    show_line = 1; /* assume no match, if -v set */
    /* The old code removed */
    if (use_pmatch) {
        for (i=0; i<regexp_cnt; i++) {
            if (pmatch(pm_pat[i], line, &pm_len)) {
                if (!(flags & FV_REVERSE)) {
                    add_output_local(&out, wt, line_count,
                        byte_count, rline);
                    continue_line(rline, fin, out, wt,
                        &line_count, &byte_count);
                }
                else {
                    show_line = 0;
                }
            } /* end of if -v flag if / else block */
            /*
            ** if we get here on ANY of the regexp targets
            ** jump out of the loop, we found a single
            ** match so do not keep looking!
            ** If name only, do not keep searching the same

```

```

        ** file, we found a single match, so close the file,
        ** print the file name and move on to the next file.
        */
        if (flags & FL_NAMEONLY)
            goto OUT_OF_LOOP;
        else
            goto OUT_AND_DONE;
    } /* end found a match if block */
} /* end of the for pat[s] loop */
}
else {
    if (bm_pmatch( bm_pat, line)) {
        if (!(flags & FV_REVERSE)) {
            add_output_local(&out,wt,line_count,byte_count,rline);
            continue_line(rline,fin,out,wt,
                          &line_count,&byte_count);
        }
        else {
            show_line = 0;
        }
        if (flags & FL_NAMEONLY)
            goto OUT_OF_LOOP;
    }
}
OUT_AND_DONE:
if ((flags & FV_REVERSE) && show_line) {
    add_output_local(&out,wt,line_count,byte_count,rline);
    show_line = 0;
}
byte_count = next_line;
}
OUT_OF_LOOP:
fclose(fin);
/*
** The search part is done, but before we give back the FD,
** and park this thread in the search thread pool, print the
** local output we have gathered.
*/
print_local_output(out,wt); /* this also frees out nodes */
out = NULL; /* for the next time around, if there is one */
ERROR:
DP(DLEVEL5,("Search done for %s\n",wt->path));
free(wt->path);
free(wt);

notrun();
pthread_mutex_lock(&search_q_lk);
if (search_pool_cnt > search_thr_limit) {
    pthread_mutex_unlock(&search_q_lk);
    DP(DLEVEL5,("Search thread exiting\n"));
    if (flags & FS_STATS) {
        pthread_mutex_lock(&stat_lk);
        st_destroy++;
        pthread_mutex_unlock(&stat_lk);
    }
    return(0);
}
else {
    search_pool_cnt++;
    while (!search_q)

```

```

        pthread_cond_wait(&search_q_cv,&search_q_lk);
        search_pool_cnt--;
        wt = search_q; /* we have work to do! */
        if (search_q->next)
            search_q = search_q->next;
        else
            search_q = NULL;
        pthread_mutex_unlock(&search_q_lk);
    }
}
/*NOTREACHED*/
}

/*
 * Continue line: Special case search with the -C flag set. If you are
 * searching files like Makefiles, some lines might have escape char's to
 * continue the line on the next line. So the target string can be found, but
 * no data is displayed. This function continues to print the escaped line
 * until there are no more "\" chars found.
 */
int
continue_line(char *rline, FILE *fin, out_t *out, work_t *wt,
              int *lc, long *bc)
{
    int len;
    int cnt = 0;
    char *line;
    char nline[128];

    if (!(flags & FC_LINE))
        return(0);

    line = rline;
AGAIN:
    len = strlen(line);
    if (line[len-1] == '\\') {
        if ((fgets(nline,127,fin)) == NULL) {
            return(cnt);
        }
        line = nline;
        len = strlen(line);
        if (line[len-1] == '\n')
            line[len-1] = '\0';
        *bc = *bc + len;
        *lc++;
        add_output_local(&out,wt,*lc,*bc,line);
        cnt++;
        goto AGAIN;
    }
    return(cnt);
}

/*
 * cascade: This thread is started by the main thread when directory names
 * are found on the work Q. The thread reads all the new file, and directory
 * names from the directory it was started when and adds the names to the
 * work Q. (it finds more work!)
 */
void *
```

```

cascade(void *arg) /* work_t *arg */
{
    char          fullpath[1025];
    int           restart_cnt = 10;
    DIR           *dp;

    char          dir_buf[sizeof(struct dirent) + PATH_MAX];
    struct dirent *dent = (struct dirent *)dir_buf;
    struct stat    sbuf;
    char          *fpath;
    work_t        *wt;
    int           fl = 0, dl = 0;
    int           pm_file_len = 0;

    pthread_yield_np(); /* try to give control back to main thread */
    wt = (work_t *)arg;

    while(1) {
        fl = 0;
        dl = 0;
        restart_cnt = 10;
        pm_file_len = 0;

        pthread_mutex_lock(&running_lk);
        running++;
        pthread_mutex_unlock(&running_lk);
        pthread_mutex_lock(&work_q_lk);
        tglimit--;
        pthread_mutex_unlock(&work_q_lk);

        if (!wt) {
            if (!(flags & FS_NOERROR))
                fprintf(stderr, "tgrep: Bad work node passed to cascade\n");
            goto DONE;
        }
        fpath = (char *)wt->path;
        if (!fpath) {
            if (!(flags & FS_NOERROR))
                fprintf(stderr, "tgrep: Bad path name passed to cascade\n");
            goto DONE;
        }
        DP(DLEVEL3, ("Cascading on %s\n", fpath));
        if ((dp = opendir(fpath)) == NULL) {
            if (!(flags & FS_NOERROR))
                fprintf(stderr, "tgrep: Can't open dir %s, %s. Ignored.\n",
                    fpath, strerror(errno));
            goto DONE;
        }
        while ((readdir_r(dp, dent)) != NULL) {
            restart_cnt = 10; /* only try to restart the interrupted 10 X */

            if (dent->d_name[0] == '.') {
                if (dent->d_name[1] == '.' && dent->d_name[2] == '\0')
                    continue;
                if (dent->d_name[1] == '\0')
                    continue;
            }

            fl = strlen(fpath);
            dl = strlen(dent->d_name);
        }
    }
}

```

```

    if ((fl + 1 + dl) > 1024) {
        fprintf(stderr, "tgrep: Path %s/%s is too long. "
            "MaxPath = 1024\n",
            fpath, dent->d_name);
        continue; /* try the next name in this directory */
    }
    strcpy(fullpath, fpath);
    strcat(fullpath, "/");
    strcat(fullpath, dent->d_name);

RESTART_STAT:
    if (stat(fullpath, &sbuf)) {
        if (errno == EINTR) {
            if (--restart_cnt)
                goto RESTART_STAT;
        }
        if (!(flags & FS_NOERROR))
            fprintf(stderr, "tgrep: Can't stat file/dir %s, %s. "
                "Ignored.\n",
                fullpath, strerror(errno));
        goto ERROR;
    }

    switch (sbuf.st_mode & S_IFMT) {
    case S_IFREG :
        if (flags & TG_FILEPAT) {
            if (pmatch(pm_file_pat, dent->d_name, &pm_file_len)) {
                DP(DLEVEL3, ("file pat match (cascade) %s\n",
                    dent->d_name));
                add_work(fullpath, FILET);
            }
        }
        else {
            add_work(fullpath, FILET);
            DP(DLEVEL3, ("cascade added file (MATCH) %s to Work Q\n",
                fullpath));
        }
        break;

    case S_IFDIR :
        DP(DLEVEL3, ("cascade added dir %s to Work Q\n", fullpath));
        add_work(fullpath, DIRT);
        break;
    }
}

ERROR:
    closedir(dp);

DONE:
    free(wt->path);
    free(wt);
    notrun();
    pthread_mutex_lock(&cascade_q_lk);
    if (cascade_pool_cnt > cascade_thr_limit) {
        pthread_mutex_unlock(&cascade_q_lk);
        DP(DLEVEL5, ("Cascade thread exiting\n"));
        if (flags & FS_STATS) {
            pthread_mutex_lock(&stat_lk);
            st_cascade_destroy++;
        }
    }

```

```

        pthread_mutex_unlock(&stat_lk);
    }
    return(0); /* pthread_exit */
}
else {
    DP(DLEVEL5,("Cascade thread waiting in pool\n"));
    cascade_pool_cnt++;
    while (!cascade_q)
        pthread_cond_wait(&cascade_q_cv,&cascade_q_lk);
    cascade_pool_cnt--;
    wt = cascade_q; /* we have work to do! */
    if (cascade_q->next)
        cascade_q = cascade_q->next;
    else
        cascade_q = NULL;
    pthread_mutex_unlock(&cascade_q_lk);
}
}
/*NOTREACHED*/
}

/*
 * Print Local Output: Called by the search thread after it is done searching
 * a single file. If any output was saved (matching lines), the lines are
 * displayed as a group on stdout.
 */
int
print_local_output(out_t *out, work_t *wt)
{
    out_t      *pp, *op;
    int         out_count = 0;
    int         printed = 0;

    pp = out;
    pthread_mutex_lock(&output_print_lk);
    if (pp && (flags & TG_PROGRESS)) {
        progress++;
        if (progress >= progress_offset) {
            progress = 0;
            fprintf(stderr, ".");
        }
    }
    while (pp) {
        out_count++;
        if (!(flags & FC_COUNT)) {
            if (flags & FL_NAMEONLY) { /* Print name ONLY ! */
                if (!printed) {
                    printed = 1;
                    printf("%s\n",wt->path);
                }
            }
            else { /* We are printing more then just the name */
                if (!(flags & FH_HOLDNAME))
                    printf("%s :",wt->path);
                if (flags & FB_BLOCK)
                    printf("%ld:",pp->byte_count/512+1);
                if (flags & FN_NUMBER)
                    printf("%d:",pp->line_count);
                printf("%s\n",pp->line);
            }
        }
    }
}

```

```

    }
    op = pp;
    pp = pp->next;
    /* free the nodes as we go down the list */
    free(op->line);
    free(op);
}

pthread_mutex_unlock(&output_print_lk);
pthread_mutex_lock(&global_count_lk);
global_count += out_count;
pthread_mutex_unlock(&global_count_lk);
return(0);
}

/*
 * add output local: is called by a search thread as it finds matching lines.
 * the matching line, its byte offset, line count, etc. are stored until the
 * search thread is done searching the file, then the lines are printed as
 * a group. This way the lines from more than a single file are not mixed
 * together.
 */

int
add_output_local(out_t **out, work_t *wt,int lc, long bc, char *line)
{
    out_t      *ot,*oo, *op;

    if (( ot = (out_t *)malloc(sizeof(out_t))) == NULL)
        goto ERROR;
    if (( ot->line = (char *)malloc(strlen(line)+1)) == NULL)
        goto ERROR;

    strcpy(ot->line,line);
    ot->line_count = lc;
    ot->byte_count = bc;

    if (!*out) {
        *out = ot;
        ot->next = NULL;
        return(0);
    }
    /* append to the END of the list; keep things sorted! */
    op = oo = *out;
    while(oo) {
        op = oo;
        oo = oo->next;
    }
    op->next = ot;
    ot->next = NULL;
    return(0);

ERROR:
    if (!(flags & FS_NOERROR))
        fprintf(stderr,"tgrep: Output lost. No space. "
            "[%s: line %d byte %d match : %s\n",
            wt->path,lc,bc,line);
    return(1);
}

```

```

/*
 * print stats: If the -S flag is set, after ALL files have been searched,
 * main thread calls this function to print the stats it keeps on how the
 * search went.
 */

void
prnt_stats(void)
{
    float a,b,c;
    float t = 0.0;
    time_t st_end = 0;
    char    tl[80];

    st_end = time(NULL); /* stop the clock */
    printf("\n----- Tgrep Stats. ----- \n");
    printf("Number of directories searched:      %d\n",st_dir_search);
    printf("Number of files searched:                %d\n",st_file_search);
    c = (float)(st_dir_search + st_file_search) / (float)(st_end - st_start);
    printf("Dir/files per second:                      %3.2f\n",c);
    printf("Number of lines searched:                  %d\n",st_line_search);
    printf("Number of matching lines to target:        %d\n",global_count);

    printf("Number of cascade threads created:          %d\n",st_cascade);
    printf("Number of cascade threads from pool:        %d\n",st_cascade_pool);
    a = st_cascade_pool; b = st_dir_search;
    printf("Cascade thread pool hit rate:              %3.2f%%\n",((a/b)*100));
    printf("Cascade pool overall size:                  %d\n",cascade_pool_cnt);
    printf("Number of search threads created:           %d\n",st_search);
    printf("Number of search threads from pool:         %d\n",st_pool);
    a = st_pool; b = st_file_search;
    printf("Search thread pool hit rate:                %3.2f%%\n",((a/b)*100));
    printf("Search pool overall size:                   %d\n",search_pool_cnt);
    printf("Search pool size limit:                     %d\n",search_thr_limit);
    printf("Number of search threads destroyed:         %d\n",st_destroy);

    printf("Max # of threads running concurrently:      %d\n",st_maxrun);
    printf("Total run time, in seconds.                  %d\n",
        (st_end - st_start));

    /* Why did we wait ? */
    a = st_workfds; b = st_dir_search+st_file_search;
    c = (a/b)*100; t += c;
    printf("Work stopped due to no FD's:  (%.3d)      %d Times, %3.2f%%\n",
        search_thr_limit,st_workfds,c);
    a = st_worknull; b = st_dir_search+st_file_search;
    c = (a/b)*100; t += c;
    printf("Work stopped due to no work on Q:          %d Times, %3.2f%%\n",
        st_worknull,c);
    if (tglimit == UNLIMITED)
        strcpy(tl,"Unlimited");
    else
        sprintf(tl,"  %.3d  ",tglimit);
    a = st_worklimit; b = st_dir_search+st_file_search;
    c = (a/b)*100; t += c;
    printf("Work stopped due to TGLIMIT:  (%.9s) %d Times, %3.2f%%\n",
        tl,st_worklimit,c);
    printf("Work continued to be handed out:          %3.2f%%\n",100.00-t);
    printf("----- \n");
}

```

```

/*
 * not running: A glue function to track if any search threads or cascade
 * threads are running. When the count is zero, and the work Q is NULL,
 * we can safely say, WE ARE DONE.
 */
void
notrun (void)
{
    pthread_mutex_lock(&work_q_lk);
    work_cnt--;
    tglimit++;
    current_open_files++;
    pthread_mutex_lock(&running_lk);
    if (flags & FS_STATS) {
        pthread_mutex_lock(&stat_lk);
        if (running > st_maxrun) {
            st_maxrun = running;
            DP(DLEVEL6, ("Max Running has increased to %d\n", st_maxrun));
        }
        pthread_mutex_unlock(&stat_lk);
    }
    running--;
    if (work_cnt == 0 && running == 0) {
        all_done = 1;
        DP(DLEVEL6, ("Setting ALL_DONE flag to TRUE.\n"));
    }
    pthread_mutex_unlock(&running_lk);
    pthread_cond_signal(&work_q_cv);
    pthread_mutex_unlock(&work_q_lk);
}

/*
 * uncase: A glue function. If the -i (case insensitive) flag is set, the
 * target string and the read in line is converted to lower case before
 * comparing them.
 */
void
uncase(char *s)
{
    char      *p;

    for (p = s; *p != NULL; p++)
        *p = (char)tolower(*p);
}

/*
 * usage: Have to have one of these.
 */

void
usage(void)
{
    fprintf(stderr, "usage: tgrep <options> pattern <{file,dir}>...\n");
    fprintf(stderr, "\n");
    fprintf(stderr, "Where:\n");
#ifdef DEBUG
    fprintf(stderr, "Debug      -d = debug level -d <levels> (-d0 for usage)\n");
    fprintf(stderr, "Debug      -f = block fd's from use (-f #)\n");
#endif
    fprintf(stderr, "          -b = show block count (512 byte block)\n");
}

```



```

fprintf(stderr,"    @(n)c - match <c> at column <n>\n");
fprintf(stderr,"    | - match either pattern\n");
fprintf(stderr,"    \\ - escape any special characters\n");
fprintf(stderr,"    \\c - escape any special characters\n");
fprintf(stderr,"    \\o - turn on any special characters\n");
fprintf(stderr,"\\n");
fprintf(stderr,"To match two different patterns in the same command\n");
fprintf(stderr,"Use the or function. \\n"
        "ie: tgrep -e \"(pat1)|(pat2)\" file\n"
        "This will match any line with \"pat1\" or \"pat2\" in it.\n");
fprintf(stderr,"You can also use up to %d -e expressions\n",MAXREGEXP);
fprintf(stderr,"RegExp Pattern matching brought to you by Marc Staveley\n");
exit(0);
}

/*
 * debug usage: If compiled with -DDEBUG, turn it on, and tell the world
 * how to get tgrep to print debug info on different threads.
 */

#ifdef DEBUG
void
debug_usage(void)
{
    int i = 0;

    fprintf(stderr,"DEBUG usage and levels:\n");
    fprintf(stderr,"-----\n");
    fprintf(stderr,"Level                code\n");
    fprintf(stderr,"-----\n");
    fprintf(stderr,"0                    This message.\n");
    for (i=0; i<9; i++) {
        fprintf(stderr,"%d                    %s\n",i+1,debug_set[i].name);
    }
    fprintf(stderr,"-----\n");
    fprintf(stderr,"You can or the levels together like -d134 for levels\n");
    fprintf(stderr,"1 and 3 and 4.\n");
    fprintf(stderr,"\\n");
    exit(0);
}
#endif

/* Pthreads NP functions */

#ifdef __sun
void
pthread_setconcurrency_np(int con)
{
    thr_setconcurrency(con);
}

int
pthread_getconcurrency_np(void)
{
    return(thr_getconcurrency());
}

void
pthread_yield_np(void)
{

```

```

/*      In Solaris 2.4, these functions always return - 1 and set errno to ENOSYS */
    if (sched_yield()) /* call UI interface if we are older than 2.5 */
        thr_yield();
}

#else
void
pthread_setconcurrency_np(int con)
{
    return;
}

int
pthread_getconcurrency_np(void)
{
    return(0);
}

void
pthread_yield_np(void)
{
    return;
}
#endif

```


Solaris Threads Example: barrier.c

The `barrier.c` program demonstrates an implementation of a barrier for Solaris threads. (See “Parallelizing a Loop on a Shared-Memory Parallel Computer” on page 240 for a definition of barriers.)

CODE EXAMPLE B-1 Solaris Threads Example: barrier.c

```
#define _REENTRANT

/* Include Files */

#include <thread.h>
#include <errno.h>

/* Constants & Macros */

/* Data Declarations */

typedef struct {
    int    maxcnt;    /* maximum number of runners */
    struct _sb {
        cond_t wait_cv; /* cv for waiters at barrier */
        mutex_t wait_lk; /* mutex for waiters at barrier */
        int    runners; /* number of running threads */
    } sb[2];
    struct _sb *sbp; /* current sub-barrier */
} barrier_t;

/*
 * barrier_init - initialize a barrier variable.
 */

int
barrier_init( barrier_t *bp, int count, int type, void *arg ) {
    int n;
```

```

    int i;

    if (count < 1)
        return(EINVAL);

    bp->maxcnt = count;
    bp->sbp = &bp->sb[0];

    for (i = 0; i < 2; ++i) {
#ifdef __cplusplus
        struct barrier_t::_sb *sbp = &( bp->sb[i] );
#else
        struct _sb *sbp = &( bp->sb[i] );
#endif
        sbp->runners = count;

        if (n = mutex_init(&sbp->wait_lk, type, arg))
            return(n);

        if (n = cond_init(&sbp->wait_cv, type, arg))
            return(n);
    }
    return(0);
}

/*
 * barrier_wait - wait at a barrier for everyone to arrive.
 */

int
barrier_wait(register barrier_t *bp) {
#ifdef __cplusplus
    register struct barrier_t::_sb *sbp = bp->sbp;
#else
    register struct _sb *sbp = bp->sbp;
#endif
    mutex_lock(&sbp->wait_lk);

    if (sbp->runners == 1) { /* last thread to reach barrier */
        if (bp->maxcnt != 1) {
            /* reset runner count and switch sub-barriers */
            sbp->runners = bp->maxcnt;
            bp->sbp = (bp->sbp == &bp->sb[0])
                ? &bp->sb[1] : &bp->sb[0];

            /* wake up the waiters */
            cond_broadcast(&sbp->wait_cv);
        }
    } else {
        sbp->runners--; /* one less runner */

        while (sbp->runners != bp->maxcnt)
            cond_wait( &sbp->wait_cv, &sbp->wait_lk);
    }

    mutex_unlock(&sbp->wait_lk);

    return(0);
}

```

```

/*
 * barrier_destroy - destroy a barrier variable.
 */

int
barrier_destroy(barrier_t *bp) {
    int    n;
    int    i;

    for (i=0; i < 2; ++ i) {
        if (n = cond_destroy(&bp->sb[i].wait_cv))
            return( n );

        if (n = mutex_destroy( &bp->sb[i].wait_lk))
            return(n);
    }

    return(0);
}

#define NTHR    4
#define NCOMPUTATION 2
#define NITER   1000
#define NSQRT   1000

void *
compute(barrier_t *ba )
{
    int count = NCOMPUTATION;

    while (count--) {
        barrier_wait( ba );
        /* do parallel computation */
    }
}

main( int argc, char *argv[] ) {
    int    i;
    int    niter;
    int    nthr;
    barrier_t    ba;
    double    et;
    thread_t    *tid;

    switch ( argc ) {
        default:
            case 3 :    niter    = atoi( argv[1] );
                       nthr     = atoi( argv[2] );
                       break;

            case 2 :    niter    = atoi( argv[1] );
                       nthr     = NTHR;
                       break;

            case 1 :    niter    = NITER;
                       nthr     = NTHR;
                       break;
    }
}

```

```

    }

    barrier_init( &ba, nthr + 1, USYNC_THREAD, NULL );
    tid = (thread_t *) calloc(nthr, sizeof(thread_t));

    for (i = 0; i < nthr; ++i) {
        int n;

        if (n = thr_create(NULL, 0,
            (void (*)( void *)) compute,
            &ba, NULL, &tid[i])) {
            errno = n;
            perror("thr_create");
            exit(1);
        }
    }

    for (i = 0; i < NCOMPUTATION; i++) {
        barrier_wait(&ba );
        /* do parallel algorithm */
    }

    for (i = 0; i < nthr; i++) {
        thr_join(tid[i], NULL, NULL);
    }
}

```

MT Safety Levels: Library Interfaces

Table C-1 lists the safety levels for interfaces from Section 3 of the `man(3)` Library Routines (see “MT Interface Safety Levels” on page 170 for explanations of the safety categories).

TABLE C-1 MT Safety Levels of Library Routines

Library Routine	Safety Level
<code>a64l(3C)</code>	MT-Safe
<code>abort(3C)</code>	Safe
<code>abs(3C)</code>	MT-Safe
<code>accept(3N)</code>	Safe
<code>acos(3M)</code>	MT-Safe
<code>acosh(3M)</code>	MT-Safe
<code>addch(3X)</code>	Unsafe
<code>addchnstr(3X)</code>	Unsafe
<code>addchstr(3X)</code>	Unsafe
<code>addnstr(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>addnwstr(3X)</code>	Unsafe
<code>addsev(3C)</code>	MT-Safe
<code>addseverity(3C)</code>	Safe
<code>addstr(3X)</code>	Unsafe
<code>addwch(3X)</code>	Unsafe
<code>addwchnstr(3X)</code>	Unsafe
<code>addwchstr(3X)</code>	Unsafe
<code>addwstr(3X)</code>	Unsafe
<code>adjcurspos(3X)</code>	Unsafe
<code>advance(3G)</code>	MT-Safe
<code>aiocancel(3)</code>	Unsafe
<code>aioread(3)</code>	Unsafe
<code>aiowait(3)</code>	Unsafe
<code>aiowrite(3)</code>	Unsafe
<code>aio_cancel(3R)</code>	MT-Safe
<code>aio_error(3R)</code>	Async-Signal-Safe
<code>aio_fsync(3R)</code>	MT-Safe
<code>aio_read(3R)</code>	MT-Safe
<code>aio_return(3R)</code>	Async-Signal-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>aio_suspend(3R)</code>	Async-Signal-Safe
<code>aio_write(3R)</code>	MT-Safe
<code>alloca(3C)</code>	Safe
<code>arc(3)</code>	Safe
<code>ascftime(3C)</code>	MT-Safe
<code>asctime(3C)</code>	Unsafe, use <code>asctime_r()</code>
<code>asin(3M)</code>	MT-Safe
<code>asinh(3M)</code>	MT-Safe
<code>assert(3C)</code>	Safe
<code>atan(3M)</code>	MT-Safe
<code>atan2(3M)</code>	MT-Safe
<code>atanh(3M)</code>	MT-Safe
<code>atexit(3C)</code>	Safe
<code>atof(3C)</code>	MT-Safe
<code>atoi(3C)</code>	MT-Safe
<code>atol(3C)</code>	MT-Safe
<code>atoll(3C)</code>	MT-Safe
<code>attroff(3X)</code>	Unsafe
<code>attron(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>attrset(3X)</code>	Unsafe
<code>authdes_create(3N)</code>	Unsafe
<code>authdes_getucred(3N)</code>	MT-Safe
<code>authdes_seccreate(3N)</code>	MT-Safe
<code>authkerb_getucred(3N)</code>	Unsafe
<code>authkerb_seccreate(3N)</code>	Unsafe
<code>authnone_create(3N)</code>	MT-Safe
<code>authsys_create(3N)</code>	MT-Safe
<code>authsys_create_default(3N)</code>	MT-Safe
<code>authunix_create(3N)</code>	Unsafe
<code>authunix_create_default(3N)</code>	Unsafe
<code>auth_destroy(3N)</code>	MT-Safe
<code>au_close(3)</code>	Safe
<code>au_open(3)</code>	Safe
<code>au_user_mask(3)</code>	MT-Safe
<code>au_write(3)</code>	Safe
<code>basename(3G)</code>	MT-Safe
<code>baudrate(3X)</code>	Unsafe
<code>beep(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
bessel(3M)	MT-Safe
bgets(3G)	MT-Safe
bind(3N)	Safe
bindtextdomain(3I)	Safe with exceptions
bkgd(3X)	Unsafe
bkgdset(3X)	Unsafe
border(3X)	Unsafe
bottom_panel(3X)	Unsafe
box(3)	Safe
box(3X)	Unsafe
bsearch(3C)	Safe
bufsplit(3G)	MT-Safe
byteorder(3N)	Safe
calloc(3C)	Safe
calloc(3X)	Safe
callrpc(3N)	Unsafe
cancellation(3T)	MT-Safe
can_change_color(3X)	Unsafe
catclose(3C)	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>catgets(3C)</code>	MT-Safe
<code>catopen(3C)</code>	MT-Safe
<code>cbc_crypt(3)</code>	MT-Safe
<code>cbreak(3X)</code>	Unsafe
<code>cbrt(3M)</code>	MT-Safe
<code>ceil(3M)</code>	MT-Safe
<code>cfgetispeed(3)</code>	MT-Safe, Async-Signal-Safe
<code>cfgetospeed(3)</code>	MT-Safe, Async-Signal-Safe
<code>cfree(3X)</code>	Safe
<code>cfsetispeed(3)</code>	MT-Safe, Async-Signal-Safe
<code>cfsetospeed(3)</code>	MT-Safe, Async-Signal-Safe
<code>cftime(3C)</code>	MT-Safe
<code>circle(3)</code>	Safe
<code>clear(3X)</code>	Unsafe
<code>clearerr(3S)</code>	MT-Safe
<code>clearok(3X)</code>	Unsafe
<code>clntraw_create(3N)</code>	Unsafe
<code>clnttcp_create(3N)</code>	Unsafe
<code>clntudp_bufcreate(3N)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>clntudp_create(3N)</code>	Unsafe
<code>clnt_broadcast(3N)</code>	Unsafe
<code>clnt_call(3N)</code>	MT-Safe
<code>clnt_control(3N)</code>	MT-Safe
<code>clnt_create(3N)</code>	MT-Safe
<code>clnt_create_timed(3N)</code>	MT-Safe
<code>clnt_create_vers(3N)</code>	MT-Safe
<code>clnt_destroy(3N)</code>	MT-Safe
<code>clnt_dg_create(3N)</code>	MT-Safe
<code>clnt_freeres(3N)</code>	MT-Safe
<code>clnt_geterr(3N)</code>	MT-Safe
<code>clnt_pcreateerror(3N)</code>	MT-Safe
<code>clnt_perrno(3N)</code>	MT-Safe
<code>clnt_perror(3N)</code>	MT-Safe
<code>clnt_raw_create(3N)</code>	MT-Safe
<code>clnt_spcreateerror(3N)</code>	MT-Safe
<code>clnt_sperrno(3N)</code>	MT-Safe
<code>clnt_sperror(3N)</code>	MT-Safe
<code>clnt_tli_create(3N)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>clnt_tp_create(3N)</code>	MT-Safe
<code>clnt_tp_create_timed(3N)</code>	MT-Safe
<code>clnt_vc_create(3N)</code>	MT-Safe
<code>clock(3C)</code>	MT-Safe
<code>clock_gettime(3R)</code>	Async-Signal-Safe
<code>closedir(3C)</code>	Safe
<code>closelog(3)</code>	Safe
<code>closepl(3)</code>	Safe
<code>closevt(3)</code>	Safe
<code>clrtobot(3X)</code>	Unsafe
<code>clrtoeol(3X)</code>	Unsafe
<code>color_content(3X)</code>	Unsafe
<code>compile(3G)</code>	MT-Safe
<code>condition(3T)</code>	MT-Safe
<code>cond_broadcast(3T)</code>	MT-Safe
<code>cond_destroy(3T)</code>	MT-Safe
<code>cond_init(3T)</code>	MT-Safe
<code>cond_signal(3T)</code>	MT-Safe
<code>cond_timedwait(3T)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>cond_wait(3T)</code>	MT-Safe
<code>confstr(3C)</code>	MT-Safe
<code>connect(3N)</code>	Safe
<code>cont(3)</code>	Safe
<code>conv(3C)</code>	MT-Safe with exceptions
<code>copylist(3G)</code>	MT-Safe
<code>copysign(3M)</code>	MT-Safe
<code>copywin(3X)</code>	Unsafe
<code>cos(3M)</code>	MT-Safe
<code>cosh(3M)</code>	MT-Safe
<code>crypt(3C)</code>	Safe
<code>crypt(3X)</code>	Unsafe
<code>cset(3I)</code>	MT-Safe with exceptions
<code>csetcol(3I)</code>	MT-Safe with exceptions
<code>csetlen(3I)</code>	MT-Safe with exceptions
<code>csetno(3I)</code>	MT-Safe with exceptions
<code>ctermid(3S)</code>	Unsafe, use <code>ctermid_r()</code>
<code>ctime(3C)</code>	Unsafe, use <code>ctime_r()</code>
<code>ctype(3C)</code>	MT-Safe with exceptions

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>current_field(3X)</code>	Unsafe
<code>current_item(3X)</code>	Unsafe
<code>curses(3X)</code>	Unsafe
<code>curs_addch(3X)</code>	Unsafe
<code>curs_addchstr(3X)</code>	Unsafe
<code>curs_addstr(3X)</code>	Unsafe
<code>curs_addwch(3X)</code>	Unsafe
<code>curs_addwchstr(3X)</code>	Unsafe
<code>curs_addwstr(3X)</code>	Unsafe
<code>curs_alecompat(3X)</code>	Unsafe
<code>curs_attr(3X)</code>	Unsafe
<code>curs_beep(3X)</code>	Unsafe
<code>curs_bkgd(3X)</code>	Unsafe
<code>curs_border(3X)</code>	Unsafe
<code>curs_clear(3X)</code>	Unsafe
<code>curs_color(3X)</code>	Unsafe
<code>curs_delch(3X)</code>	Unsafe
<code>curs_deleteln(3X)</code>	Unsafe
<code>curs_getch(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>curs_getstr(3X)</code>	Unsafe
<code>curs_getwch(3X)</code>	Unsafe
<code>curs_getwstr(3X)</code>	Unsafe
<code>curs_getyx(3X)</code>	Unsafe
<code>curs_inch(3X)</code>	Unsafe
<code>curs_inchstr(3X)</code>	Unsafe
<code>curs_initscr(3X)</code>	Unsafe
<code>curs_inopts(3X)</code>	Unsafe
<code>curs_insch(3X)</code>	Unsafe
<code>curs_insstr(3X)</code>	Unsafe
<code>curs_instr(3X)</code>	Unsafe
<code>curs_inswch(3X)</code>	Unsafe
<code>curs_inswstr(3X)</code>	Unsafe
<code>curs_inwch(3X)</code>	Unsafe
<code>curs_inwchstr(3X)</code>	Unsafe
<code>curs_inwstr(3X)</code>	Unsafe
<code>curs_kernel(3X)</code>	Unsafe
<code>curs_move(3X)</code>	Unsafe
<code>curs_outopts(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>curs_overlay(3X)</code>	Unsafe
<code>curs_pad(3X)</code>	Unsafe
<code>curs_printw(3X)</code>	Unsafe
<code>curs_refresh(3X)</code>	Unsafe
<code>curs_scanw(3X)</code>	Unsafe
<code>curs_scroll(3X)</code>	Unsafe
<code>curs_scr_dump(3X)</code>	Unsafe
<code>curs_set(3X)</code>	Unsafe
<code>curs_slk(3X)</code>	Unsafe
<code>curs_termattrs(3X)</code>	Unsafe
<code>curs_termcap(3X)</code>	Unsafe
<code>curs_terminfo(3X)</code>	Unsafe
<code>curs_touch(3X)</code>	Unsafe
<code>curs_util(3X)</code>	Unsafe
<code>curs_window(3X)</code>	Unsafe
<code>cuserid(3S)</code>	MT-Safe
<code>data_ahead(3X)</code>	Unsafe
<code>data_behind(3X)</code>	Unsafe
<code>dbm_clearerr(3)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>dbm_close(3)</code>	Unsafe
<code>dbm_delete(3)</code>	Unsafe
<code>dbm_error(3)</code>	Unsafe
<code>dbm_fetch(3)</code>	Unsafe
<code>dbm_firstkey(3)</code>	Unsafe
<code>dbm_nextkey(3)</code>	Unsafe
<code>dbm_open(3)</code>	Unsafe
<code>dbm_store(3)</code>	Unsafe
<code>db_add_entry(3N)</code>	Unsafe
<code>db_checkpoint(3N)</code>	Unsafe
<code>db_create_table(3N)</code>	Unsafe
<code>db_destroy_table(3N)</code>	Unsafe
<code>db_first_entry(3N)</code>	Unsafe
<code>db_free_result(3N)</code>	Unsafe
<code>db_initialize(3N)</code>	Unsafe
<code>db_list_entries(3N)</code>	Unsafe
<code>db_next_entry(3N)</code>	Unsafe
<code>db_remove_entry(3N)</code>	Unsafe
<code>db_reset_next_entry(3N)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>db_standby(3N)</code>	Unsafe
<code>db_table_exists(3N)</code>	Unsafe
<code>db_unload_table(3N)</code>	Unsafe
<code>dcgettext(3I)</code>	Safe with exceptions
<code>decimal_to_double(3)</code>	MT-Safe
<code>decimal_to_extended(3)</code>	MT-Safe
<code>decimal_to_floating(3)</code>	MT-Safe
<code>decimal_to_quadruple(3)</code>	MT-Safe
<code>decimal_to_single(3)</code>	MT-Safe
<code>def_prog_mode(3X)</code>	Unsafe
<code>def_shell_mode(3X)</code>	Unsafe
<code>delay_output(3X)</code>	Unsafe
<code>delch(3X)</code>	Unsafe
<code>deleteln(3X)</code>	Unsafe
<code>delscreen(3X)</code>	Unsafe
<code>delwin(3X)</code>	Unsafe
<code>del_curterm(3X)</code>	Unsafe
<code>del_panel(3X)</code>	Unsafe
<code>derwin(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>des_crypt(3)</code>	MT-Safe
<code>DES_FAILED(3)</code>	MT-Safe
<code>des_failed(3)</code>	MT-Safe
<code>des_setparity(3)</code>	MT-Safe
<code>dgettext(3I)</code>	Safe with exceptions
<code>dial(3N)</code>	Unsafe
<code>difftime(3C)</code>	MT-Safe
<code>dirname(3G)</code>	MT-Safe
<code>div(3C)</code>	MT-Safe
<code>dladdr(3X)</code>	MT-Safe
<code>dlclose(3X)</code>	MT-Safe
<code>dlderror(3X)</code>	MT-Safe
<code>dlopen(3X)</code>	MT-Safe
<code>dlsym(3X)</code>	MT-Safe
<code>dn_comp(3N)</code>	Unsafe
<code>dn_expand(3N)</code>	Unsafe
<code>doconfig(3N)</code>	Unsafe
<code>double_to_decimal(3)</code>	MT-Safe
<code>doupdate(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>drand48(3C)</code>	Safe
<code>dup2(3C)</code>	Unsafe, Async-Signal-Safe
<code>dupwin(3X)</code>	Unsafe
<code>dup_field(3X)</code>	Unsafe
<code>dynamic_field_info(3X)</code>	Unsafe
<code>ecb_crypt(3)</code>	MT-Safe
<code>echo(3X)</code>	Unsafe
<code>echochar(3X)</code>	Unsafe
<code>echowchar(3X)</code>	Unsafe
<code>econvert(3)</code>	MT-Safe
<code>ecvt(3)</code>	MT-Safe
<code>ecvt(3C)</code>	Unsafe
<code>el(32_fsize.3E)</code>	Unsafe
<code>el(32_getehdr.3E)</code>	Unsafe
<code>el(32_getshdr.3E)</code>	Unsafe
<code>el(32_newehdr.3E)</code>	Unsafe
<code>el(32_newphdr.3E)</code>	Unsafe
<code>el(32_xlatetof.3E)</code>	Unsafe
<code>el(32_xlatetom.3E)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>elf(3E)</code>	Unsafe
<code>elf_begin(3E)</code>	Unsafe
<code>elf_cntl(3E)</code>	Unsafe
<code>elf_end(3E)</code>	Unsafe
<code>elf_errmsg(3E)</code>	Unsafe
<code>elf_errno(3E)</code>	Unsafe
<code>elf_fill(3E)</code>	Unsafe
<code>elf_flagdata(3E)</code>	Unsafe
<code>elf_flagehdr(3E)</code>	Unsafe
<code>elf_flagelf(3E)</code>	Unsafe
<code>elf_flagphdr(3E)</code>	Unsafe
<code>elf_flagscn(3E)</code>	Unsafe
<code>elf_flagshdr(3E)</code>	Unsafe
<code>elf_getarhdr(3E)</code>	Unsafe
<code>elf_getarsym(3E)</code>	Unsafe
<code>elf_getbase(3E)</code>	Unsafe
<code>elf_getdata(3E)</code>	Unsafe
<code>elf_getident(3E)</code>	Unsafe
<code>elf_getscn(3E)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>elf_hash(3E)</code>	Unsafe
<code>elf_kind(3E)</code>	Unsafe
<code>elf_memory(3E)</code>	Unsafe
<code>elf_ndxscn(3E)</code>	Unsafe
<code>elf_newdata(3E)</code>	Unsafe
<code>elf_newscn(3E)</code>	Unsafe
<code>elf_next(3E)</code>	Unsafe
<code>elf_nextscn(3E)</code>	Unsafe
<code>elf_rand(3E)</code>	Unsafe
<code>elf_rawdata(3E)</code>	Unsafe
<code>elf_rawfile(3E)</code>	Unsafe
<code>elf_strptr(3E)</code>	Unsafe
<code>elf_update(3E)</code>	Unsafe
<code>elf_version(3E)</code>	Unsafe
<code>encrypt(3C)</code>	Safe
<code>endac(3)</code>	Safe
<code>endauclass(3)</code>	MT-Safe
<code>endauevent(3)</code>	MT-Safe
<code>endauuser(3)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
endnetconfig(3N)	MT-Safe
endnetpath(3N)	MT-Safe
endutent(3C)	Unsafe
endutxent(3C)	Unsafe
endwin(3X)	Unsafe
erand48(3C)	Safe
erase(3)	Safe
erase(3X)	Unsafe
erasechar(3X)	Unsafe
erf(3M)	MT-Safe
erfc(3M)	MT-Safe
errno(3C)	MT-Safe
ethers(3N)	MT-Safe
ether_aton(3N)	MT-Safe
ether_hostton(3N)	MT-Safe
ether_line(3N)	MT-Safe
ether_ntoa(3N)	MT-Safe
ether_ntohost(3N)	MT-Safe
euccol(3I)	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
euclen(3I)	Safe
eucscol(3I)	Safe
exit(3C)	Safe
exp(3M)	MT-Safe
expm1(3M)	MT-Safe
extended_to_decimal(3)	MT-Safe
fabs(3M)	MT-Safe
fattach(3C)	MT-Safe
fclose(3S)	MT-Safe
fconvert(3)	MT-Safe
fcvt(3)	MT-Safe
fcvt(3C)	Unsafe
fdatasync(3R)	Async-Signal-Safe
fdetach(3C)	Unsafe
fdopen(3S)	MT-Safe
feof(3S)	MT-Safe
ferror(3S)	MT-Safe
fflush(3S)	MT-Safe
ffs(3C)	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fgetc(3S)</code>	MT-Safe
<code>fgetgrent(3C)</code>	Unsafe, use <code>fgetgrent_r()</code>
<code>fgetpos(3C)</code>	MT-Safe
<code>fgetpwent(3C)</code>	Unsafe, use <code>fgetpwent_r()</code>
<code>fgets(3S)</code>	MT-Safe
<code>fgetspent(3C)</code>	Unsafe, use <code>fgetspent_r()</code>
<code>fgetwc(3I)</code>	MT-Safe
<code>fgetws(3I)</code>	MT-Safe
<code>field_arg(3X)</code>	Unsafe
<code>field_back(3X)</code>	Unsafe
<code>field_buffer(3X)</code>	Unsafe
<code>field_count(3X)</code>	Unsafe
<code>field_fore(3X)</code>	Unsafe
<code>field_index(3X)</code>	Unsafe
<code>field_info(3X)</code>	Unsafe
<code>field_init(3X)</code>	Unsafe
<code>field_just(3X)</code>	Unsafe
<code>field_opts(3X)</code>	Unsafe
<code>field_opts_off(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>field_opts_on(3X)</code>	Unsafe
<code>field_pad(3X)</code>	Unsafe
<code>field_status(3X)</code>	Unsafe
<code>field_term(3X)</code>	Unsafe
<code>field_type(3X)</code>	Unsafe
<code>field_userptr(3X)</code>	Unsafe
<code>fileno(3S)</code>	MT-Safe
<code>file_to_decimal(3)</code>	MT-Safe
<code>filter(3X)</code>	Unsafe
<code>finite(3C)</code>	MT-Safe
<code>flash(3X)</code>	Unsafe
<code>floating_to_decimal(3)</code>	MT-Safe
<code>flockfile(3S)</code>	MT-Safe
<code>floor(3M)</code>	MT-Safe
<code>flushinp(3X)</code>	Unsafe
<code>fmod(3M)</code>	MT-Safe
<code>fmtmsg(3C)</code>	Safe
<code>fnmatch(3C)</code>	MT-Safe
<code>fn_attribute_add(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_attribute_assign(3N)</code>	Safe
<code>fn_attribute_copy(3N)</code>	Safe
<code>fn_attribute_create(3N)</code>	Safe
<code>fn_attribute_destroy(3N)</code>	Safe
<code>fn_attribute_first(3N)</code>	Safe
<code>fn_attribute_identifier(3N)</code>	Safe
<code>fn_attribute_next(3N)</code>	Safe
<code>fn_attribute_remove(3N)</code>	Safe
<code>fn_attribute_syntax(3N)</code>	Safe
<code>FN_attribute_t(3N)</code>	Safe
<code>fn_attribute_valuecount(3N)</code>	Safe
<code>fn_attrmodlist_add(3N)</code>	Safe
<code>fn_attrmodlist_assign(3N)</code>	Safe
<code>fn_attrmodlist_copy(3N)</code>	Safe
<code>fn_attrmodlist_count(3N)</code>	Safe
<code>fn_attrmodlist_create(3N)</code>	Safe
<code>fn_attrmodlist_destroy(3N)</code>	Safe
<code>fn_attrmodlist_first(3N)</code>	Safe
<code>fn_attrmodlist_next(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
FN_attrmodlist_t(3N)	Safe
fn_attrset_add(3N)	Safe
fn_attrset_assign(3N)	Safe
fn_attrset_copy(3N)	Safe
fn_attrset_count(3N)	Safe
fn_attrset_create(3N)	Safe
fn_attrset_destroy(3N)	Safe
fn_attrset_first(3N)	Safe
fn_attrset_get(3N)	Safe
fn_attrset_next(3N)	Safe
fn_attrset_remove(3N)	Safe
FN_attrset_t(3N)	Safe
fn_attr_get(3N)	Safe
fn_attr_get_ids(3N)	Safe
fn_attr_get_values(3N)	Safe
fn_attr_modify(3N)	Safe
fn_attr_multi_get(3N)	Safe
fn_attr_multi_modify(3N)	Safe
fn_bindinglist_destroy(3N)	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_bindinglist_next(3N)</code>	Safe
<code>FN_bindinglist_t(3N)</code>	Safe
<code>fn_composite_name_append_comp(3N)</code>	Safe
<code>fn_composite_name_append_name(3N)</code>	Safe
<code>fn_composite_name_assign(3N)</code>	Safe
<code>fn_composite_name_copy(3N)</code>	Safe
<code>fn_composite_name_count(3N)</code>	Safe
<code>fn_composite_name_create(3N)</code>	Safe
<code>fn_composite_name_delete_comp(3N)</code>	Safe
<code>fn_composite_name_destroy(3N)</code>	Safe
<code>fn_composite_name_first(3N)</code>	Safe
<code>fn_composite_name_from_string(3N)</code>	Safe
<code>fn_composite_name_insert_comp(3N)</code>	Safe
<code>fn_composite_name_insert_name(3N)</code>	Safe
<code>fn_composite_name_is_empty(3N)</code>	Safe
<code>fn_composite_name_is_equal(3N)</code>	Safe
<code>fn_composite_name_is_prefix(3N)</code>	Safe
<code>fn_composite_name_is_suffix(3N)</code>	Safe
<code>fn_composite_name_last(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_composite_name_next(3N)</code>	Safe
<code>fn_composite_name_prefix(3N)</code>	Safe
<code>fn_composite_name_prepend_comp(3N)</code>	Safe
<code>fn_composite_name_prepend_name(3N)</code>	Safe
<code>fn_composite_name_prev(3N)</code>	Safe
<code>fn_composite_name_suffix(3N)</code>	Safe
<code>FN_composite_name_t(3N)</code>	Safe
<code>fn_compound_name_append_comp(3N)</code>	Safe
<code>fn_compound_name_assign(3N)</code>	Safe
<code>fn_compound_name_copy(3N)</code>	Safe
<code>fn_compound_name_count(3N)</code>	Safe
<code>fn_compound_name_delete_all(3N)</code>	Safe
<code>fn_compound_name_delete_comp(3N)</code>	Safe
<code>fn_compound_name_destroy(3N)</code>	Safe
<code>fn_compound_name_first(3N)</code>	Safe
<code>fn_compound_name_from_syntax_attrs</code>	Safe
<code>fn_compound_name_get_syntax_attrs(3N)</code>	Safe
<code>fn_compound_name_insert_comp(3N)</code>	Safe
<code>fn_compound_name_is_empty(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_compound_name_is_equal(3N)</code>	Safe
<code>fn_compound_name_is_prefix(3N)</code>	Safe
<code>fn_compound_name_is_suffix(3N)</code>	Safe
<code>fn_compound_name_last(3N)</code>	Safe
<code>fn_compound_name_next(3N)</code>	Safe
<code>fn_compound_name_prefix(3N)</code>	Safe
<code>fn_compound_name_prepend_comp(3N)</code>	Safe
<code>fn_compound_name_prev(3N)</code>	Safe
<code>fn_compound_name_suffix(3N)</code>	Safe
<code>FN_compound_name_t(3N)</code>	Safe
<code>fn_ctx_bind(3N)</code>	Safe
<code>fn_ctx_create_subcontext(3N)</code>	Safe
<code>fn_ctx_destroy_subcontext(3N)</code>	Safe
<code>fn_ctx_get_ref(3N)</code>	Safe
<code>fn_ctx_get_syntax_attrs(3N)</code>	Safe
<code>fn_ctx_handle_destroy(3N)</code>	Safe
<code>fn_ctx_handle_from_initial(3N)</code>	MT-Safe
<code>fn_ctx_handle_from_ref(3N)</code>	Safe
<code>fn_ctx_list_bindings(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_ctx_list_names(3N)</code>	Safe
<code>fn_ctx_lookup(3N)</code>	Safe
<code>fn_ctx_lookup_link(3N)</code>	Safe
<code>fn_ctx_rename(3N)</code>	Safe
<code>FN_ctx_t(3N)</code>	Safe
<code>fn_ctx_unbind(3N)</code>	Safe
<code>fn_multigetlist_destroy(3N)</code>	Safe
<code>fn_multigetlist_next(3N)</code>	Safe
<code>FN_multigetlist_t(3N)</code>	Safe
<code>fn_namelist_destroy(3N)</code>	Safe
<code>fn_namelist_next(3N)</code>	Safe
<code>FN_namelist_t(3N)</code>	Safe
<code>fn_ref_addrcount(3N)</code>	Safe
<code>fn_ref_addr_assign(3N)</code>	Safe
<code>fn_ref_addr_copy(3N)</code>	Safe
<code>fn_ref_addr_create(3N)</code>	Safe
<code>fn_ref_addr_data(3N)</code>	Safe
<code>fn_ref_addr_description(3N)</code>	Safe
<code>fn_ref_addr_destroy(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_ref_addr_length(3N)</code>	Safe
<code>FN_ref_addr_t(3N)</code>	Safe
<code>fn_ref_addr_type(3N)</code>	Safe
<code>fn_ref_append_addr(3N)</code>	Safe
<code>fn_ref_assign(3N)</code>	Safe
<code>fn_ref_copy(3N)</code>	Safe
<code>fn_ref_create(3N)</code>	Safe
<code>fn_ref_create_link(3N)</code>	Safe
<code>fn_ref_delete_addr(3N)</code>	Safe
<code>fn_ref_delete_all(3N)</code>	Safe
<code>fn_ref_description(3N)</code>	Safe
<code>fn_ref_destroy(3N)</code>	Safe
<code>fn_ref_first(3N)</code>	Safe
<code>fn_ref_insert_addr(3N)</code>	Safe
<code>fn_ref_is_link(3N)</code>	Safe
<code>fn_ref_link_name(3N)</code>	Safe
<code>fn_ref_next(3N)</code>	Safe
<code>fn_ref_prepend_addr(3N)</code>	Safe
<code>FN_ref_t(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_ref_type(3N)</code>	Safe
<code>fn_status_advance_by_name(3N)</code>	Safe
<code>fn_status_append_remaining_name(3N)</code>	Safe
<code>fn_status_append_resolved_name(3N)</code>	Safe
<code>fn_status_assign(3N)</code>	Safe
<code>fn_status_code(3N)</code>	Safe
<code>fn_status_copy(3N)</code>	Safe
<code>fn_status_create(3N)</code>	Safe
<code>fn_status_description(3N)</code>	Safe
<code>fn_status_destroy(3N)</code>	Safe
<code>fn_status_diagnostic_message(3N)</code>	Safe
<code>fn_status_is_success(3N)</code>	Safe
<code>fn_status_link_code(3N)</code>	Safe
<code>fn_status_link_diagnostic_message(3N)</code>	Safe
<code>fn_status_link_remaining_name(3N)</code>	Safe
<code>fn_status_link_resolved_name(3N)</code>	Safe
<code>fn_status_link_resolved_ref(3N)</code>	Safe
<code>fn_status_remaining_name(3N)</code>	Safe
<code>fn_status_resolved_name(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_status_resolved_ref(3N)</code>	Safe
<code>fn_status_set(3N)</code>	Safe
<code>fn_status_set_code(3N)</code>	Safe
<code>fn_status_set_diagnostic_message(3N)</code>	Safe
<code>fn_status_set_link_code(3N)</code>	Safe
<code>fn_status_set_link_diagnostic_message</code>	Safe
<code>fn_status_set_link_remaining_name(3N)</code>	Safe
<code>fn_status_set_link_resolved_name(3N)</code>	Safe
<code>fn_status_set_link_resolved_ref(3N)</code>	Safe
<code>fn_status_set_remaining_name(3N)</code>	Safe
<code>fn_status_set_resolved_name(3N)</code>	Safe
<code>fn_status_set_resolved_ref(3N)</code>	Safe
<code>fn_status_set_success(3N)</code>	Safe
<code>FN_status_t(3N)</code>	Safe
<code>fn_string_assign(3N)</code>	Safe
<code>fn_string_bytecount(3N)</code>	Safe
<code>fn_string_charcount(3N)</code>	Safe
<code>fn_string_code_set(3N)</code>	Safe
<code>fn_string_compare(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fn_string_compare_substring(3N)</code>	Safe
<code>fn_string_contents(3N)</code>	Safe
<code>fn_string_copy(3N)</code>	Safe
<code>fn_string_create(3N)</code>	Safe
<code>fn_string_destroy(3N)</code>	Safe
<code>fn_string_from_composite_name(3N)</code>	Safe
<code>fn_string_from_compound_name(3N)</code>	Safe
<code>fn_string_from_contents(3N)</code>	Safe
<code>fn_string_from_str(3N)</code>	Safe
<code>fn_string_from_strings(3N)</code>	Safe
<code>fn_string_from_str_n(3N)</code>	Safe
<code>fn_string_from_substring(3N)</code>	Safe
<code>fn_string_is_empty(3N)</code>	Safe
<code>fn_string_next_substring(3N)</code>	Safe
<code>fn_string_prev_substring(3N)</code>	Safe
<code>fn_string_str(3N)</code>	Safe
<code>FN_string_t(3N)</code>	Safe
<code>fn_valuelist_destroy(3N)</code>	Safe
<code>fn_valuelist_next(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
FN_valuelist_t(3N)	Safe
fopen(3S)	MT-Safe
forms(3X)	Unsafe
form_cursor(3X)	Unsafe
form_data(3X)	Unsafe
form_driver(3X)	Unsafe
form_field(3X)	Unsafe
form_fields(3X)	Unsafe
form_fieldtype(3X)	Unsafe
form_field_attributes(3X)	Unsafe
form_field_buffer(3X)	Unsafe
form_field_info(3X)	Unsafe
form_field_just(3X)	Unsafe
form_field_new(3X)	Unsafe
form_field_opts(3X)	Unsafe
form_field_userptr(3X)	Unsafe
form_field_validation(3X)	Unsafe
form_hook(3X)	Unsafe
form_init(3X)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>form_new(3X)</code>	Unsafe
<code>form_new_page(3X)</code>	Unsafe
<code>form_opts(3X)</code>	Unsafe
<code>form_opts_off(3X)</code>	Unsafe
<code>form_opts_on(3X)</code>	Unsafe
<code>form_page(3X)</code>	Unsafe
<code>form_post(3X)</code>	Unsafe
<code>form_sub(3X)</code>	Unsafe
<code>form_term(3X)</code>	Unsafe
<code>form_userptr(3X)</code>	Unsafe
<code>form_win(3X)</code>	Unsafe
<code>fpclass(3C)</code>	MT-Safe
<code>fpgetmask(3C)</code>	MT-Safe
<code>fpgetround(3C)</code>	MT-Safe
<code>fpgetsticky(3C)</code>	MT-Safe
<code>fprintf(3S)</code>	MT-Safe except with <code>setlocale()</code>
<code>fpsetmask(3C)</code>	MT-Safe
<code>fpsetround(3C)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fpsetsticky(3C)</code>	MT-Safe
<code>fputc(3S)</code>	MT-Safe
<code>fputs(3S)</code>	MT-Safe
<code>fputwc(3I)</code>	MT-Safe
<code>fputws(3I)</code>	MT-Safe
<code>fread(3S)</code>	MT-Safe
<code>free(3C)</code>	Safe
<code>free(3X)</code>	Safe
<code>freenetconfignt(3N)</code>	MT-Safe
<code>free_field(3X)</code>	Unsafe
<code>free_fieldtype(3X)</code>	Unsafe
<code>free_form(3X)</code>	Unsafe
<code>free_item(3X)</code>	Unsafe
<code>free_menu(3X)</code>	Unsafe
<code>freopen(3S)</code>	MT-Safe
<code>frexp(3C)</code>	MT-Safe
<code>fscanf(3S)</code>	MT-Safe
<code>fseek(3S)</code>	MT-Safe
<code>fsetpos(3C)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>fsync(3C)</code>	Async-Signal-Safe
<code>ftell(3S)</code>	MT-Safe
<code>ftok(3C)</code>	MT-Safe
<code>ftruncate(3C)</code>	MT-Safe
<code>ftrylockfile(3S)</code>	MT-Safe
<code>ftw(3C)</code>	Safe
<code>func_to_decimal(3)</code>	MT-Safe
<code>funlockfile(3S)</code>	MT-Safe
<code>fwrite(3S)</code>	MT-Safe
<code>gconvert(3)</code>	MT-Safe
<code>gcv(3)</code>	MT-Safe
<code>gcv(3C)</code>	Unsafe
<code>getacdir(3)</code>	Safe
<code>getacflg(3)</code>	Safe
<code>getacinfo(3)</code>	Safe
<code>getacmin(3)</code>	Safe
<code>getacna(3)</code>	Safe
<code>getauclassent(3)</code>	Unsafe
<code>getauclassent_r(3)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
getauclassnam(3)	Unsafe
getauclassnam_r(3)	MT-Safe
getauditflags(3)	MT-Safe
getauditflagsbin(3)	MT-Safe
getauditflagschar(3)	MT-Safe
getauevent(3)	Unsafe
getauevent_r(3)	MT-Safe
getauevnam(3)	Unsafe
getauevnam_r(3)	MT-Safe
getauevnonam(3)	MT-Safe
getauevnum(3)	Unsafe
getauevnum_r(3)	MT-Safe
getauuserent(3)	Unsafe
getauusernam(3)	Unsafe
getbegyx(3X)	Unsafe
getc(3S)	MT-Safe
getch(3X)	Unsafe
getchar(3S)	MT-Safe
getcwd(3C)	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>getdate(3C)</code>	MT-Safe
<code>getenv(3C)</code>	Safe
<code>getfauditflags(3)</code>	MT-Safe
<code>getgrent(3C)</code>	Unsafe, use <code>getgrent_r()</code>
<code>getgrgid(3C)</code>	Unsafe, use <code>getgrgid_r()</code>
<code>getgrnam(3C)</code>	Unsafe, use <code>getgrnam_r()</code>
<code>gethostbyaddr(3N)</code>	Unsafe, use <code>gethostbyaddr_r()</code>
<code>gethostbyname(3N)</code>	Unsafe, use <code>gethostbyname_r()</code>
<code>gethrtime(3C)</code>	MT-Safe
<code>gethrvtime(3C)</code>	MT-Safe
<code>getlogin(3C)</code>	Unsafe, use <code>getlogin_r()</code>
<code>getmaxyx(3X)</code>	Unsafe
<code>getmntany(3C)</code>	Safe
<code>getmntent(3C)</code>	Safe
<code>getnetbyaddr(3N)</code>	Unsafe, use <code>getnetbyaddr_r()</code>
<code>getnetbyname(3N)</code>	Unsafe, use <code>getnetbyname_r()</code>
<code>getnetconfig(3N)</code>	MT-Safe
<code>getnetconfigent(3N)</code>	MT-Safe
<code>getnetgrent(3N)</code>	Unsafe, use <code>getnetgrent_r()</code>

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>getnetname(3N)</code>	MT-Safe
<code>getnetpath(3N)</code>	MT-Safe
<code>getnwstr(3X)</code>	Unsafe
<code>getopt(3C)</code>	Unsafe
<code>getparyx(3X)</code>	Unsafe
<code>getpass(3C)</code>	Unsafe
<code>getpeername(3N)</code>	Safe
<code>getprotobyname(3N)</code>	Unsafe, use <code>getprotobyname_r()</code>
<code>getprotobynumber(3N)</code>	Unsafe, use <code>getprotobynumber_r()</code>
<code>getprotoent(3N)</code>	Unsafe, use <code>getprotoent_r()</code>
<code>getpublickey(3N)</code>	Safe
<code>getpw(3C)</code>	Safe
<code>getpwent(3C)</code>	Unsafe, use <code>getpwent_r()</code>
<code>getpwnam(3C)</code>	Unsafe, use <code>getpwnam_r()</code>
<code>getpwuid(3C)</code>	Unsafe, use <code>getpwuid_r()</code>
<code>getrpcbyname(3N)</code>	Unsafe, use <code>getrpcbyname_r()</code>
<code>getrpcbynumber(3N)</code>	Unsafe, use <code>getrpcbynumber_r()</code>
<code>getrpcent(3N)</code>	Unsafe, use <code>getrpcent_r()</code>

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>getrpcport(3N)</code>	Unsafe
<code>gets(3S)</code>	MT-Safe
<code>getsecretkey(3N)</code>	Safe
<code>getservbyname(3N)</code>	Unsafe, use <code>getservbyname_r()</code>
<code>getservbyport(3N)</code>	Unsafe, use <code>getservbyport_r()</code>
<code>getservent(3N)</code>	Unsafe, use <code>getservent_r()</code>
<code>getsockname(3N)</code>	Safe
<code>getsockopt(3N)</code>	Safe
<code>getspent(3C)</code>	Unsafe, use <code>getspent_r()</code>
<code>getspnam(3C)</code>	Unsafe, use <code>getspnam_r()</code>
<code>getstr(3X)</code>	Unsafe
<code>getsubopt(3C)</code>	MT-Safe
<code>getsyx(3X)</code>	Unsafe
<code>gettext(3I)</code>	Safe with exceptions
<code>gettimeofday(3C)</code>	MT-Safe
<code>gettxt(3C)</code>	Safe with exceptions
<code>getutent(3C)</code>	Unsafe
<code>getutid(3C)</code>	Unsafe
<code>getutline(3C)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
getutmp(3C)	Unsafe
getutmpx(3C)	Unsafe
getutxent(3C)	Unsafe
getutxid(3C)	Unsafe
getutxline(3C)	Unsafe
getvfsany(3C)	Safe
getvfsent(3C)	Safe
getvfssfile(3C)	Safe
getvfsspec(3C)	Safe
getw(3S)	MT-Safe
getwc(3I)	MT-Safe
getwch(3X)	Unsafe
getwchar(3I)	MT-Safe
getwidth(3I)	MT-Safe with exceptions
getwin(3X)	Unsafe
getws(3I)	MT-Safe
getwstr(3X)	Unsafe
getyx(3X)	Unsafe
get_myaddress(3N)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>gmatch(3G)</code>	MT-Safe
<code>gmtime(3C)</code>	Unsafe, use <code>gmtime_r()</code>
<code>grantpt(3C)</code>	Safe
<code>gsignal(3C)</code>	Unsafe
<code>halfdelay(3X)</code>	Unsafe
<code>hasmntopt(3C)</code>	Safe
<code>has_colors(3X)</code>	Unsafe
<code>has_ic(3X)</code>	Unsafe
<code>has_il(3X)</code>	Unsafe
<code>havedisk(3N)</code>	MT-Safe
<code>hcreate(3C)</code>	Safe
<code>hdestroy(3C)</code>	Safe
<code>hide_panel(3X)</code>	Unsafe
<code>host2netname(3N)</code>	MT-Safe
<code>hsearch(3C)</code>	Safe
<code>htonl(3N)</code>	Safe
<code>htons(3N)</code>	Safe
<code>hyperbolic(3M)</code>	MT-Safe
<code>hypot(3M)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>iconv(3)</code>	MT-Safe
<code>iconv_close(3)</code>	MT-Safe
<code>iconv_open(3)</code>	MT-Safe
<code>idcok(3X)</code>	Unsafe
<code>idlck(3X)</code>	Unsafe
<code>ieee_functions(3M)</code>	MT-Safe
<code>ieee_test(3M)</code>	MT-Safe
<code>ilogb(3M)</code>	MT-Safe
<code>immedok(3X)</code>	Unsafe
<code>inch(3X)</code>	Unsafe
<code>inchnstr(3X)</code>	Unsafe
<code>inchstr(3X)</code>	Unsafe
<code>inet(3N)</code>	Safe
<code>inet_addr(3N)</code>	Safe
<code>inet_lnaof(3N)</code>	Safe
<code>inet_makeaddr(3N)</code>	Safe
<code>inet_netof(3N)</code>	Safe
<code>inet_network(3N)</code>	Safe
<code>inet_ntoa(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>initgroups(3C)</code>	Unsafe
<code>initscr(3X)</code>	Unsafe
<code>init_color(3X)</code>	Unsafe
<code>init_pair(3X)</code>	Unsafe
<code>innstr(3X)</code>	Unsafe
<code>innwstr(3X)</code>	Unsafe
<code>insch(3X)</code>	Unsafe
<code>insdelln(3X)</code>	Unsafe
<code>insertln(3X)</code>	Unsafe
<code>insnstr(3X)</code>	Unsafe
<code>insnwstr(3X)</code>	Unsafe
<code>insque(3C)</code>	Unsafe
<code>insstr(3X)</code>	Unsafe
<code>instr(3X)</code>	Unsafe
<code>inswch(3X)</code>	Unsafe
<code>inswstr(3X)</code>	Unsafe
<code>intrflush(3X)</code>	Unsafe
<code>inwch(3X)</code>	Unsafe
<code>inwchnstr(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>inwchstr(3X)</code>	Unsafe
<code>inwstr(3X)</code>	Unsafe
<code>isalnum(3C)</code>	MT-Safe with exceptions
<code>isalpha(3C)</code>	MT-Safe with exceptions
<code>isascii(3C)</code>	MT-Safe with exceptions
<code>isastream(3C)</code>	MT-Safe
<code>iscntrl(3C)</code>	MT-Safe with exceptions
<code>isdigit(3C)</code>	MT-Safe with exceptions
<code>isencrypt(3G)</code>	MT-Safe
<code>isendwin(3X)</code>	Unsafe
<code>isenglish(3I)</code>	MT-Safe with exceptions
<code>isgraph(3C)</code>	MT-Safe with exceptions
<code>isideogram(3I)</code>	MT-Safe with exceptions
<code>islower(3C)</code>	MT-Safe with exceptions
<code>isnan(3C)</code>	MT-Safe
<code>isnan(3M)</code>	MT-Safe
<code>isnand(3C)</code>	MT-Safe
<code>isnanf(3C)</code>	MT-Safe
<code>isnumber(3I)</code>	MT-Safe with exceptions

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>isphonogram(3I)</code>	MT-Safe with exceptions
<code>isprint(3C)</code>	MT-Safe with exceptions
<code>ispunct(3C)</code>	MT-Safe with exceptions
<code>isspace(3C)</code>	MT-Safe with exceptions
<code>isspecial(3I)</code>	MT-Safe with exceptions
<code>isupper(3C)</code>	MT-Safe with exceptions
<code>iswalnum(3I)</code>	MT-Safe with exceptions
<code>iswalpha(3I)</code>	MT-Safe with exceptions
<code>iswascii(3I)</code>	MT-Safe with exceptions
<code>iswcntrl(3I)</code>	MT-Safe with exceptions
<code>iswctype(3I)</code>	MT-Safe
<code>iswdigit(3I)</code>	MT-Safe with exceptions
<code>iswgraph(3I)</code>	MT-Safe with exceptions
<code>iswlower(3I)</code>	MT-Safe with exceptions
<code>iswprint(3I)</code>	MT-Safe with exceptions
<code>iswpunct(3I)</code>	MT-Safe with exceptions
<code>iswspace(3I)</code>	MT-Safe with exceptions
<code>iswupper(3I)</code>	MT-Safe with exceptions
<code>iswxdigit(3I)</code>	MT-Safe with exceptions

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>isxdigit(3C)</code>	MT-Safe with exceptions
<code>is_linetouched(3X)</code>	Unsafe
<code>is_wintouched(3X)</code>	Unsafe
<code>item_count(3X)</code>	Unsafe
<code>item_description(3X)</code>	Unsafe
<code>item_index(3X)</code>	Unsafe
<code>item_init(3X)</code>	Unsafe
<code>item_name(3X)</code>	Unsafe
<code>item_opts(3X)</code>	Unsafe
<code>item_opts_off(3X)</code>	Unsafe
<code>item_opts_on(3X)</code>	Unsafe
<code>item_term(3X)</code>	Unsafe
<code>item_userptr(3X)</code>	Unsafe
<code>item_value(3X)</code>	Unsafe
<code>item_visible(3X)</code>	Unsafe
<code>j0(3M)</code>	MT-Safe
<code>j1(3M)</code>	MT-Safe
<code>jn(3M)</code>	MT-Safe
<code>jrand48(3C)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
kerberos(3N)	Unsafe
kerberos_rpc(3N)	Unsafe
keyname(3X)	Unsafe
keypad(3X)	Unsafe
key_decryptsession(3N)	MT-Safe
key_encryptsession(3N)	MT-Safe
key_gendes(3N)	MT-Safe
key_secretkey_is_set(3N)	MT-Safe
key_setsecret(3N)	MT-Safe
killchar(3X)	Unsafe
krb_get_admhst(3N)	Unsafe
krb_get_cred(3N)	Unsafe
krb_get_krbhst(3N)	Unsafe
krb_get_lrealm(3N)	Unsafe
krb_get_phost(3N)	Unsafe
krb_kntoln(3N)	Unsafe
krb_mk_err(3N)	Unsafe
krb_mk_req(3N)	Unsafe
krb_mk_safe(3N)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>krb_net_read(3N)</code>	Unsafe
<code>krb_net_write(3N)</code>	Unsafe
<code>krb_rd_err(3N)</code>	Unsafe
<code>krb_rd_req(3N)</code>	Unsafe
<code>krb_rd_safe(3N)</code>	Unsafe
<code>krb_realmofhost(3N)</code>	Unsafe
<code>krb_recvauth(3N)</code>	Unsafe
<code>krb_sendauth(3N)</code>	Unsafe
<code>krb_set_key(3N)</code>	Unsafe
<code>krb_set_tkt_string(3N)</code>	Unsafe
<code>kvm_close(3K)</code>	Unsafe
<code>kvm_getcmd(3K)</code>	Unsafe
<code>kvm_getproc(3K)</code>	Unsafe
<code>kvm_getu(3K)</code>	Unsafe
<code>kvm_kread(3K)</code>	Unsafe
<code>kvm_kwrite(3K)</code>	Unsafe
<code>kvm_nextproc(3K)</code>	Unsafe
<code>kvm_nlist(3K)</code>	Unsafe
<code>kvm_open(3K)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
kvm_read(3K)	Unsafe
kvm_setproc(3K)	Unsafe
kvm_uread(3K)	Unsafe
kvm_uwrite(3K)	Unsafe
kvm_write(3K)	Unsafe
l64a(3C)	MT-Safe
label(3)	Safe
labs(3C)	MT-Safe
lckpddf(3C)	MT-Safe
lcong48(3C)	Safe
ldexp(3C)	MT-Safe
ldiv(3C)	MT-Safe
leaveok(3X)	Unsafe
lfind(3C)	Safe
lfmt(3C)	MT-Safe
lgamma(3M)	Unsafe, use <code>lgamma_r()</code>
libpthread(3T)	Fork1-Safe,MT-Safe,Async-Signal-Safe
libthread(3T)	Fork1-Safe,MT-Safe,Async-Signal-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>line(3)</code>	Safe
<code>link_field(3X)</code>	Unsafe
<code>link_fieldtype(3X)</code>	Unsafe
<code>linmod(3)</code>	Safe
<code>lio_listio(3R)</code>	MT-Safe
<code>listen(3N)</code>	Safe
<code>llabs(3C)</code>	MT-Safe
<code>lldiv(3C)</code>	MT-Safe
<code>lltostr(3C)</code>	MT-Safe
<code>localeconv(3C)</code>	Safe with exceptions
<code>localtime(3C)</code>	Unsafe, use <code>localtime_r()</code>
<code>lockf(3C)</code>	MT-Safe
<code>log(3M)</code>	MT-Safe
<code>log10(3M)</code>	MT-Safe
<code>log1p(3M)</code>	MT-Safe
<code>logb(3C)</code>	MT-Safe
<code>logb(3M)</code>	MT-Safe
<code>longjmp(3C)</code>	Unsafe
<code>longname(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>lrand48(3C)</code>	Safe
<code>lsearch(3C)</code>	Safe
<code>madvise(3)</code>	MT-Safe
<code>maillock(3X)</code>	Unsafe
<code>major(3C)</code>	MT-Safe
<code>makecontext(3C)</code>	MT-Safe
<code>makedev(3C)</code>	MT-Safe
<code>mallinfo(3X)</code>	Safe
<code>malloc(3C)</code>	Safe
<code>malloc(3X)</code>	Safe
<code>mallopt(3X)</code>	Safe
<code>mapmalloc(3X)</code>	Safe
<code>matherr(3M)</code>	MT-Safe
<code>mbchar(3C)</code>	MT-Safe with exceptions
<code>mblen(3C)</code>	MT-Safe with exceptions
<code>mbstowcs(3C)</code>	MT-Safe with exceptions
<code>mbstring(3C)</code>	MT-Safe with exceptions
<code>mbtowc(3C)</code>	MT-Safe with exceptions
<code>media_findname(3X)</code>	MT-Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>media_getattr(3X)</code>	MT-Safe
<code>media_setattr(3X)</code>	MT-Safe
<code>memalign(3C)</code>	Safe
<code>memccpy(3C)</code>	MT-Safe
<code>memchr(3C)</code>	MT-Safe
<code>memcmp(3C)</code>	MT-Safe
<code>memcpy(3C)</code>	MT-Safe
<code>memmove(3C)</code>	MT-Safe
<code>memory(3C)</code>	MT-Safe
<code>memset(3C)</code>	MT-Safe
<code>menus(3X)</code>	Unsafe
<code>menu_attributes(3X)</code>	Unsafe
<code>menu_back(3X)</code>	Unsafe
<code>menu_cursor(3X)</code>	Unsafe
<code>menu_driver(3X)</code>	Unsafe
<code>menu_fore(3X)</code>	Unsafe
<code>menu_format(3X)</code>	Unsafe
<code>menu_grey(3X)</code>	Unsafe
<code>menu_hook(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>menu_init(3X)</code>	Unsafe
<code>menu_items(3X)</code>	Unsafe
<code>menu_item_current(3X)</code>	Unsafe
<code>menu_item_name(3X)</code>	Unsafe
<code>menu_item_new(3X)</code>	Unsafe
<code>menu_item_opts(3X)</code>	Unsafe
<code>menu_item_userptr(3X)</code>	Unsafe
<code>menu_item_value(3X)</code>	Unsafe
<code>menu_item_visible(3X)</code>	Unsafe
<code>menu_mark(3X)</code>	Unsafe
<code>menu_new(3X)</code>	Unsafe
<code>menu_opts(3X)</code>	Unsafe
<code>menu_opts_off(3X)</code>	Unsafe
<code>menu_opts_on(3X)</code>	Unsafe
<code>menu_pad(3X)</code>	Unsafe
<code>menu_pattern(3X)</code>	Unsafe
<code>menu_post(3X)</code>	Unsafe
<code>menu_sub(3X)</code>	Unsafe
<code>menu_term(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>menu_userptr(3X)</code>	Unsafe
<code>menu_win(3X)</code>	Unsafe
<code>meta(3X)</code>	Unsafe
<code>minor(3C)</code>	MT-Safe
<code>mkdirp(3G)</code>	MT-Safe
<code>mkfifo(3C)</code>	MT-Safe, Async-Signal-Safe
<code>mktemp(3C)</code>	Safe
<code>mktime(3C)</code>	Unsafe
<code>mlock(3C)</code>	MT-Safe
<code>monitor(3C)</code>	Safe
<code>move(3)</code>	Safe
<code>move(3X)</code>	Unsafe
<code>movenextch(3X)</code>	Unsafe
<code>moveprevch(3X)</code>	Unsafe
<code>move_field(3X)</code>	Unsafe
<code>move_panel(3X)</code>	Unsafe
<code>mq_close(3R)</code>	MT-Safe
<code>mq_getattr(3R)</code>	MT-Safe
<code>mq_notify(3R)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>mq_open(3R)</code>	MT-Safe
<code>mq_receive(3R)</code>	MT-Safe
<code>mq_send(3R)</code>	MT-Safe
<code>mq_setattr(3R)</code>	MT-Safe
<code>mq_unlink(3R)</code>	MT-Safe
<code>rand48(3C)</code>	Safe
<code>msync(3C)</code>	MT-Safe
<code>munlock(3C)</code>	MT-Safe
<code>munlockall(3C)</code>	MT-Safe
<code>mutex(3T)</code>	MT-Safe
<code>mutex_destroy(3T)</code>	MT-Safe
<code>mutex_init(3T)</code>	MT-Safe
<code>mutex_lock(3T)</code>	MT-Safe
<code>mutex_trylock(3T)</code>	MT-Safe
<code>mutex_unlock(3T)</code>	MT-Safe
<code>mvaddch(3X)</code>	Unsafe
<code>mvaddchnstr(3X)</code>	Unsafe
<code>mvaddchstr(3X)</code>	Unsafe
<code>mvaddnstr(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>mvaddnwstr(3X)</code>	Unsafe
<code>mvaddstr(3X)</code>	Unsafe
<code>mvaddwch(3X)</code>	Unsafe
<code>mvaddwchnstr(3X)</code>	Unsafe
<code>mvaddwchstr(3X)</code>	Unsafe
<code>mvaddwstr(3X)</code>	Unsafe
<code>mvcur(3X)</code>	Unsafe
<code>mvdelch(3X)</code>	Unsafe
<code>mvderwin(3X)</code>	Unsafe
<code>mvgetch(3X)</code>	Unsafe
<code>mvgetnwstr(3X)</code>	Unsafe
<code>mvgetstr(3X)</code>	Unsafe
<code>mvgetwch(3X)</code>	Unsafe
<code>mvgetwstr(3X)</code>	Unsafe
<code>mvinch(3X)</code>	Unsafe
<code>mvinchnstr(3X)</code>	Unsafe
<code>mvinchstr(3X)</code>	Unsafe
<code>mvinnstr(3X)</code>	Unsafe
<code>mvinnwstr(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>mvinsch(3X)</code>	Unsafe
<code>mvinsnstr(3X)</code>	Unsafe
<code>mvinsnWSTR(3X)</code>	Unsafe
<code>mvinsstr(3X)</code>	Unsafe
<code>mvinstr(3X)</code>	Unsafe
<code>mvinswch(3X)</code>	Unsafe
<code>mvinsWSTR(3X)</code>	Unsafe
<code>mvinwch(3X)</code>	Unsafe
<code>mvinwchnstr(3X)</code>	Unsafe
<code>mvinwchstr(3X)</code>	Unsafe
<code>mvinWSTR(3X)</code>	Unsafe
<code>mvprintw(3X)</code>	Unsafe
<code>mvscanw(3X)</code>	Unsafe
<code>mvwaddch(3X)</code>	Unsafe
<code>mvwaddchnstr(3X)</code>	Unsafe
<code>mvwaddchstr(3X)</code>	Unsafe
<code>mvwaddnstr(3X)</code>	Unsafe
<code>mvwaddnwstr(3X)</code>	Unsafe
<code>mvwaddstr(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>mvwaddwch(3X)</code>	Unsafe
<code>mvwaddwchnstr(3X)</code>	Unsafe
<code>mvwaddwchstr(3X)</code>	Unsafe
<code>mvwaddwstr(3X)</code>	Unsafe
<code>mvwdelch(3X)</code>	Unsafe
<code>mvwgetch(3X)</code>	Unsafe
<code>mvwgetnwstr(3X)</code>	Unsafe
<code>mvwgetstr(3X)</code>	Unsafe
<code>mvwgetwch(3X)</code>	Unsafe
<code>mvwgetwstr(3X)</code>	Unsafe
<code>mvwin(3X)</code>	Unsafe
<code>mvwinch(3X)</code>	Unsafe
<code>mvwinchnstr(3X)</code>	Unsafe
<code>mvwinchstr(3X)</code>	Unsafe
<code>mvwinnstr(3X)</code>	Unsafe
<code>mvwinnwstr(3X)</code>	Unsafe
<code>mvwinsch(3X)</code>	Unsafe
<code>mvwinsnstr(3X)</code>	Unsafe
<code>mvwinsnwstr(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>mvwinsstr(3X)</code>	Unsafe
<code>mvwinstr(3X)</code>	Unsafe
<code>mvwinswch(3X)</code>	Unsafe
<code>mvwinswstr(3X)</code>	Unsafe
<code>mvwinwch(3X)</code>	Unsafe
<code>mvwinwchnstr(3X)</code>	Unsafe
<code>mvwinwchstr(3X)</code>	Unsafe
<code>mvwinwstr(3X)</code>	Unsafe
<code>mvwprintw(3X)</code>	Unsafe
<code>mvwscanw(3X)</code>	Unsafe
<code>nanosleep(3R)</code>	MT-Safe
<code>napms(3X)</code>	Unsafe
<code>nc_perror(3N)</code>	MT-Safe
<code>nc_serror(3N)</code>	MT-Safe
<code>ndbm(3)</code>	Unsafe
<code>netdir(3N)</code>	MT-Safe
<code>netdir_free(3N)</code>	MT-Safe
<code>netdir_getbyaddr(3N)</code>	MT-Safe
<code>netdir_getbyname(3N)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>netdir_mergeaddr(3N)</code>	MT-Safe
<code>netdir_options(3N)</code>	MT-Safe
<code>netdir_perror(3N)</code>	MT-Safe
<code>netdir_sperror(3N)</code>	MT-Safe
<code>netname2host(3N)</code>	MT-Safe
<code>netname2user(3N)</code>	MT-Safe
<code>newpad(3X)</code>	Unsafe
<code>newterm(3X)</code>	Unsafe
<code>newwin(3X)</code>	Unsafe
<code>new_field(3X)</code>	Unsafe
<code>new_fieldtype(3X)</code>	Unsafe
<code>new_form(3X)</code>	Unsafe
<code>new_item(3X)</code>	Unsafe
<code>new_menu(3X)</code>	Unsafe
<code>new_page(3X)</code>	Unsafe
<code>new_panel(3X)</code>	Unsafe
<code>nextafter(3C)</code>	MT-Safe
<code>nextafter(3M)</code>	MT-Safe
<code>nftw(3C)</code>	Safe with exceptions

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>nis_add(3N)</code>	MT-Safe
<code>nis_clone_object(3N)</code>	Safe
<code>nis_creategroup(3N)</code>	MT-Safe
<code>nis_db(3N)</code>	Unsafe
<code>nis_destroygroup(3N)</code>	MT-Safe
<code>nis_destroy_object(3N)</code>	Safe
<code>nis_dir_cmp(3N)</code>	Safe
<code>nis_domain_of(3N)</code>	Safe
<code>nis_error(3N)</code>	Safe
<code>nis_first_entry(3N)</code>	MT-Safe
<code>nis_freenames(3N)</code>	Safe
<code>nis_freeresult(3N)</code>	MT-Safe
<code>nis_freeservlist(3N)</code>	MT-Safe
<code>nis_freetags(3N)</code>	MT-Safe
<code>nis_getnames(3N)</code>	Safe
<code>nis_getservlist(3N)</code>	MT-Safe
<code>nis_groups(3N)</code>	MT-Safe
<code>nis_ismember(3N)</code>	MT-Safe
<code>nis_leaf_of(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>nis_terror(3N)</code>	Safe
<code>nis_list(3N)</code>	MT-Safe
<code>nis_local_directory(3N)</code>	MT-Safe
<code>nis_local_group(3N)</code>	MT-Safe
<code>nis_local_host(3N)</code>	MT-Safe
<code>nis_local_names(3N)</code>	MT-Safe
<code>nis_local_principal(3N)</code>	MT-Safe
<code>nis_lookup(3N)</code>	MT-Safe
<code>nis_map_group(3N)</code>	MT-Safe
<code>nis_mkdir(3N)</code>	MT-Safe
<code>nis_modify(3N)</code>	MT-Safe
<code>nis_modify_entry(3N)</code>	MT-Safe
<code>nis_names(3N)</code>	MT-Safe
<code>nis_name_of(3N)</code>	Safe
<code>nis_next_entry(3N)</code>	MT-Safe
<code>nis_perror(3N)</code>	Safe
<code>nis_ping(3N)</code>	MT-Safe
<code>nis_print_group_entry(3N)</code>	MT-Safe
<code>nis_print_object(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>nis_remove(3N)</code>	MT-Safe
<code>nis_removemember(3N)</code>	MT-Safe
<code>nis_remove_entry(3N)</code>	MT-Safe
<code>nis_rmdir(3N)</code>	MT-Safe
<code>nis_server(3N)</code>	MT-Safe
<code>nis_servstate(3N)</code>	MT-Safe
<code>nis_sperrno(3N)</code>	Safe
<code>nis_sperror(3N)</code>	Safe
<code>nis_sperror_r(3N)</code>	Safe
<code>nis_stats(3N)</code>	MT-Safe
<code>nis_subr(3N)</code>	Safe
<code>nis_tables(3N)</code>	MT-Safe
<code>nis_verifygroup(3N)</code>	MT-Safe
<code>nl(3X)</code>	Unsafe
<code>nlist(3E)</code>	Safe
<code>nlsgetcall(3N)</code>	Unsafe
<code>nlsprovider(3N)</code>	Unsafe
<code>nlsrequest(3N)</code>	Unsafe
<code>nl_langinfo(3C)</code>	Safe with exceptions

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>nocbreak(3X)</code>	Unsafe
<code>nodeLAY(3X)</code>	Unsafe
<code>noecho(3X)</code>	Unsafe
<code>nonl(3X)</code>	Unsafe
<code>noqiflush(3X)</code>	Unsafe
<code>noraw(3X)</code>	Unsafe
<code>NOTE(3X)</code>	Safe
<code>notimeout(3X)</code>	Unsafe
<code>nrand48(3C)</code>	Safe
<code>ntohl(3N)</code>	Safe
<code>ntohs(3N)</code>	Safe
<code>offsetof(3C)</code>	MT-Safe
<code>opendir(3C)</code>	Safe
<code>openlog(3)</code>	Safe
<code>openpl(3)</code>	Safe
<code>openvt(3)</code>	Safe
<code>overlay(3X)</code>	Unsafe
<code>overwrite(3X)</code>	Unsafe
<code>p2close(3G)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>p2open(3G)</code>	Unsafe
<code>pair_content(3X)</code>	Unsafe
<code>panels(3X)</code>	Unsafe
<code>panel_above(3X)</code>	Unsafe
<code>panel_below(3X)</code>	Unsafe
<code>panel_hidden(3X)</code>	Unsafe
<code>panel_move(3X)</code>	Unsafe
<code>panel_new(3X)</code>	Unsafe
<code>panel_show(3X)</code>	Unsafe
<code>panel_top(3X)</code>	Unsafe
<code>panel_update(3X)</code>	Unsafe
<code>panel_userptr(3X)</code>	Unsafe
<code>panel_window(3X)</code>	Unsafe
<code>pathfind(3G)</code>	MT-Safe
<code>pclose(3S)</code>	Unsafe
<code>pechochar(3X)</code>	Unsafe
<code>pechowchar(3X)</code>	Unsafe
<code>perror(3C)</code>	MT-Safe
<code>pfmt(3C)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>plot(3)</code>	Safe
<code>pmap_getmaps(3N)</code>	Unsafe
<code>pmap_getport(3N)</code>	Unsafe
<code>pmap_rmtcall(3N)</code>	Unsafe
<code>pmap_set(3N)</code>	Unsafe
<code>pmap_unset(3N)</code>	Unsafe
<code>pnoutrefresh(3X)</code>	Unsafe
<code>point(3)</code>	Safe
<code>popen(3S)</code>	Unsafe
<code>post_form(3X)</code>	Unsafe
<code>post_menu(3X)</code>	Unsafe
<code>pos_form_cursor(3X)</code>	Unsafe
<code>pos_menu_cursor(3X)</code>	Unsafe
<code>pow(3M)</code>	MT-Safe
<code>prefresh(3X)</code>	Unsafe
<code>printf(3S)</code>	MT-Safe except with <code>setlocale()</code>
<code>printw(3X)</code>	Unsafe
<code>psiginfo(3C)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>psignal(3C)</code>	Safe
<code>pthread(3T)</code>	Fork1-Safe, MT-Safe, Async-Signal-Safe
<code>pthread_atfork(3T)</code>	MT-Safe
<code>pthread_attr_destroy(3T)</code>	MT-Safe
<code>pthread_attr_getdetachstate(3T)</code>	MT-Safe
<code>pthread_attr_getinheritsched(3T)</code>	MT-Safe
<code>pthread_attr_getschedparam(3T)</code>	MT-Safe
<code>pthread_attr_getschedpolicy(3T)</code>	MT-Safe
<code>pthread_attr_getscope(3T)</code>	MT-Safe
<code>pthread_attr_getstackaddr(3T)</code>	MT-Safe
<code>pthread_attr_getstacksize(3T)</code>	MT-Safe
<code>pthread_attr_init(3T)</code>	MT-Safe
<code>pthread_attr_setdetachstate(3T)</code>	MT-Safe
<code>pthread_attr_setscope(3T)</code>	MT-Safe
<code>pthread_attr_setstackaddr(3T)</code>	MT-Safe
<code>pthread_attr_setstacksize(3T)</code>	MT-Safe
<code>pthread_cancel(3T)</code>	MT-Safe
<code>pthread_cleanup_pop(3T)</code>	MT-Safe
<code>pthread_cleanup_push(3T)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>pthread_condattr_destroy(3T)</code>	MT-Safe
<code>pthread_condattr_getpshared(3T)</code>	MT-Safe
<code>pthread_condattr_init(3T)</code>	MT-Safe
<code>pthread_condattr_setpshared(3T)</code>	MT-Safe
<code>pthread_cond_broadcast(3T)</code>	MT-Safe
<code>pthread_cond_destroy(3T)</code>	MT-Safe
<code>pthread_cond_init(3T)</code>	MT-Safe
<code>pthread_cond_signal(3T)</code>	MT-Safe
<code>pthread_cond_timedwait(3T)</code>	MT-Safe
<code>pthread_cond_wait(3T)</code>	MT-Safe
<code>pthread_create(3T)</code>	MT-Safe
<code>pthread_detach(3T)</code>	MT-Safe
<code>pthread_equal(3T)</code>	MT-Safe
<code>pthread_exit(3T)</code>	MT-Safe
<code>pthread_getschedparam(3T)</code>	MT-Safe
<code>pthread_getspecific(3T)</code>	MT-Safe
<code>pthread_join(3T)</code>	MT-Safe
<code>pthread_key_create(3T)</code>	MT-Safe
<code>pthread_key_delete(3T)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>pthread_kill(3T)</code>	MT-Safe, Async-Signal-Safe
<code>pthread_mutexattr_destroy(3T)</code>	MT-Safe
<code>pthread_mutexattr_getprioceiling(3T)</code>	MT-Safe
<code>pthread_mutexattr_getprotocol(3T)</code>	MT-Safe
<code>pthread_mutexattr_getpshared(3T)</code>	MT-Safe
<code>pthread_mutexattr_init(3T)</code>	MT-Safe
<code>pthread_mutexattr_setprioceiling(3T)</code>	MT-Safe
<code>pthread_mutexattr_setprotocol(3T)</code>	MT-Safe
<code>pthread_mutexattr_setpshared(3T)</code>	MT-Safe
<code>pthread_mutex_destroy(3T)</code>	MT-Safe
<code>pthread_mutex_getprioceiling(3T)</code>	MT-Safe
<code>pthread_mutex_init(3T)</code>	MT-Safe
<code>pthread_mutex_lock(3T)</code>	MT-Safe
<code>pthread_mutex_setprioceiling(3T)</code>	MT-Safe
<code>pthread_mutex_trylock(3T)</code>	MT-Safe
<code>pthread_mutex_unlock(3T)</code>	MT-Safe
<code>pthread_once(3T)</code>	MT-Safe
<code>pthread_self(3T)</code>	MT-Safe
<code>pthread_setcancelstate(3T)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>pthread_setcanceltype(3T)</code>	MT-Safe
<code>pthread_setschedparam(3T)</code>	MT-Safe
<code>pthread_setspecific(3T)</code>	MT-Safe
<code>pthread_sigmask(3T)</code>	MT-Safe, Async-Signal-Safe
<code>pthread_testcancel(3T)</code>	MT-Safe
<code>ptsname(3C)</code>	Safe
<code>publickey(3N)</code>	Safe
<code>putc(3S)</code>	MT-Safe
<code>putchar(3S)</code>	MT-Safe
<code>putenv(3C)</code>	Safe
<code>putmntent(3C)</code>	Safe
<code>putp(3X)</code>	Unsafe
<code>putpwent(3C)</code>	Unsafe
<code>puts(3S)</code>	MT-Safe
<code>putspent(3C)</code>	Unsafe
<code>pututline(3C)</code>	Unsafe
<code>pututxline(3C)</code>	Unsafe
<code>putw(3S)</code>	MT-Safe
<code>putwc(3I)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>putwchar(3I)</code>	MT-Safe
<code>putwin(3X)</code>	Unsafe
<code>putws(3I)</code>	MT-Safe
<code>qeconvert(3)</code>	MT-Safe
<code>qfconvert(3)</code>	MT-Safe
<code>qgconvert(3)</code>	MT-Safe
<code>qiflush(3X)</code>	Unsafe
<code>qsort(3C)</code>	Safe
<code>quadruple_to_decimal(3)</code>	MT-Safe
<code>rac_drop(3N)</code>	Unsafe
<code>rac_poll(3N)</code>	Unsafe
<code>rac_recv(3N)</code>	Unsafe
<code>rac_send(3N)</code>	Unsafe
<code>raise(3C)</code>	MT-Safe
<code>rand(3C)</code>	Unsafe, use <code>rand_r()</code>
<code>random(3C)</code>	Unsafe
<code>raw(3X)</code>	Unsafe
<code>rcmd(3N)</code>	Unsafe
<code>readdir(3C)</code>	Unsafe, use <code>readdir_r()</code>

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>read_vtoc(3X)</code>	Unsafe
<code>realloc(3C)</code>	Safe
<code>realloc(3X)</code>	Safe
<code>realpath(3C)</code>	MT-Safe
<code>recv(3N)</code>	Safe
<code>recvfrom(3N)</code>	Safe
<code>recvmsg(3N)</code>	Safe
<code>redrawwin(3X)</code>	Unsafe
<code>refresh(3X)</code>	Unsafe
<code>regcmp(3G)</code>	MT-Safe
<code>regcomp(3C)</code>	MT-Safe
<code>regerror(3C)</code>	MT-Safe
<code>regex(3G)</code>	MT-Safe
<code>regexexec(3C)</code>	MT-Safe
<code>regexpr(3G)</code>	MT-Safe
<code>regfree(3C)</code>	MT-Safe
<code>registerrpc(3N)</code>	Unsafe
<code>remainder(3M)</code>	MT-Safe
<code>remove(3C)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>remque(3C)</code>	Unsafe
<code>replace_panel(3X)</code>	Unsafe
<code>resetty(3X)</code>	Unsafe
<code>reset_prog_mode(3X)</code>	Unsafe
<code>reset_shell_mode(3X)</code>	Unsafe
<code>resolver(3N)</code>	Unsafe
<code>restartterm(3X)</code>	Unsafe
<code>res_init(3N)</code>	Unsafe
<code>res_mkquery(3N)</code>	Unsafe
<code>res_search(3N)</code>	Unsafe
<code>res_send(3N)</code>	Unsafe
<code>rewind(3S)</code>	MT-Safe
<code>rewinddir(3C)</code>	Safe
<code>rexec(3N)</code>	Unsafe
<code>rint(3M)</code>	MT-Safe
<code>ripcoffline(3X)</code>	Unsafe
<code>rmdirp(3G)</code>	MT-Safe
<code>rnusers(3N)</code>	MT-Safe
<code>rpc(3N)</code>	MT-Safe with exceptions

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>rpcbind(3N)</code>	MT-Safe
<code>rpcb_getaddr(3N)</code>	MT-Safe
<code>rpcb_getmaps(3N)</code>	MT-Safe
<code>rpcb_gettime(3N)</code>	MT-Safe
<code>rpcb_rmtcall(3N)</code>	MT-Safe
<code>rpc_broadcast_exp(3N)</code>	MT-Safe
<code>rpc_call(3N)</code>	MT-Safe
<code>rpc_clnt_auth(3N)</code>	MT-Safe
<code>rpc_clnt_calls(3N)</code>	MT-Safe
<code>rpc_clnt_create(3N)</code>	MT-Safe
<code>rpc_control(3N)</code>	MT-Safe
<code>rpc_createerr(3N)</code>	MT-Safe
<code>rpc_rac(3N)</code>	Unsafe
<code>rpc_reg(3N)</code>	MT-Safe
<code>rpc_soc(3N)</code>	Unsafe
<code>rpc_svc_create(3N)</code>	MT-Safe
<code>rpc_svc_err(3N)</code>	MT-Safe
<code>rpc_svc_reg(3N)</code>	MT-Safe
<code>rpc_xdr(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>rresvport(3N)</code>	Unsafe
<code>rstat(3N)</code>	MT-Safe
<code>ruserok(3N)</code>	Unsafe
<code>rusers(3N)</code>	MT-Safe
<code>rwall(3N)</code>	MT-Safe
<code>rwlock(3T)</code>	MT-Safe
<code>rwlock_destroy(3T)</code>	MT-Safe
<code>rwlock_init(3T)</code>	MT-Safe
<code>rw_rdlock(3T)</code>	MT-Safe
<code>rw_tryrdlock(3T)</code>	MT-Safe
<code>rw_trywrlock(3T)</code>	MT-Safe
<code>rw_unlock(3T)</code>	MT-Safe
<code>rw_wrlock(3T)</code>	MT-Safe
<code>savetty(3X)</code>	Unsafe
<code>scalb(3C)</code>	MT-Safe
<code>scalb(3M)</code>	MT-Safe
<code>scalbn(3M)</code>	MT-Safe
<code>scale_form(3X)</code>	Unsafe
<code>scale_menu(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>scanf(3S)</code>	MT-Safe
<code>scanw(3X)</code>	Unsafe
<code>sched_getparam(3R)</code>	MT-Safe
<code>sched_getscheduler(3R)</code>	MT-Safe
<code>sched_get_priority_max(3R)</code>	MT-Safe
<code>sched_get_priority_min(3R)</code>	MT-Safe
<code>sched_rr_get_interval(3R)</code>	MT-Safe
<code>sched_setparam(3R)</code>	MT-Safe
<code>sched_setscheduler(3R)</code>	MT-Safe
<code>sched_yield(3R)</code>	MT-Safe
<code>scr1(3X)</code>	Unsafe
<code>scroll(3X)</code>	Unsafe
<code>scrollok(3X)</code>	Unsafe
<code>scr_dump(3X)</code>	Unsafe
<code>scr_init(3X)</code>	Unsafe
<code>scr_restore(3X)</code>	Unsafe
<code>scr_set(3X)</code>	Unsafe
<code>seconvert(3)</code>	MT-Safe
<code>secure_rpc(3N)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>seed48(3C)</code>	Safe
<code>seekdir(3C)</code>	Safe
<code>select(3C)</code>	MT-Safe
<code>sema_destroy(3T)</code>	MT-Safe
<code>sema_init(3T)</code>	MT-Safe
<code>sema_post(3T)</code>	MT-Safe, Async-Signal-Safe
<code>sema_trywait(3T)</code>	MT-Safe
<code>sema_wait(3T)</code>	MT-Safe
<code>sem_close(3R)</code>	MT-Safe
<code>sem_destroy(3R)</code>	MT-Safe
<code>sem_getvalue(3R)</code>	MT-Safe
<code>sem_init(3R)</code>	MT-Safe
<code>sem_open(3R)</code>	MT-Safe
<code>sem_post(3R)</code>	Async-Signal-Safe
<code>sem_trywait(3R)</code>	MT-Safe
<code>sem_unlink(3R)</code>	MT-Safe
<code>sem_wait(3R)</code>	MT-Safe
<code>send(3N)</code>	Safe
<code>sendmsg(3N)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>sendto(3N)</code>	Safe
<code>setac(3)</code>	Safe
<code>setauclass(3)</code>	MT-Safe
<code>setauevent(3)</code>	MT-Safe
<code>setauuser(3)</code>	MT-Safe
<code>setbuf(3S)</code>	MT-Safe
<code>setcat(3C)</code>	MT-Safe
<code>setjmp(3C)</code>	Unsafe
<code>setkey(3C)</code>	Safe
<code>setlabel(3C)</code>	MT-Safe
<code>setlocale(3C)</code>	Safe with exceptions
<code>setlogmask(3)</code>	Safe
<code>setnetconfig(3N)</code>	MT-Safe
<code>setnetpath(3N)</code>	MT-Safe
<code>setscrreg(3X)</code>	Unsafe
<code>setsockopt(3N)</code>	Safe
<code>setsyx(3X)</code>	Unsafe
<code>setterm(3X)</code>	Unsafe
<code>settimeofday(3C)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>setupterm(3X)</code>	Unsafe
<code>setutent(3C)</code>	Unsafe
<code>setvbuf(3S)</code>	MT-Safe
<code>set_current_field(3X)</code>	Unsafe
<code>set_current_item(3X)</code>	Unsafe
<code>set_curterm(3X)</code>	Unsafe
<code>set_fielddtype_arg(3X)</code>	Unsafe
<code>set_fielddtype_choice(3X)</code>	Unsafe
<code>set_field_back(3X)</code>	Unsafe
<code>set_field_buffer(3X)</code>	Unsafe
<code>set_field_fore(3X)</code>	Unsafe
<code>set_field_init(3X)</code>	Unsafe
<code>set_field_just(3X)</code>	Unsafe
<code>set_field_opts(3X)</code>	Unsafe
<code>set_field_pad(3X)</code>	Unsafe
<code>set_field_status(3X)</code>	Unsafe
<code>set_field_term(3X)</code>	Unsafe
<code>set_field_type(3X)</code>	Unsafe
<code>set_field_userptr(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
set_form_fields(3X)	Unsafe
set_form_init(3X)	Unsafe
set_form_opts(3X)	Unsafe
set_form_page(3X)	Unsafe
set_form_sub(3X)	Unsafe
set_form_term(3X)	Unsafe
set_form_userptr(3X)	Unsafe
set_form_win(3X)	Unsafe
set_item_init(3X)	Unsafe
set_item_opts(3X)	Unsafe
set_item_term(3X)	Unsafe
set_item_userptr(3X)	Unsafe
set_item_value(3X)	Unsafe
set_max_field(3X)	Unsafe
set_menu_back(3X)	Unsafe
set_menu_init(3X)	Unsafe
set_menu_items(3X)	Unsafe
set_menu_mark(3X)	Unsafe
set_menu_opts(3X)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>set_menu_pad(3X)</code>	Unsafe
<code>set_menu_pattern(3X)</code>	Unsafe
<code>set_menu_sub(3X)</code>	Unsafe
<code>set_menu_term(3X)</code>	Unsafe
<code>set_menu_userptr(3X)</code>	Unsafe
<code>set_menu_win(3X)</code>	Unsafe
<code>set_new_page(3X)</code>	Unsafe
<code>set_panel_userptr(3X)</code>	Unsafe
<code>set_term(3X)</code>	Unsafe
<code>set_top_row(3X)</code>	Unsafe
<code>sfconvert(3)</code>	MT-Safe
<code>sgconvert(3)</code>	MT-Safe
<code>shm_open(3R)</code>	MT-Safe
<code>shm_unlink(3R)</code>	MT-Safe
<code>show_panel(3X)</code>	Unsafe
<code>shutdown(3N)</code>	Safe
<code>sigaddset(3C)</code>	MT-Safe, Async-Signal-Safe
<code>sigdelset(3C)</code>	MT-Safe, Async-Signal-Safe
<code>sigemptyset(3C)</code>	MT-Safe, Async-Signal-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>sigfillset(3C)</code>	MT-Safe, Async-Signal-Safe
<code>sigfpe(3)</code>	Safe
<code>sigismember(3C)</code>	MT-Safe, Async-Signal-Safe
<code>siglongjmp(3C)</code>	Unsafe
<code>significand(3M)</code>	MT-Safe
<code>sigqueue(3R)</code>	Async-Signal-Safe
<code>sigsetjmp(3C)</code>	Unsafe
<code>sigsetops(3C)</code>	MT-Safe, Async-Signal-Safe
<code>sigtimedwait(3R)</code>	Async-Signal-Safe
<code>sigwaitinfo(3R)</code>	Async-Signal-Safe
<code>sin(3M)</code>	MT-Safe
<code>single_to_decimal(3)</code>	MT-Safe
<code>sinh(3M)</code>	MT-Safe
<code>sleep(3B)</code>	Async-Signal-Safe
<code>sleep(3C)</code>	Safe
<code>slk_attroff(3X)</code>	Unsafe
<code>slk_attron(3X)</code>	Unsafe
<code>slk_attrset(3X)</code>	Unsafe
<code>slk_clear(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>slk_init(3X)</code>	Unsafe
<code>slk_label(3X)</code>	Unsafe
<code>slk_noutrefresh(3X)</code>	Unsafe
<code>slk_refresh(3X)</code>	Unsafe
<code>slk_restore(3X)</code>	Unsafe
<code>slk_set(3X)</code>	Unsafe
<code>slk_touch(3X)</code>	Unsafe
<code>socket(3N)</code>	Safe
<code>socketpair(3N)</code>	Safe
<code>space(3)</code>	Safe
<code>spray(3N)</code>	Unsafe
<code>sprintf(3S)</code>	MT-Safe
<code>sqrt(3M)</code>	MT-Safe
<code>srand(3C)</code>	Unsafe
<code>srand48(3C)</code>	Safe
<code>srandom(3C)</code>	Unsafe
<code>sscanf(3S)</code>	MT-Safe
<code>ssignal(3C)</code>	Unsafe
<code>standend(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>standout(3X)</code>	Unsafe
<code>start_color(3X)</code>	Unsafe
<code>step(3G)</code>	MT-Safe
<code>str(3G)</code>	MT-Safe
<code>strcadd(3G)</code>	MT-Safe
<code>strcasecmp(3C)</code>	Safe
<code>strcat(3C)</code>	Safe
<code>strccpy(3G)</code>	MT-Safe
<code>strchr(3C)</code>	Safe
<code>strcmp(3C)</code>	Safe
<code>strcoll(3C)</code>	Safe with exceptions
<code>strcpy(3C)</code>	Safe
<code>strcspn(3C)</code>	Safe
<code>strdup(3C)</code>	Safe
<code>streadd(3G)</code>	MT-Safe
<code>strecpy(3G)</code>	MT-Safe
<code>strerror(3C)</code>	Safe
<code>strfind(3G)</code>	MT-Safe
<code>strfmon(3C)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>strftime(3C)</code>	MT-Safe
<code>string(3C)</code>	Safe
<code>string_to_decimal(3)</code>	MT-Safe
<code>strlen(3C)</code>	Safe
<code>strncasecmp(3C)</code>	Safe
<code>strncat(3C)</code>	Safe
<code>strncmp(3C)</code>	Safe
<code>strncpy(3C)</code>	Safe
<code>strpbrk(3C)</code>	Safe
<code>strptime(3C)</code>	MT-Safe
<code>strrchr(3C)</code>	Safe
<code>strrspn(3G)</code>	MT-Safe
<code>strsignal(3C)</code>	Safe
<code>strspn(3C)</code>	Safe
<code>strstr(3C)</code>	Safe
<code>strtod(3C)</code>	MT-Safe
<code>strtok(3C)</code>	Unsafe, use <code>strtok_r()</code>
<code>strtol(3C)</code>	MT-Safe
<code>strtoll(3C)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>strtoul(3C)</code>	MT-Safe
<code>strtoull(3C)</code>	MT-Safe
<code>strtrns(3G)</code>	MT-Safe
<code>strxfrm(3C)</code>	Safe with exceptions
<code>subpad(3X)</code>	Unsafe
<code>subwin(3X)</code>	Unsafe
<code>svcerr_auth(3N)</code>	MT-Safe
<code>svcerr_decode(3N)</code>	MT-Safe
<code>svcerr_noproc(3N)</code>	MT-Safe
<code>svcerr_noprogram(3N)</code>	MT-Safe
<code>svcerr_progvers(3N)</code>	MT-Safe
<code>svcerr_systemerr(3N)</code>	MT-Safe
<code>svcerr_weakauth(3N)</code>	MT-Safe
<code>svcfld_create(3N)</code>	Unsafe
<code>svcrow_create(3N)</code>	Unsafe
<code>svctcp_create(3N)</code>	Unsafe
<code>svcudp_bufcreate(3N)</code>	Unsafe
<code>svcudp_create(3N)</code>	Unsafe
<code>svc_auth_reg(3N)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>svc_control(3N)</code>	MT-Safe
<code>svc_create(3N)</code>	MT-Safe
<code>svc_destroy(3N)</code>	MT-Safe
<code>svc_dg_create(3N)</code>	MT-Safe
<code>svc_fds(3N)</code>	Unsafe
<code>svc_fd_create(3N)</code>	MT-Safe
<code>svc_getcaller(3N)</code>	Unsafe
<code>svc_reg(3N)</code>	MT-Safe
<code>svc_register(3N)</code>	Unsafe
<code>svc_tli_create(3N)</code>	MT-Safe
<code>svc_tp_create(3N)</code>	MT-Safe
<code>svc_unreg(3N)</code>	MT-Safe
<code>svc_unregister(3N)</code>	Unsafe
<code>svc_vc_create(3N)</code>	MT-Safe
<code>swab(3C)</code>	MT-Safe
<code>swapcontext(3C)</code>	MT-Safe
<code>syncok(3X)</code>	Unsafe
<code>sysconf(3C)</code>	MT-Safe, Async-Signal-Safe
<code>syslog(3)</code>	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>system(3S)</code>	Unsafe
<code>taddr2uaddr(3N)</code>	MT-Safe
<code>tan(3M)</code>	MT-Safe
<code>tanh(3M)</code>	MT-Safe
<code>tcdrain(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcflow(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcflush(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcgetattr(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcgetpgrp(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcgetsid(3)</code>	MT-Safe
<code>tcsendbreak(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcsetattr(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcsetpgrp(3)</code>	MT-Safe, Async-Signal-Safe
<code>tcsetpgrp(3C)</code>	MT-Safe
<code>tdelete(3C)</code>	Safe
<code>telldir(3C)</code>	Safe
<code>tempnam(3S)</code>	Safe
<code>termattrs(3X)</code>	Unsafe
<code>termname(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>textdomain(3I)</code>	Safe with exceptions
<code>tfind(3C)</code>	Safe
<code>tgetent(3X)</code>	Unsafe
<code>tgetflag(3X)</code>	Unsafe
<code>tgetnum(3X)</code>	Unsafe
<code>tgetstr(3X)</code>	Unsafe
<code>tgoto(3X)</code>	Unsafe
<code>threads(3T)</code>	Fork1-Safe, MT-Safe, Async-Signal-Safe
<code>thr_continue(3T)</code>	MT-Safe
<code>thr_create(3T)</code>	MT-Safe
<code>thr_exit(3T)</code>	MT-Safe
<code>thr_getconcurrency(3T)</code>	MT-Safe
<code>thr_getprio(3T)</code>	MT-Safe
<code>thr_getspecific(3T)</code>	MT-Safe
<code>thr_join(3T)</code>	MT-Safe
<code>thr_keycreate(3T)</code>	MT-Safe
<code>thr_kill(3T)</code>	MT-Safe, Async-Signal-Safe
<code>thr_main(3T)</code>	MT-Safe
<code>thr_min_stack(3T)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>thr_self(3T)</code>	MT-Safe
<code>thr_setconcurrency(3T)</code>	MT-Safe
<code>thr_setprio(3T)</code>	MT-Safe
<code>thr_setspecific(3T)</code>	MT-Safe
<code>thr_sigsetmask(3T)</code>	MT-Safe, Async-Signal-Safe
<code>thr_stksegment(3T)</code>	MT-Safe
<code>thr_suspend(3T)</code>	MT-Safe
<code>thr_yield(3T)</code>	MT-Safe
<code>tigetflag(3X)</code>	Unsafe
<code>tigetnum(3X)</code>	Unsafe
<code>tigetstr(3X)</code>	Unsafe
<code>timeout(3X)</code>	Unsafe
<code>timer_create(3R)</code>	MT-Safe with exceptions
<code>timer_delete(3R)</code>	MT-Safe with exceptions
<code>timer_getoverrun(3R)</code>	Async-Signal-Safe
<code>timer_gettime(3R)</code>	Async-Signal-Safe
<code>timer_settime(3R)</code>	Async-Signal-Safe
<code>tmpfile(3S)</code>	Safe
<code>tmpnam(3S)</code>	Unsafe, use <code>tmpnam_r()</code>

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
TNF_DECLARE_RECORD(3X)	MT-Safe
TNF_DEFINE_RECORD(3.3X)	MT-Safe
TNF_DEFINE_RECORD_1(3X)	MT-Safe
TNF_DEFINE_RECORD_2(3X)	MT-Safe
TNF_DEFINE_RECORD_4(3X)	MT-Safe
TNF_DEFINE_RECORD_5(3X)	MT-Safe
TNF_PROBE(3.3X)	MT-Safe
TNF_PROBE(3X)	MT-Safe
TNF_PROBE_0(3X)	MT-Safe
TNF_PROBE_1(3X)	MT-Safe
TNF_PROBE_2(3X)	MT-Safe
TNF_PROBE_4(3X)	MT-Safe
TNF_PROBE_5(3X)	MT-Safe
tnf_process_disable(3X)	MT-Safe
tnf_process_enable(3X)	MT-Safe
tnf_thread_disable(3X)	MT-Safe
tnf_thread_enable(3X)	MT-Safe
toascii(3C)	MT-Safe with exceptions
tolower(3C)	MT-Safe with exceptions

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>top_panel(3X)</code>	Unsafe
<code>top_row(3X)</code>	Unsafe
<code>touchline(3X)</code>	Unsafe
<code>touchwin(3X)</code>	Unsafe
<code>toupper(3C)</code>	MT-Safe with exceptions
<code>towlower(3I)</code>	MT-Safe with exceptions
<code>towupper(3I)</code>	MT-Safe with exceptions
<code>tparm(3X)</code>	Unsafe
<code>tputs(3X)</code>	Unsafe
<code>trig(3M)</code>	MT-Safe
<code>truncate(3C)</code>	MT-Safe
<code>tsearch(3C)</code>	Safe
<code>ttynname(3C)</code>	Unsafe, use <code>ttynname_r()</code>
<code>ttyslot(3C)</code>	Safe
<code>twalk(3C)</code>	Safe
<code>typeahead(3X)</code>	Unsafe
<code>t_accept(3N)</code>	MT-Safe
<code>t_alloc(3N)</code>	MT-Safe
<code>t_bind(3N)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>t_close(3N)</code>	MT-Safe
<code>t_connect(3N)</code>	MT-Safe
<code>t_error(3N)</code>	MT-Safe
<code>t_free(3N)</code>	MT-Safe
<code>t_getinfo(3N)</code>	MT-Safe
<code>t_getstate(3N)</code>	MT-Safe
<code>t_listen(3N)</code>	MT-Safe
<code>t_look(3N)</code>	MT-Safe
<code>t_open(3N)</code>	MT-Safe
<code>t_optmgmt(3N)</code>	MT-Safe
<code>t_rcv(3N)</code>	MT-Safe
<code>t_rcvconnect(3N)</code>	MT-Safe
<code>t_rcvdis(3N)</code>	MT-Safe
<code>t_rcvrel(3N)</code>	MT-Safe
<code>t_rcvudata(3N)</code>	MT-Safe
<code>t_rcvuderr(3N)</code>	MT-Safe
<code>t_snd(3N)</code>	MT-Safe
<code>t_snddis(3N)</code>	MT-Safe
<code>t_sync(3N)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>t_unbind(3N)</code>	MT-Safe
<code>uaddr2taddr(3N)</code>	MT-Safe
<code>ulckpddf(3C)</code>	MT-Safe
<code>ulltostr(3C)</code>	MT-Safe
<code>unctrl(3X)</code>	Unsafe
<code>ungetc(3S)</code>	MT-Safe
<code>ungetch(3X)</code>	Unsafe
<code>ungetwc(3I)</code>	MT-Safe
<code>ungetwch(3X)</code>	Unsafe
<code>unlockpt(3C)</code>	Safe
<code>unordered(3C)</code>	MT-Safe
<code>unpost_form(3X)</code>	Unsafe
<code>unpost_menu(3X)</code>	Unsafe
<code>untouchwin(3X)</code>	Unsafe
<code>update_panels(3X)</code>	Unsafe
<code>updwtmp(3C)</code>	Unsafe
<code>updwtmpx(3C)</code>	Unsafe
<code>user2netname(3N)</code>	MT-Safe
<code>use_env(3X)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
utmpname(3C)	Unsafe
utmpxname(3C)	Unsafe
valloc(3C)	Safe
vfprintf(3S)	Async-Signal-Safe
vidattr(3X)	Unsafe
vidputs(3X)	Unsafe
vlfmt(3C)	MT-Safe
volmgt_check(3X)	MT-Safe
volmgt_inuse(3X)	MT-Safe
volmgt_root(3X)	MT-Safe
volmgt_running(3X)	MT-Safe
volmgt_symdev(3X)	MT-Safe
volmgt_symname(3X)	MT-Safe
vpfmt(3C)	MT-Safe
vprintf(3S)	Async-Signal-Safe
vsprintf(3S)	MT-Safe
vsyslog(3)	Safe
vwprintw(3X)	Unsafe
vwscanw(3X)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
waddch(3X)	Unsafe
waddchnstr(3X)	Unsafe
waddchstr(3X)	Unsafe
waddnstr(3X)	Unsafe
waddnwstr(3X)	Unsafe
waddstr(3X)	Unsafe
waddwch(3X)	Unsafe
waddwchnstr(3X)	Unsafe
waddwchstr(3X)	Unsafe
waddwstr(3X)	Unsafe
wadjcurspos(3X)	Unsafe
watof(3I)	MT-Safe
watoi(3I)	MT-Safe
watol(3I)	MT-Safe
watoll(3I)	MT-Safe
wattroff(3X)	Unsafe
wattron(3X)	Unsafe
wattrset(3X)	Unsafe
wbkgd(3X)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
wbkgdset(3X)	Unsafe
wborder(3X)	Unsafe
wclear(3X)	Unsafe
wclrtoebot(3X)	Unsafe
wclrtoeol(3X)	Unsafe
wconv(3I)	MT-Safe with exceptions
wscat(3I)	MT-Safe
wcschr(3I)	MT-Safe
wscmp(3I)	MT-Safe
wscoll(3I)	MT-Safe
wscopy(3I)	MT-Safe
wscspn(3I)	MT-Safe
wcsetno(3I)	MT-Safe with exceptions
wcslen(3I)	MT-Safe
wcsncat(3I)	MT-Safe
wcsncmp(3I)	MT-Safe
wcsncpy(3I)	MT-Safe
wcspbrk(3I)	MT-Safe
wcsrchr(3I)	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
wcsspn(3I)	MT-Safe
wcstod(3I)	MT-Safe
wcstok(3I)	MT-Safe
wcstol(3I)	MT-Safe
wcstombs(3C)	MT-Safe with exceptions
wcstoul(3I)	MT-Safe
wcstring(3I)	MT-Safe
wcswcs(3I)	MT-Safe
wcswidth(3I)	MT-Safe
wcsxfrm(3I)	MT-Safe
wctomb(3C)	MT-Safe with exceptions
wctype(3I)	MT-Safe
wcursyncup(3X)	Unsafe
wcwidth(3I)	MT-Safe
wdelch(3X)	Unsafe
wdeleteln(3X)	Unsafe
wechochar(3X)	Unsafe
wechowchar(3X)	Unsafe
werase(3X)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
wgetch(3X)	Unsafe
wgetnstr(3X)	Unsafe
wgetnwstr(3X)	Unsafe
wgetstr(3X)	Unsafe
wgetwch(3X)	Unsafe
wgetwstr(3X)	Unsafe
whline(3X)	Unsafe
winch(3X)	Unsafe
winchnstr(3X)	Unsafe
winchstr(3X)	Unsafe
windex(3I)	MT-Safe
winnstr(3X)	Unsafe
winnwstr(3X)	Unsafe
winsch(3X)	Unsafe
winsdelln(3X)	Unsafe
winertln(3X)	Unsafe
winsnstr(3X)	Unsafe
winsnswstr(3X)	Unsafe
winsstr(3X)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>winstr(3X)</code>	Unsafe
<code>winswch(3X)</code>	Unsafe
<code>winswstr(3X)</code>	Unsafe
<code>winwch(3X)</code>	Unsafe
<code>winwchnstr(3X)</code>	Unsafe
<code>winwchstr(3X)</code>	Unsafe
<code>winwstr(3X)</code>	Unsafe
<code>wmove(3X)</code>	Unsafe
<code>wmovenextch(3X)</code>	Unsafe
<code>wmoveprevch(3X)</code>	Unsafe
<code>wprintw(3X)</code>	Unsafe
<code>wredrawln(3X)</code>	Unsafe
<code>wrefresh(3X)</code>	Unsafe
<code>wrindex(3I)</code>	MT-Safe
<code>write_vtoc(3X)</code>	Unsafe
<code>wscanw(3X)</code>	Unsafe
<code>wscasecmp(3I)</code>	MT-Safe
<code>wscat(3I)</code>	MT-Safe
<code>wschr(3I)</code>	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
wscmp(3I)	MT-Safe
wscol(3I)	MT-Safe
wscoll(3I)	MT-Safe
wscpy(3I)	MT-Safe
wscr1(3X)	Unsafe
wscspn(3I)	MT-Safe
wsdup(3I)	MT-Safe
wsetscrreg(3X)	Unsafe
wslen(3I)	MT-Safe
wsncasecmp(3I)	MT-Safe
wsncat(3I)	MT-Safe
wsncmp(3I)	MT-Safe
wsncpy(3I)	MT-Safe
wsprk(3I)	MT-Safe
wsprintf(3I)	MT-Safe
wsrchr(3I)	MT-Safe
wsscanf(3I)	MT-Safe
wssp(3I)	MT-Safe
wstandend(3X)	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
wstandout(3X)	Unsafe
wstod(3I)	MT-Safe
wstok(3I)	MT-Safe
wstol(3I)	MT-Safe
wstring(3I)	MT-Safe
wsxfrm(3I)	MT-Safe
wsyncdown(3X)	Unsafe
wsyncup(3X)	Unsafe
wtimeout(3X)	Unsafe
wtouchln(3X)	Unsafe
wvline(3X)	Unsafe
xdr(3N)	Safe
xdrmem_create(3N)	MT-Safe
xdrrec_create(3N)	MT-Safe
xdrrec_endofrecord(3N)	Safe
xdrrec_eof(3N)	Safe
xdrrec_readbytes(3N)	Safe
xdrrec_skiprecord(3N)	Safe
xdrstdio_create(3N)	MT-Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
xdr_accepted_reply(3N)	Safe
xdr_admin(3N)	Safe
xdr_array(3N)	Safe
xdr_authsys_parms(3N)	Safe
xdr_authunix_parms(3N)	Unsafe
xdr_bool(3N)	Safe
xdr_bytes(3N)	Safe
xdr_callhdr(3N)	Safe
xdr_callmsg(3N)	Safe
xdr_char(3N)	Safe
xdr_complex(3N)	Safe
xdr_control(3N)	Safe
xdr_create(3N)	MT-Safe
xdr_destroy(3N)	MT-Safe
xdr_double(3N)	Safe
xdr_enum(3N)	Safe
xdr_float(3N)	Safe
xdr_free(3N)	Safe
xdr_getpos(3N)	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
xdr_hyper(3N)	Safe
xdr_inline(3N)	Safe
xdr_int(3N)	Safe
xdr_long(3N)	Safe
xdr_longlong_t(3N)	Safe
xdr_opaque(3N)	Safe
xdr_opaque_auth(3N)	Safe
xdr_pointer(3N)	Safe
xdr_quadruple(3N)	Safe
xdr_reference(3N)	Safe
xdr_rejected_reply(3N)	Safe
xdr_replymsg(3N)	Safe
xdr_setpos(3N)	Safe
xdr_short(3N)	Safe
xdr_simple(3N)	Safe
xdr_sizeof(3N)	Safe
xdr_string(3N)	Safe
xdr_union(3N)	Safe
xdr_u_char(3N)	Safe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>xdr_u_hyper(3N)</code>	Safe
<code>xdr_u_int(3N)</code>	Safe
<code>xdr_u_long(3N)</code>	Safe
<code>xdr_u_longlong_t(3N)</code>	Safe
<code>xdr_u_short(3N)</code>	Safe
<code>xdr_vector(3N)</code>	Safe
<code>xdr_void(3N)</code>	Safe
<code>xdr_wrapstring(3N)</code>	Safe
<code>xprt_register(3N)</code>	MT-Safe
<code>xprt_unregister(3N)</code>	MT-Safe
<code>y0(3M)</code>	MT-Safe
<code>y1(3M)</code>	MT-Safe
<code>yn(3M)</code>	MT-Safe
<code>ypclnt(3N)</code>	Unsafe
<code>yperr_string(3N)</code>	Unsafe
<code>ypprot_err(3N)</code>	Unsafe
<code>yp_all(3N)</code>	Unsafe
<code>yp_bind(3N)</code>	Unsafe
<code>yp_first(3N)</code>	Unsafe

TABLE C-1 MT Safety Levels of Library Routines *(continued)*

Library Routine	Safety Level
<code>yp_get_default_domain(3N)</code>	Unsafe
<code>yp_master(3N)</code>	Unsafe
<code>yp_match(3N)</code>	Unsafe
<code>yp_next(3N)</code>	Unsafe
<code>yp_order(3N)</code>	Unsafe
<code>yp_unbind(3N)</code>	Unsafe
<code>yp_update(3N)</code>	Unsafe
<code>_NOTE(3X)</code>	Safe
<code>_tolower(3C)</code>	MT-Safe with exceptions
<code>_toupper(3C)</code>	MT-Safe with exceptions
<code>__nis_map_group(3N)</code>	MT-Safe

Index

Numbers

- 32-bit architectures 76
- 64-bit environment
 - large file support 26
 - large virtual address space 25
 - libraries 26
 - /proc restrictions 25
 - registers 26

A

- Ada 160
- adb 182
- adding
 - to LWP pool 206
 - signals to mask 43
- aiocancel(3) 165, 166
- aioread(3) 165, 166
- aiowait(3) 166
- aiowrite(3) 165, 166
- aio_errno 165
- AIO_INPROGRESS 165
- aio_result_t 165, 166
- algorithms
 - faster with MT 19
 - parallel 240
 - sequential 240
- alternate signal stacks 24, 154
- ANSI C 183
- application-level threads 18
- architecture
 - multiprocessor 237
 - SPARC 76, 237, 239

- assert statement 120, 121, 229
- Async-Signal-Safe
 - signal handlers 160
 - functions 158, 173
- asynchronous
 - event notification 124
 - I/O 164 to 166
 - semaphore use 124
 - signals 154, 158
- atomic, defined 76
- automatic
 - LWP number adjustments 152
 - stack allocation 71

B

- binary semaphores 122
- binding
 - reasons to bind 24, 153, 235, 236
 - threads to LWPs 206
 - values to keys 33, 211
- bottlenecks 231
- bound threads 18, 22, 153, 235, 234
 - alternate signal stacks 154
 - concurrency 236
 - defined 18
 - mixing with unbound threads 234
 - no LWP caching 235
 - priority 151
 - reasons to bind 24, 153
 - scheduling class 151

C

- C++ 183
- cache, defined 237
- caching
 - not for bound thread LWPs 235
 - threads data structure 232
- changing the signal mask 43, 209
- coarse-grained locking 229
- code lock 228, 229
- code monitor 228, 230
- compile flag
 - D_POSIX_C_SOURCE 177
 - D_REENTRANT 177
 - _REENTRANT 177
 - single-threaded application 177
- completion semantics 159
- concurrency 235, 236, 229
 - level 206
 - unbound threads 196
- condition variables 76, 106, 121, 162
- cond_broadcast 218, 219
- cond_destroy 217
- cond_init 216
- cond_init(3T) 222
- cond_signal 218
- cond_timedwait 218
- cond_wait 217
- cond_wait(3T) 163
- contention 231
- continue execution 195
- coroutine linkage 234
- counting semaphores 18, 122
- creating
 - stacks 72, 73, 206, 208
 - thread-specific keys 33 to 35, 211, 212
 - threads 28, 30, 236, 232
- critical section 239
- custom stack 72, 208

D

- daemon threads 207
- data
 - global 33
 - local 33
 - lock 228, 229
 - profile 149
 - races 169

- shared 22, 238
- data, *see* thread-specific data
 - thread specific
- dbx 183
- deadlock 230, 231
- debugging 181, 185
 - adb 182
 - asynchronous signals 181
 - dbx 183
 - deadlocks 181
 - hidden gap in synchronization 181
 - inadequate stack size 182
 - large automatic arrays 182
 - long-jumping without releasing mutex lock 181
 - no synchronization of global memory 181
 - passing pointer to caller's stack 181
 - recursive deadlock 181
 - reevaluate conditions after return from condition wait 181
- deleting signals from mask 43
- destructor function 33, 38
- detached threads 31, 58, 206
- Dijkstra, E. W. 122

E

- EAGAIN 29, 34, 95, 97, 100, 112, 127, 197, 207
- EBUSY 95, 100, 112, 201, 202
- EDEADLK 30, 98
- EFAULT 199 to 203
- EINTR 127, 147, 155, 162, 163
- EINVAL 29, 30, 32, 34, 35, 40 to 42, 44, 48, 49, 57 to 59, 61, 62, 64 to 68, 70, 73, 74, 79 to 81, 85 to 90, 92, 93, 95, 96, 101, 108 to 110, 112, 114, 116 to 118, 124, 126 to 128, 131 to 133, 135 to 137, 139, 197, 199 to 203, 207
- ENOMEM 34, 35, 79, 98, 101, 108, 112, 207
- ENOSPC 124
- ENOSYS 40, 89, 90, 92, 93, 96
- ENOTRECOVERABLE 98, 100
- ENOTSUP 41, 64, 66, 92, 93
- EOWNERDEAD 98, 100
- EPERM 85 to 87, 89 to 91, 99, 125

errno 36, 177, 179, 226
 global variables 226
errno.h 175
__errno 179
error checking 42
ESRCH 30, 32, 42, 43, 47, 195, 196
ETIME 116
event notification 124
examining the signal mask 43, 209
exec(2) 144, 146 to 148
execution resources 196, 197, 235
exit(2) 148, 207
exit(3C) 45

F

finding
 minimum stack size 208
 thread concurrency level 197
 thread priority 213
fine-grained locking 229
flags to thr_create() 206
flockfile(3S) 167
flowchart of compile options 178
fork 218
fork(2) 146, 147
fork1(2) 146, 147
FORTRAN 183
funlockfile(3S) 167

G

getc(3S) 167
getc_unlocked(3S) 167
gethostbyname(3N) 226
gethostbyname_r(3N) 227
getrusage(3B) 150
global
 data 228
 side effects 232
 state 228
 variables 36, 37, 225

H

heap, malloc(3C) storage from 31

I

I/O
 asynchronous 164, 165
 nonsequential 166
 standard 167
 synchronous 164
inheriting priority 205
interrupt 154
interval timer 235
invariants 121, 229

J

joining threads 29, 58, 210

K

kernel context switching 21
keys
 bind value to key 211
 get specific key 35, 212
 global into private 37
 storing value of 35, 212
kill(2) 154, 156

L

-lc 178
ld 178
libc 173
libC 174
libc 175, 178
libdl_stubs 173
libintl 174, 175
libm 174, 175
libmalloc 174, 175
libmapmalloc 174, 176
libnsl 174, 175, 179
libpthread 175, 178
library
 C routines 225
 MT safety 173
 threads 175, 233
libresolv 174
librt 175
libsocket 174, 175
libthread 21, 175, 178, 233
 alternate library 180

- libw 174, 175
- libX11 174
- lightweight processes 22, 150, 153, 233, 234
 - adding an LWP 207
 - creation 234
 - debugging 182
 - defined 18
 - independence 234
 - SunOS 4.0 22
 - system calls 235
 - multiplexing 234
 - not supported 22
 - profile state 149
 - shortage 153
 - special capabilities 234
- limits, resources 150
- limits.h 175
- linking 175
- linking with alternate libthread library 180
 - LD_LIBRARY_PATH 180
 - LD_LIBRARY_PATH_64 180
- linking with libpthread
 - lc 178
 - ld 178
 - lpthread 178
- linking with libthread
 - lc 178
 - ld 178
 - lthread 178
- local variable 227
- lock hierarchy 230
- locking 228
 - coarse grained 229, 231
 - code 228
 - conditional 104
 - data 228
 - fine-grained 229, 231
 - guidelines 231
 - invariants 229
- locks 76
 - mutual exclusion 76, 106, 146, 162
 - read-write 204
 - readers/writer 76
- longjmp(3C) 150, 160
- lposix4 library
 - POSIX 1003.1c semaphore 179
- lpthread 178
- lseek(2) 166

- lthread 178
- LWPs , *see* lightweight processes

M

- main() 232
- malloc(3C) 31
- MAP_NORESERVE 71
- MAP_SHARED 147
- memory
 - consistency 236
 - ordering, relaxed 238
 - strongly ordered 238
- mmap(2) 71, 147
- monitor, code 228, 230
- mprotect(2) 72, 208
- MT-Safe libraries 173
- multiple-readers, single-writer locks 204
- multiplexing with LWPs 234
- multiprocessors 236, 240
- multithreading
 - defined 18
- mutex
 - PTHREAD_MUTEX_ERRORCHECK 97
 - PTHREAD_MUTEX_NORMAL 97
 - PTHREAD_MUTEX_RECURSIVE 97
- mutex scope 80
- mutex, mutual exclusion locks 230
- mutex_destroy 215
- mutex_init 213
- mutex_init(3T) 222
- mutex_lock 215
- mutex_trylock 216
- mutex_trylock(3T) 231
- mutex_unlock 215
- mutual exclusion locks 76, 106, 146, 162
 - attributes 78
 - deadlock 102
 - default attributes 76
 - scope, Solaris and POSIX 77
 - type attribute 81

N

- NDEBUG 120
- netdir 174
- netselect 174

- nice(2) 151, 152
- nondetached threads 31, 44
- nonsequential I/O 166
- null
 - threads 72, 207, 208
- null procedures
 - libpthread stub 178
 - libthread stub 178
- P**
- parallel
 - algorithms 240
 - array computation 234
- Pascal 183
- PC
 - program counter 22
- PC_GETCID 151
- PC_GETCLINFO 151
- PC_GETPARMS 151
- PC_SETPARMS 151
- per-process signal handler 154
- per-thread signal handler 154
- Peterson's Algorithm 239
- PL/1 language 156
- portability 76
- POSIX 1003.4a 19
- pread(2) 165, 166
- printf problem 227
- printf(3S) 160
- prctl(2) 151, 152
- priority 22, 150 to 152, 234
 - finding for a thread 213
 - inheritance 205, 212, 213
 - and scheduling 212
 - range 212
 - setting for a thread 212
- priority inversion 84
- process
 - terminating 45
 - traditional UNIX 17
- producer/consumer problem 140, 222, 237
- profil(2) 149
- profiling an LWP 149
- programmer-allocated stack 72, 208
- prolgen
 - semaphore, P operation 122
- pthread.h 175

- pthread_atfork 44
- pthread_attr_getdetachstate 58
- pthread_attr_getguardsize 60
- pthread_attr_getinheritsched 66
- pthread_attr_getschedparam 68
- pthread_attr_getschedpolicy 65
- pthread_attr_getscope 61
- pthread_attr_getstackaddr 74
- pthread_attr_getstacksize 70
- pthread_attr_init 55
 - attribute values 55
- pthread_attr_setdetachstate 57
- pthread_attr_setguardsize 59
- pthread_attr_setinheritsched 65
- pthread_attr_setschedparam 67
- pthread_attr_setschedpolicy 63
- pthread_attr_setscope 60
- pthread_attr_setstackaddr 72
- pthread_attr_setstacksize 69
- pthread_cancel 47
- pthread_cleanup_pop 50
- pthread_cleanup_push 50
- pthread_condattr_destroy 108
- pthread_condattr_getpshared 110
- pthread_condattr_init 107
- pthread_condattr_setpshared 109
- pthread_cond_broadcast 113, 116, 118
 - example 117
- pthread_cond_broadcast(3T) 155
- pthread_cond_destroy 118
- pthread_cond_init 111
- pthread_cond_signal 113, 114, 118, 120
 - example 115
- pthread_cond_signal(3T) 155
- pthread_cond_timedwait 115
 - example 116
- pthread_cond_timedwait(3T) 162
- pthread_cond_wait 112, 119
 - example 115
- pthread_cond_wait(3T) 155, 162
- pthread_create 28
- pthread_detach 32
- pthread_equal 39
- pthread_exit 44, 45
- pthread_getconcurrency 63
- pthread_getschedparam 41
- pthread_getspecific 35, 37, 38

pthread_join 29, 56
 pthread_join(3T) 71, 164
 pthread_keycreate 33, 38
 pthread_keycreate(3T)
 example 38
 pthread_keydelete 34
 pthread_kill 42
 pthread_kill(3T) 156
 pthread_mutexattr_destroy 79
 pthread_mutexattr_getprioceiling
 get priority ceiling of mutex attribute 88
 pthread_mutexattr_getprotocol
 get protocol of mutex attribute 86
 pthread_mutexattr_getpshared 81
 pthread_mutexattr_getrobust_np
 get robust attribute of mutex 92
 pthread_mutexattr_gettype 83
 pthread_mutexattr_init 78
 pthread_mutexattr_setprioceiling
 set priority ceiling of mutex attribute 87
 pthread_mutexattr_setprotocol
 set protocol of mutex attribute 83
 pthread_mutexattr_setpshared 80
 pthread_mutexattr_setrobust_np
 set robust attribute of mutex 91
 pthread_mutexattr_settype 81
 pthread_mutex_consistent_np 95
 pthread_mutex_destroy 101
 pthread_mutex_getprioceiling
 get priority ceiling of mutex 90
 pthread_mutex_init 94
 pthread_mutex_lock 97
 example 102, 105
 pthread_mutex_lock(3T)
 example 106
 pthread_mutex_setprioceiling
 set priority ceiling of mutex 89
 pthread_mutex_trylock 99, 104
 pthread_mutex_unlock 98
 example 102, 105
 pthread_mutex_unlock(3T)
 example 106
 pthread_once 39
 PTHREAD_PRIO_INHERIT 84
 PTHREAD_PRIO_NONE 83
 PTHREAD_PRIO_PROTECT 84
 pthread_rwlockattr_destroy 131
 pthread_rwlockattr_getpshared 133
 pthread_rwlockattr_init 131
 pthread_rwlockattr_setpshared 132
 pthread_rwlock_destroy 139
 pthread_rwlock_init 134
 pthread_rwlock_rdlock 135
 pthread_rwlock_tryrdlock 136
 pthread_rwlock_trywrlock 137
 pthread_rwlock_unlock 138
 pthread_rwlock_wrlock 136
 PTHREAD_SCOPE_PROCESS 24, 60
 PTHREAD_SCOPE_SYSTEM 24, 60
 pthread_self 38
 pthread_setcancelstate 48
 pthread_setcanceltype 48
 pthread_setconcurrency 62
 pthread_setprio(3T) 151, 153
 pthread_setschedparam 40
 pthread_setspecific 35, 38
 pthread_setspecific(3T)
 example 37
 pthread_sigmask 43
 pthread_sigsetmask(3T) 156
 PTHREAD_STACK_MIN() 72
 pthread_testcancel 49
 pthread_yield 40
 putc(3S) 167
 putc_unlocked(3S) 167
 pwrite(2) 165, 166

R

read(2) 166
 read-write locks 133, 204
 attribute 132
 attributes 130
 readers/writer locks 76
 realtime 235
 scheduling 150, 152
 red zone 71, 72, 208
 reentrant 228
 described 228
 functions 172
 strategies for making 228
 register state 22
 relaxed memory ordering 238
 remote procedure call RPC 20
 replacing signal mask 43

- resume execution 195
- RPC 20, 174, 233
- RT, , *see* realtime
- rwlock_destroy 203
- rwlock_init(3T) 198, 222
- rw_rdlock(3T) 200
- rw_tryrdlock 200
- rw_trywrlock 202
- rw_unlock(3T) 202
- rw_wrlock 201
- _r 227

S

- safety, threads interfaces 169, 174
- SA_RESTART 163
- scheduling
 - class 150, 153
 - compute-bound threads 197
 - priorities 212
 - realtime 150, 152
 - system class 150
 - timeshare 150, 151
- semaphores 76, 122, 141
 - binary 122
 - counting 122
 - counting, defined 18
 - decrement semaphore value 122
 - increment semaphore value 122
 - interprocess 125
 - intraprocess 125
 - named 125
- sema_destroy 221
- sema_init 219
- sema_init(3T) 222
- sema_post 220
- sema_post(3T) 173
- sema_trywait 221
- sema_wait 220
- sem_destroy 128
- sem_init 124
 - example 129
- sem_post 122, 126
 - example 129
- sem_trywait 122, 127
- sem_wait 122, 126
 - example 129
- sending signal to thread 42, 209

- sequential algorithms 240
- setjmp(3C) 150, 159, 160
- shared data 22, 228
- shared-memory multiprocessor 238
- sigaction(2) 154, 155, 163
- sigaltstack(2) 154
- SIGFPE 155, 160
- SIGILL 155
- SIGINT 155, 159, 163
- SIGIO 155, 166
- siglongjmp(3C) 160
- signal(2) 154
- signal(5) 154
- signal.h 42, 43, 175, 209, 210
- signals
 - access mask 43, 209
 - add to mask 43
 - asynchronous 154, 158
 - delete from mask 43
 - handler 154, 158
 - inheritance 205
 - masks 22
 - pending 195, 205
 - replace current mask 43
 - send to thread 42, 209
 - SIGSEGV 71
 - SIG_BLOCK 43
 - SIG_SETMASK 43
 - SIG_UNBLOCK 43
 - stack 154
 - unmasked and caught 162
- sigprocmask(2) 156
- SIGPROF 148
- SIGSEGV 71, 155
- sigsend(2) 154
- sigsetjmp(3C) 160
- sigtimedwait(2) 158
- SIGVTALRM 148
- sigwait(2) 157, 158, 160
- SIGWAITING 153
- SIG_BLOCK 43
- SIG_DFL 154
- SIG_IGN 154
- SIG_SETMASK 43
- SIG_UNBLOCK 43
- single-threaded
 - assumptions 225

- code 76
- defined 18
- processes 148
- size of stack 70, 72, 206, 208
- stack 235, 232
 - address 73, 206
 - boundaries 71
 - creation 73, 206
 - custom 208
 - deallocation 208
 - minimum size 72, 208
 - overflows 71
 - parameters 31
 - pointer 22
 - programmer-allocated 72, 208
 - red zone 71, 72, 208
 - returning a pointer to 171
 - size 70, 72, 206, 208
- stack_base 73, 206
- stack_size 70, 206
- standard I/O 167
- standards 19
- start_routine 206
- static storage 179, 225
- stdio 36, 177
- store buffer 239
- storing thread key value 35, 212
- streaming a tape drive 165
- strongly ordered memory 238
- strtoaddr 174
- suspending a new thread 206
- swap space 71
- synchronization objects 75, 141
 - condition variables 76, 106, 121
 - mutex locks 76, 106
 - read-write locks 204
 - semaphores 76, 122, 140, 219, 223
- synchronous I/O 164, 165
- system calls
 - handling errors 226
 - and LWPs 235
- system scheduling class 150

T

- tape drive, streaming 165
- terminating
 - process 45

- threads 30
- thread synchronization
 - condition variables 25
 - mutex locks 24
 - mutual exclusion locks 76
 - read-write locks 130
 - semaphores 25, 122
- thread-directed signal 158
- thread-private storage 22
- thread-specific data 33, 38
 - global 36 to 38
 - global into private 37
 - new storage class 226
 - private 36
- thread.h 175
- threads
 - compute-bound 197
 - concurrencyconcurrency 196
 - creating 28, 30, 205, 207, 236, 232
 - daemon 207
 - defined 18
 - detached 31, 58, 206
 - exit codes 30
 - exit status 29
 - identifiers 30, 38, 39, 44, 206, 207, 209
 - initial 45
 - joining 29, 45, 210
 - keyskeys 211
 - library 175, 233
 - lightweight processes 22
 - nondetached 31, 44
 - null 72, 207, 208
 - prioritypriority 205
 - private data 33
 - safety 169, 174
 - signalssignals 162
 - stacksstack 171
 - suspended 195
 - suspending 206
 - synchronizing 76, 141
 - terminating 30, 44, 210
 - thread-specific data 226
 - unbound threads 23
 - user-level 18, 21
- THR_BOUND 206, 207
- thr_continue 195
- thr_continue(3T) 206

- thr_create 212
- thr_create(3T) 205, 208
- THR_DAEMON 207
- THR_DETACHED 206
- thr_exit 210
- thr_exit(3T) 207
- thr_getconcurrency 197
- thr_getconcurrency(3T) 197
- thr_getprio 213
- thr_getspecific 212
- thr_join 210
- thr_keycreate 211
- thr_kill 209
- thr_kill(3T) 173
- thr_min_stack(3T) 206, 207
- THR_NEW_LWP 197, 206, 235
- thr_self(3T) 209
- thr_setconcurrency 196
- thr_setconcurrency(3T) 196, 206, 235
- thr_setprio 212
- thr_setspecific 211
- thr_sigsetmask(3T) 173
- THR_SUSPENDED 206
- thr_yield(3T) 209, 231
- time slicing 152
- time-out 116, 218
- timeshare scheduling class 150 to 152
- tiuser.h 179
- TLI 174, 179
- tools
 - adb 182
 - dbx 183
 - debugger 183
- total store order 239
- trap 154
- TS, , *see* timeshare scheduling class
- TSD, *see* thread-specific data
- __t_errno 179

U

- unbound threads 150

- alternate signal stacks 154
- caching 232
- concurrency 196, 236
- defined 18
- disadvantage 234
- mixing with bound threads 234
- priorities 150, 212
- reasons not to bind 235, 233
- and scheduling 150, 152, 153
- and thr_setconcurrency(3T) 196, 235
- and pthread_setprio(3T) 151, 153
- unistd.h 175
- UNIX 17, 19, 21, 155, 164, 166, 226
- user space 21
- user-level threads 18, 21
- USYNC_PROCESS 199, 214, 216, 219, 221, 222, 236
- USYNC_PROCESS_ROBUST 214
- USYNC_THREAD 199, 214, 217, 220, 221

V

- variables
 - condition 76, 106, 121, 141
 - global 225
 - primitive 76
- verhogen
 - semaphore, V operation 122
- vfork(2) 146

W

- write(2) 166

X

- XDR 174