



SunSoft, Inc.
2550 Garcia Avenue
Mountain View, CA 94043

For more information, call 1 800 227-9227.

Printed in USA 9/91 91025-0 1.5K

SunSoft and SunTech Enterprises plan to provide many of these components, as well as work with third party vendors for solutions. Components that are already being developed include the following:

- **Threads libraries.** SunSoft will provide a threads library that will allow C programmers to create their own threads and communicate between threads. These interfaces will be conformant with the POSIX 1003.4 specifications.
- **Multithread safe and multithread hot libraries.** Many of the standard libraries that programmers use are being made multithread safe. Doing so will ensure that program correctness is maintained even if several threads try to execute the library concurrently. Some libraries are even being made multithread hot or being optimized for multithreaded execution.
- **Multithreaded debugger.** Perhaps the most important tool for programmers writing multithreaded applications is a debugger that can aid multithreaded program development. SunSoft plans to provide such a debugger.
- **Threads extensions to other languages.** Various language compilers, including FORTRAN and Ada, will be enhanced to support the explicit use of threads. These will then be extended later on to support automatic parallelization.

SunSoft and SunTech Enterprises recognize multithreading as a powerful programming paradigm. They are currently developing software themselves and working with third party vendors to provide an environment that will benefit a wide variety of software developers.

While Mailtool may not be an application that really needs to be multithreaded, the previous example illustrates the power of multithreading. There are several types of applications that would benefit from multithreading and all of them have one characteristic in common: they all have multiple independent functions, which can be done concurrently. Specific examples include database engines, telecommunication switching software, and real-time control software that monitors asynchronous events.

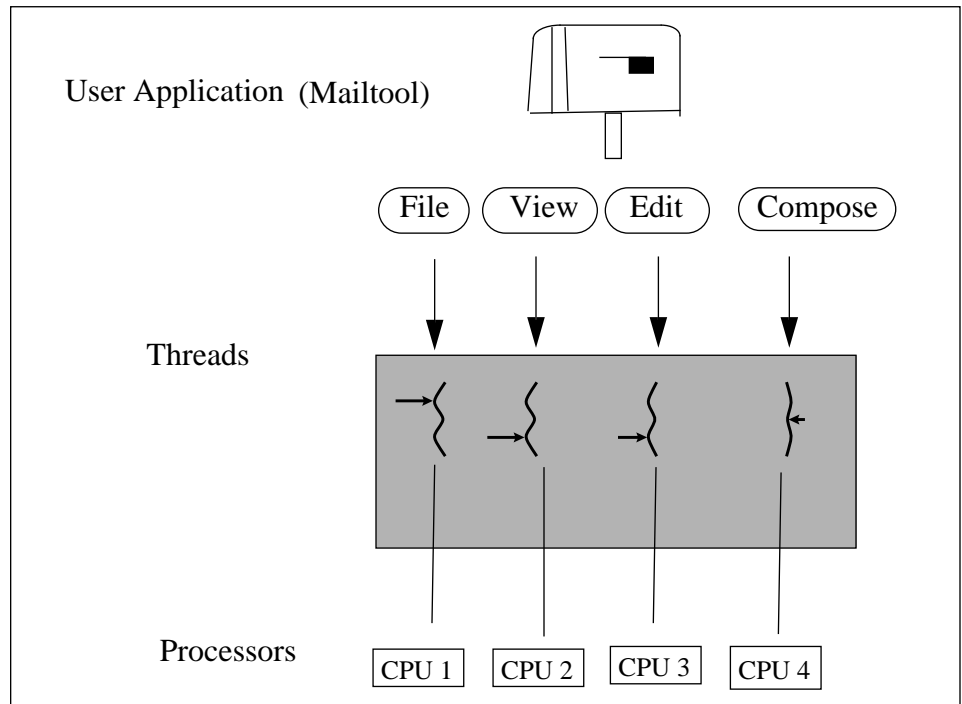


Figure 4 Multithreading a Single Application

Future Multithreading Environment in SunOS

SunSoft understands that multithreading is far more than just a multithreaded operating system. In fact, the operating system is only the beginning. In order to encourage applications programmers to multithread their applications, the software environment should have several components to aid in this process.

Compatibility

Although SunOS 5.0 has significant new technology, it remains fully compatible with UNIX SVR4. It complies with all the major industry standards including POSIX 1003.1, X/OPEN XPG3, and SVID level 3. Furthermore, SunOS 5.0 will preserve the UNIX SVR4 ABI.

Future Kernel Enhancements

SunSoft plans to continue improving the technology in SunOS while maintaining compatibility with standards. The kernel's multithreading performance will be even further improved and new real-time features added. Some of these real-time features include the following:

- Real-time file system providing high deterministic performance with improved reliability.
- Real-time I/O features that allow applications to directly interact with devices rather than use the operating system.
- Improved network performance that will allow real-time applications to make extensive use of networking while maintaining determinism and fast response time.

Multithreading User Applications

SunOS 5.0 is the first step in SunSoft's effort to introduce the multithreaded programming paradigm to its customers. In order to realize the full potential of multithreading, developers need an environment which allows them to build and execute their own multithreaded applications. By doing so they can speed up their individual applications which can execute concurrently.

For example, consider a multithreaded Mailtool where every button in the Mailtool is assigned a thread that performs the function of that particular button. Therefore, one could click on the button to retrieve new mail, and while the mail is being retrieved compose a new message to someone else. Since different threads are handling each of these functions, both functions are done concurrently. Of course, if the underlying hardware platform is a multiprocessor, each of the threads could be executing on different processors. Such a Mailtool will run several times faster than the single-threaded Mailtool we have today. (See Figure 4.)

-
- Fixed priority real-time processes (also in standard UNIX SVR4). SunOS 5.0 allows processes to be classified as real-time processes. These processes have fixed priorities that are not affected by the system's heuristic priority adjustment algorithm.
 - User process priority manipulation (also in standard UNIX SVR4). Users can increase or decrease the priority of processes by issuing a system call *priocntl()*. This gives the user control over system response.
 - High resolution timers (also in standard UNIX SVR4). The timers in SunOS 5.0 can be programmed to provide microsecond granularity notifications. This is particularly useful for monitoring real-time events.
 - Completely preemptive scheduling (SunOS 5.0 value-added). As a result of using locks for mutual exclusion, the SunOS 5.0 kernel has also been made completely preemptive. Since all critical data structures are protected by locks, interrupts can be serviced as soon as they occur and if necessary, scheduling initiated immediately after processing the interrupt.
 - Process priority inheritance (SunOS 5.0 value-added). As a result of preemptive scheduling, it is possible for processes in the kernel to suffer from starvation or indefinite postponement. In such a scenario, a high priority process could need a critical resource held by a low priority process, yet not receive that resource because the low priority process is never scheduled to run. SunOS 5.0 corrects such occurrences by making the low priority process inherit the higher priority of the waiting process until it finishes with the critical resource.
 - Deterministic and guaranteed dispatch latency (SunOS 5.0 value-added). SunOS 5.0 provides deterministic scheduling response and will guarantee various dispatch latencies on different hardware platforms.

Performance

SunOS 5.0 is designed to perform well on both multiprocessor and single processor hardware. In a multiprocessor machine, the multithreaded kernel provides high concurrency within the system, thus making maximum use of each processor. Even environments running a mix of kernel-intensive applications, such as NFS[®] servers, will see throughput increases proportional to the number of processors in the system.

Users will also see an increase in performance on their single-processor systems. The kernel takes advantage of multiple threads to asynchronously perform I/O and execute daemons. The synchronization locks used in the kernel have also been optimized to execute quickly and consume little memory.

In order to provide an efficient multithreaded kernel, free from deadlocks and starvation, many data structures and algorithms have been redesigned. Hundreds of synchronization locks have been added to the kernel to protect and arbitrate access to critical data structures. These locks use the indivisible test-and-set instructions provided by the SPARC® architecture (*swap* and *ldstwb*). Interrupt levels are no longer used to provide mutual exclusion.

The SunOS 5.0 kernel also provides the kernel support for application level multithreading. Therefore, the kernel will not have to undergo any modifications to support future multithreading interfaces (APIs).

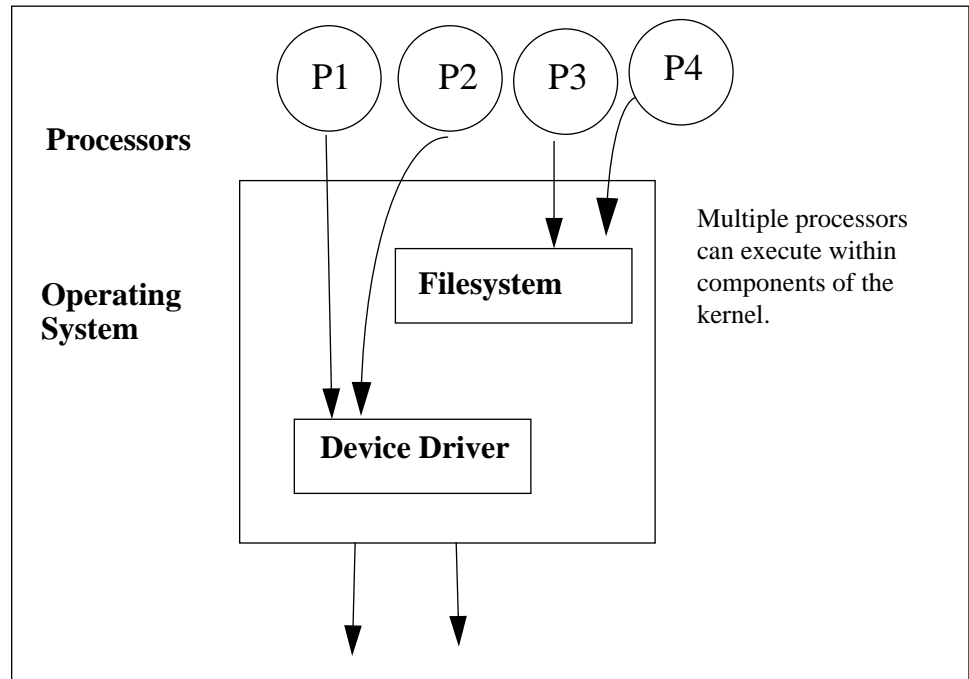


Figure 3 Multithreading in the SunOS 5.0 Kernel

Real-Time Features

SunOS 5.0 contains the real-time extensions of UNIX SVR4 as well as other important enhancements by SunSoft. All the real-time enhancements to SunOS 5.0 have been made to provide deterministic scheduling response. These features include the following:

An accurate way of judging symmetrical multiprocessors is to understand the degree of symmetry. The degree of symmetry is best measured by the amount of concurrent execution that can take place in the operating system kernel. (See Figure 2.)




Figure 2 Degree of Symmetry as a Continuum

SunOS 5.0

SunOS 5.0 is based on UNIX[®] SVR4 from Unix International. In addition to having all the functionality and interfaces provided by SVR4, it also has several value-added features above and below the standard interfaces. Perhaps the most important value-added feature is the multithreaded environment. The SunOS 5.0 kernel is the first step in providing a complete multithreading environment. The kernel provides a foundation on which to build such an environment and SunSoft and SunTech Enterprises are already in the process of developing several future software products that will exploit multithreading.

Multithreaded Kernel Architecture

SunOS 5.0 is a highly symmetric operating system which allows multiple threads to execute concurrently in the operating system. In fact, several processors can even execute concurrently within the same kernel component such as the filesystem, or a device driver. (See Figure 3.)

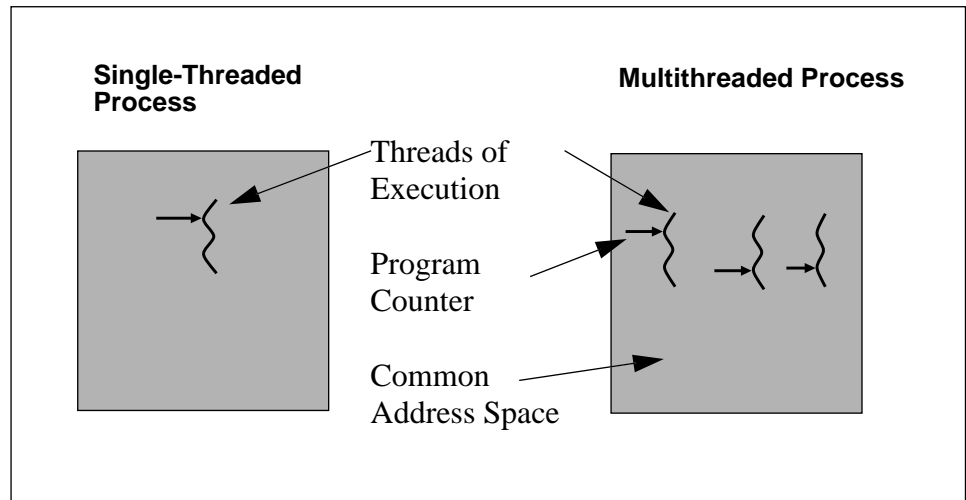


Figure 1 Single and Multithreaded Processes

ASMP (Asymmetric Multiprocessing) and SMP (Symmetric Multiprocessing)

Although these terms are commonly used to refer to two types of multiprocessor systems, they mistakenly imply that symmetric multiprocessing is a binary condition. The degree of symmetry is more accurately portrayed as a continuum and not a binary condition and various systems fall all along this continuum. (See Figure 1.)

Clearly at one end of the spectrum is ASMP, which describes an MP system where several processors exist, but only one of them is designated as master. The master performs all privileged operations such as I/O and the management of operating system resources. In an asymmetric system only one processor can execute the code of the operating system.

The throughput (performance) of asymmetric systems decreases whenever multiple processors try to access operating system resources.

Unlike ASMP, SMP is a little more difficult to define because a variety of systems can be called symmetric. However, all symmetric systems share the property that multiple processors can execute concurrently inside the operating system (thus the term symmetric).

-
- Traditional (single-threaded) applications will run unchanged in this new environment. Therefore, SunOS 5.0 will provide all benefits of UNIX SVR4 — portability across hardware, interoperability across different vendor platforms, and scalability.

Important Terminology

Multiprocessing and multithreading are gaining popularity in the workstation industry. Of course this has meant a proliferation of jargon that is usually better at confusing rather than educating the reader. The following is an explanation of some important terms used in this paper to describe multiprocessing hardware and software.

Tightly Coupled Multiprocessor

A system with multiple processors (CPUs) that share memory and run a single operating system. Most multiprocessor workstations (MP) are of this category.

Thread (Thread of Control)

A thread is a sequence of instructions or flow of executions within a program or process. Traditional processes have only one sequence of instructions that can execute at any given time and are called single-threaded processes.

Multithreading

In contrast to single-threaded processes, multithreaded processes have multiple threads that can execute concurrently. It is important to understand that both single-threaded and multithreaded processes can run on single and multiprocessor machines.

Multithreading and Real-Time

Atri Chatterjee and Jim Herriot

Executive Summary

This paper is targeted at anyone who wants to know about SunSoft's™ multithreading operating system, SunOS™ 5.0. SunOS 5.0 is the foundation of SunSoft's next generation distributed computing solution, Solaris™ 2.0. Solaris 2.0 is comprised of SunOS 5.0, enhanced ONC™, OpenWindows™ V3, DeskSet™ V3, and OPEN LOOK®.

This paper introduces various terms and concepts and puts these in the context of SunSoft products.

Multithreading is seen as an important programming paradigm in the software industry. Not only will it improve the performance of the operating system, but using it will also allow applications to improve their own performance.

Introduction

There are several important characteristics that make the multithreaded SunOS 5.0 particularly attractive:

- A fully symmetric kernel. The SunOS 5.0 kernel is a highly symmetric kernel that maximizes multiprocessor performance.
- The multithreaded kernel also runs efficiently on single processor machines and exploits parallelism wherever possible, such as when performing I/O.
- The multithreading technology has also made the kernel completely preemptive, thus giving it superior real-time performance.
- SunOS 5.0 provides the basis for providing multithreading interfaces to applications programs. The SunOS 5.0 kernel has been built to support the IEEE POSIX 1003.4a threads interfaces when the current draft becomes a standard.

Table of Contents

<i>Executive Summary</i>	<i>1</i>
<i>Introduction</i>	<i>1</i>
<i>Important Terminology</i>	<i>2</i>
<i>SunOS 5.0</i>	<i>4</i>
<i>Multithreading User Applications</i>	<i>7</i>

© 1991 by Sun Microsystems, Inc.—Printed in USA.
2550 Garcia Avenue, Mountain View, California 94043-1100

All rights reserved. No part of this work covered by copyright may be reproduced in any form or by any means—graphic, electronic or mechanical, including photocopying, recording, taping, or storage in an information retrieval system—without prior written permission of the copyright owner.

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

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 (October 1988) and FAR 52.227-19 (June 1987).

The product described in this manual may be protected by one or more U.S. patents, foreign patents, and/or pending applications.

TRADEMARKS

Sun Microsystems, the Sun Logo, NFS, NeWS and SunLink are registered trademarks, and Sun, SunSoft, the SunSoft Logo, Solaris, SunOS, AnswerBook, Catalyst, CDWare, Copilot, DeskSet, Link Manager, Online: DiskSuite, ONC, OpenWindows, SHIELD, SunView, ToolTalk and XView are trademarks of Sun Microsystems, Inc., licensed to SunSoft, Inc., a Sun Microsystems company. SPARC is a registered trademark of SPARC International, Inc. SPARCstation is a trademark of SPARC International, Inc., licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. UNIX and OPEN LOOK are registered trademarks of UNIX System Laboratories, Inc. X Window System is a product of the Massachusetts Institute of Technology.

All other products referred to in this document are identified by the trademarks of the companies who market those products.