


The Design Of The Unix Operating System PDF (Limited Copy)

Maurice J. Bach



More Free Book



Scan to Download

The Design Of The Unix Operating System

Summary

An In-Depth Guide to Unix Operating System Design Principles.

Written by New York Central Park Page Turners Books Club

More Free Book



Scan to Download

About the book

"The Design of the Unix Operating System" by Maurice J. Bach is a seminal work that delves into the architecture and principles underpinning the Unix operating system. This new international paperback edition, while possibly printed in Asia, maintains the same depth of content and insights that made the original a cornerstone in the study of operating systems. Although the cover may feature a disclaimer indicating restrictions on sale in the U.S., it remains entirely legal for use.

In this edition, readers can expect to gain comprehensive knowledge on the structure of Unix, including its processes, file systems, and the intricacies of its system calls. Bach's work emphasizes the design philosophies that contribute to Unix's efficiency and reliability, making it essential for both practitioners and students of computer science.

The economy edition offers a cost-effective opportunity for those looking to explore these fundamental principles without sacrificing the quality of information. Furthermore, every order includes a tracking number, ensuring that readers can monitor their shipment. For any inquiries or support needs, a dedicated customer service team is available to assist, allowing readers to focus entirely on enhancing their understanding of Unix and its enduring impact on modern computing.

More Free Book



Scan to Download

About the author

Maurice J. Bach was a prominent figure in the field of computer science, particularly recognized for his groundbreaking work on operating systems, with a significant focus on Unix. His extensive academic career involved teaching at multiple institutions, where he shared his deep understanding of system software and programming principles. Among his notable contributions is the seminal book "The Design of the Unix Operating System," which provides a thorough exploration of Unix's internal architecture and design principles. This work demystifies complicated concepts, making them accessible to both students and practitioners in the field. Bach's influence extends beyond his writings; his commitment to education has shaped the next generation of computer scientists, leaving a formidable legacy that continues to resonate within the realm of modern operating systems.

More Free Book



Scan to Download



Try Bookey App to read 1000+ summary of world best books

Unlock **1000+** Titles, **80+** Topics
New titles added every week

- Brand
- Leadership & Collaboration
- Time Management
- Relationship & Communication
- Business Strategy
- Creativity
- Public
- Money & Investing
- Know Yourself
- Positive Psychology
- Entrepreneurship
- World History
- Parent-Child Communication
- Self-care
- Mind & Spirituality

Insights of world best books



Free Trial with Bookey

Summary Content List

Chapter 1: General Overview of the System

Chapter 2: Introduction to the Kernel

Chapter 3: The Buffer Cache

Chapter 4: Internal Representation of Files

Chapter 5: System Calls for the File System

Chapter 6: The Structure of Processes

Chapter 7: Process Control

Chapter 8: Process Scheduling and Time

Chapter 9: Memory Management Policies

Chapter 10: The I/O Subsystem

Chapter 11: Interprocess Communication

Chapter 12: Multiprocessor Systems

Chapter 13: Distributed Unix Systems

Chapter 14: Appendix — System Calls

More Free Book



Scan to Download

Chapter 1 Summary: General Overview of the System

1. GENERAL OVERVIEW OF THE SYSTEM

The UNIX operating system, established in 1969, has become a cornerstone of computing, renowned for its versatility across various hardware platforms and a consistent execution environment. This chapter presents an overview of the UNIX system, detailing its history, structural components, and user perspectives.

1.1 HISTORY

UNIX's origins trace back to the challenges faced during the Multics project at Bell Telephone Laboratories, which sought to create a sophisticated operating environment. Pioneers Ken Thompson and Dennis Ritchie spearheaded the development of UNIX on a PDP-7, an early minicomputer. By 1971, UNIX had adapted to the more advanced PDP-11 and was already finding practical applications. The introduction of the C programming language further fueled its evolution, allowing for easier adaptation and distribution, especially in academic institutions and emerging tech companies. Despite AT&T's marketing restrictions imposed by regulatory pressures, UNIX proliferated, with installations reaching approximately 100,000 by 1984 across a diverse array of hardware.

More Free Book



Scan to Download

1.2 SYSTEM STRUCTURE

UNIX's architecture is structured around a core that includes both hardware components and the operating system, known as the kernel. The kernel plays a vital role in abstracting hardware complexities, providing essential services to applications while ensuring program portability. It executes system calls, which are crucial for the interaction between user-level programs and the operating system, signifying its central importance in UNIX's operational framework.

1.3 USER PERSPECTIVE

1.3.1 The File System

Central to user interaction is the UNIX file system, which is hierarchical and offers a uniform method of data organization. This design allows seamless file manipulation and integrates peripherals as file entities, facilitating organized navigation through a structured directory tree. A robust permission system safeguards access to files, promoting security and control.

1.3.2 Processing Environment

The UNIX environment supports multitasking through the concurrent

More Free Book



Scan to Download

execution of multiple processes. While these processes are generally independent, they can communicate and synchronize through system calls, which manage their lifecycle including creation, execution, and termination. The shell serves as the command interpreter, making interactions user-friendly and supporting both synchronous and asynchronous execution of commands.

1.3.3 Building Block Primitives

UNIX encourages modular programming through key primitives such as I/O redirection and pipes. These tools allow for efficient data transfer between processes without relying on temporary files, enabling developers to craft small, specialized programs that can be combined to accomplish complex tasks.

1.4 OPERATING SYSTEM SERVICES

The kernel is responsible for providing a variety of crucial operating system services, including process control, CPU scheduling, memory management, and device handling. Through these services, UNIX maintains a user-friendly interface while carefully obscuring the intricate details of its internal operations.

1.5 ASSUMPTIONS ABOUT HARDWARE

More Free Book



Scan to Download

UNIX functions on two operational modes: user mode and kernel mode. In kernel mode, the kernel responds to system calls from users while enforcing strict control over resource access to maintain security. Additionally, UNIX adeptly manages interrupts and exceptions to ensure system reliability, providing stability during operation.

1.6 SUMMARY

In summary, this chapter provides a foundational understanding of UNIX's structure, highlighting the differences between user and kernel modes and clarifying the kernel's assumptions regarding hardware. UNIX's design philosophy emphasizes small, efficient programs that can be effectively integrated, solidifying its reputation as a powerful and widely adopted operating system. The next chapter will explore the kernel's architecture and its implementation in greater detail, further illuminating the complexities underlying UNIX's functionality.

More Free Book



Scan to Download

Chapter 2 Summary: Introduction to the Kernel

Summary of the Chapters

Introduction to the Kernel

This chapter lays the groundwork for understanding the UNIX operating system, focusing primarily on the kernel's architecture, which serves as its core. The kernel is the vital interface that manages system resources and orchestrates communication between hardware and software.

Architecture of the UNIX Operating System

The UNIX system operates based on an intricate illusion that organizes file locations and process lifecycles. Central to this architecture are files and processes, which are intertwined through major components like the file subsystem and the process control subsystem. A block diagram illustrates this structure.

- **User and Kernel Interaction:** User programs interact with the kernel by making system calls. These requests temporarily transfer control to the

More Free Book



Scan to Download

kernel, which then communicates with hardware devices via device drivers.

- **File Subsystem:** This component is responsible for handling all file operations, including opening, reading, and writing files. It employs a buffering mechanism to streamline data transfer between the kernel and storage devices.

- **Process Control Subsystem:** This subsystem oversees the management of processes, including synchronization, communication, memory management, and scheduling. Critical system calls such as ``fork``, ``exec``, and ``exit`` facilitate the vital processes of creating and terminating functionality.

Introduction to System Concepts

In this section, the focus shifts to kernel data structures and modules that underpin UNIX's operations.

- **File Subsystem Overview:** Each file exists as an inode, which encodes its layout and attributes. File management employs three essential tables: the file table, user file descriptor table, and mode table.

- **Processes:** A process is an instance of a program in execution,

More Free Book



Scan to Download

characterized by its text, data, and stack segments. Managed by the kernel, processes are created with ``fork``, executed via ``exec``, and terminated with ``exit``. They operate in two distinct modes: user mode and kernel mode.

- **Memory Management:** The kernel dynamically allocates and manages memory for processes, allowing for efficient swapping between physical memory and secondary storage.

Process State and Context

Processes in the UNIX kernel undergo transitions through various states—executing (in either user or kernel mode), ready to run, or sleeping—based on system events.

- **Process States:** The system defines a state transition diagram illustrating how processes switch between executing, ready, and sleeping states.

- **Context Switching:** When processes change states, the kernel saves their context information, thus enabling them to resume operations seamlessly.

- **Sleep and Wakeup Mechanisms:** Process management includes

More Free Book



Scan to Download

methods for processes to voluntarily sleep while waiting for specific events, with the kernel awakening them when conditions allow.

Kernel Data Structures

Kernel data structures are typically organized using fixed-size tables, which simplifies the overall kernel design but may limit scalability in handling extensive data.

System Administration

Administrators utilize the same underlying system calls as regular user processes but operate with elevated privileges, enabling them to perform critical functions such as disk management and overall system maintenance.

Summary and Preview

In this chapter, the kernel's architecture is presented alongside its essential components: the file subsystem and process control subsystem. It emphasizes the dynamics of process states, transitions, and the kernel's commitment to maintaining data integrity. Future chapters will further

More Free Book



Scan to Download

explore kernel operations, file systems, process management, and delve into advanced operating system topics.

More Free Book



Scan to Download

Chapter 3 Summary: The Buffer Cache

Summary of Chapter 3: The Buffer Cache

In the UNIX operating system, the kernel plays a crucial role in managing files on mass storage devices, like disks, facilitating processes to read and write data efficiently. To enhance performance, the kernel employs a buffer cache, which serves as a temporary storage pool for frequently accessed disk blocks. This chapter delves into the intricacies of buffer cache management and the algorithms involved.

3.1 Buffer Headers

Each buffer in the cache consists of a memory array that contains the data from the disk and a corresponding header that uniquely identifies each buffer. This one-to-one mapping ensures that each disk block is tied to only one buffer at any time, which prevents incorrect data from being written.

3.2 Structure of the Buffer Pool

The kernel organizes these buffers using a free list and implements a least-recently-used (LRU) algorithm to manage their allocation based on usage frequency. The buffers are further arranged into hash queues, which categorize them by device and block numbers, facilitating quick access.



3.3 Scenarios for Retrieval of a Buffer

The ``getblk`` algorithm demonstrates various circumstances under which the kernel retrieves buffers:

1. A free buffer is readily available on the hash queue.
2. If not found, a free buffer is allocated from the main free list.
3. If a buffer designated for a delayed write is busy, it must complete writing to the disk before being reused.
4. In cases where no free buffers remain, the system enters a sleep state until a buffer becomes available.
5. Similarly, if the requested buffer on the hash queue is in use, a sleep state occurs, awaiting its release.

These scenarios illustrate the intricate handling and states a buffer can occupy during data retrieval.

3.4 Reading and Writing Disk Blocks

The procedures for data reading and writing begin with checking the buffer cache. If a requested block isn't already cached, the kernel communicates with the disk driver to fulfill the read request. Writing can occur in two modes: synchronous, where the write operation is completed before proceeding, or asynchronous, where the operation may be delayed to optimize overall performance.

3.5 Advantages and Disadvantages of the Buffer Cache



The buffer cache offers numerous benefits, including increased system throughput and reduced disk traffic, which contribute to more straightforward and modular system designs. However, it also presents vulnerabilities, particularly during write operations. Unexpected crashes can lead to data corruption, and the extra data copying during I/O operations can occasionally hinder performance.

3.6 Summary

In summary, this chapter highlights the buffer cache's structure and management strategies employed by the kernel to ensure efficient and reliable data access. It illustrates the optimization techniques, such as delayed writes, that are incorporated to maintain high performance levels while managing resources effectively.

3.7 Exercises

The chapter concludes with exercises designed to deepen comprehension of buffer cache mechanisms. These include tasks such as designing alternative algorithms and considering the real-world implications of different approaches to buffer management.

By exploring these facets of the buffer cache, readers gain a deeper understanding of how UNIX optimizes file access, ensuring both efficiency and data integrity.

More Free Book



Scan to Download

Chapter 4: Internal Representation of Files

Chapter 4: Internal Representation of Files

This chapter provides an in-depth exploration of how the UNIX operating system organizes and manages files at a fundamental level. It highlights the unique mode associated with each file, which includes critical information for file access, such as ownership, permissions, and the location of data.

4.1 Inodes

At the core of file representation in UNIX is the inode (index node), which retains essential static information about a file on disk, loaded into memory by the kernel. Each inode contains data regarding the file's ownership, type, access permissions, modification times, and links to the physical addresses of the file's data blocks.

When a process requires access to a file, the kernel utilizes the **iget** algorithm to retrieve the inode from in-core memory, ensuring that concurrent access is managed through proper locking and reference counting. Once a process finishes using the inode, the **iput** algorithm decrements its reference count and may update its status on disk, releasing it back to a pool of available inodes if no processes are using it.



4.2 Structure of a Regular File

Regular files in UNIX employ a block allocation strategy that allows data to be stored in non-contiguous disk locations. This approach prevents fragmentation of space and optimizes storage efficiency. The inode structure supports various addressing methods, including direct, single indirect, double indirect, and triple indirect block addressing, to effectively manage how file data is retrieved.

4.3 Directories

Directories are a special type of file designed to organize the file system hierarchically. They function as a mapping system, linking filenames to their corresponding inode numbers. The access permissions associated with directories dictate the read, write, and execute capabilities for files contained within them.

4.4 Conversion of a Path Name to an Inode

The **namei** algorithm plays a pivotal role in translating human-readable path names into their corresponding inode numbers. It systematically verifies permissions and ensures that each segment of the path corresponds with a valid entry in the directory structure.



4.5 Super Block

The super block is a vital component that stores crucial data about the file system itself, including overall sizes, available free blocks, and information about inodes. It serves as a primary resource management tool for the operating system.

4.6 Inode Assignment to a New File

When creating new files, the **ialloc** algorithm is responsible for assigning inodes from a pre-cached list of free inodes managed in the super block. This process includes safeguards against race conditions to ensure that multiple processes do not inadvertently assign the same inode simultaneously.

4.7 Allocation of Disk Blocks

Disk blocks are allocated via a linked list that exists within the super block. The **alloc** algorithm effectively manages the assignment of these blocks while addressing potential fragmentation issues to maintain efficiency.

4.8 Other File Types

More Free Book



Scan to Download

In addition to regular files and directories, UNIX supports various other file types, such as pipes and special files. Pipes are used for inter-process communication, enabling smooth data transfer between processes, while special files handle device interactions and their specifications.

Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey





Why Bookey is must have App for Book Lovers



30min Content

The deeper and clearer interpretation we provide, the better grasp of each title you have.



Text and Audio format

Absorb knowledge even in fragmented time.



Quiz

Check whether you have mastered what you just learned.



And more

Multiple Voices & fonts, Mind Map, Quotes, IdeaClips...

Free Trial with Bookey



Chapter 5 Summary: System Calls for the File System

Chapter 5: System Calls for the File System

This chapter delves into the system calls that interact with the UNIX file system, expanding on previously discussed internal data structures and algorithms. It serves as a comprehensive guide to the various operations users and applications can perform on files within the operating system.

5.1 File Access System Calls

The section introduces essential system calls—**open**, **read**, **write**, **lseek**, and **close**—which facilitate access to files. The file opening process includes checking permissions and allocating a file descriptor, a unique identifier for the file in use.

5.2 Open System Call

The **open** system call is examined in detail, highlighting its parameters: **path name** (the path to the file), **flags** (indicating the mode of access, such as read or write), and **permissions** (specifying who can access the file). This

More Free Book



Scan to Download

process includes locating the file based on its name and verifying that the necessary permissions are granted for proper access.

5.3 Read System Call

This section focuses on the **read** call, which allows data retrieval from a file. The kernel validates the given file descriptor and tracks offsets for efficient data access. Importantly, if the end of the file is reached before the requested number of bytes is read, the operation will conclude successfully, albeit with fewer bytes than requested.

5.4 Write System Call

Similar to the read operation, the **write** call is discussed, addressing how it functions to input data into a file. This call also considers cases where additional blocks may need to be allocated to accommodate the new data being written.

5.5 File and Record Locking

To manage concurrent access by multiple processes while maintaining data

More Free Book



Scan to Download

integrity, this section introduces file locking mechanisms in UNIX, key for ensuring that processes do not interfere with each other when accessing the same files.

5.6 Adjusting File I/O Position (`lseek`)

The `lseek` function allows users to manipulate the current position within a file, enabling random access to specific data points.

5.7 File Creation

The `creat` system call is summarized here, detailing its role in creating new files or truncating existing ones. It ensures that the necessary permissions and directory structure handling are observed.

5.8 Special File Creation (`mknod`)

Next, the `mknod` call is highlighted, which is used to create special files, including device files which enable communication with hardware or named pipes for inter-process communication.

More Free Book



Scan to Download

5.9 Changing Directory and Root (chdir, chroot)

This section addresses directory manipulation through the **chdir** call, which changes the current working directory, and **chroot**, which alters the root directory to provide an isolated environment for processes.

5.10 Owner and Permissions Modification (chown, chmod)

The chapter continues by exploring how to change a file's ownership and permissions using the **chown** and **chmod** system calls, which are crucial for maintaining security and accessibility within the file system.

5.11 File Status Queries (stat, fstat)

The **stat** and **fstat** calls allow users to obtain metadata about files, including their size and permissions, helping users better understand the attributes of the files they are managing.

5.12 Pipes



In this section, both unnamed and named pipes are examined, detailing their role in facilitating inter-process communication and how they are created and utilized within the operating system.

5.13 File Descriptor Duplication (**dup**)

The **dup** call is discussed here, explaining how it enables the duplication of file descriptors. This functionality allows multiple file descriptors to point to the same underlying file, effectively sharing access.

5.14 Mounting and Unmounting File Systems (**mount**, **umount**)

The chapter elaborates on the processes involved in mounting file systems onto the existing file hierarchy and the proper procedures to unmount them, ensuring the file system's integrity during these operations.

5.15 Linking Files (**link**)

This section introduces the **link** system call, which allows multiple directory entries (names) to reference the same physical file on disk, accompanied by the kernel operations that support this functionality.



5.16 Removing Directory Entries (**unlink**)

The **unlink** call is examined as a crucial operation for removing files from a directory. This includes managing link counts and ensuring that blocks allocated for the file are properly deallocated.

5.17 File System Abstractions

This section discusses higher-level abstractions that allow the UNIX file system to support various file systems, fostering flexibility and extensibility in how files are managed across different storage solutions.

5.18 File System Maintenance

The chapter details the **fsck** program, a utility for checking and repairing inconsistencies in the file system following improper shutdowns, thus ensuring reliability and functionality in data access.

5.19 Summary

More Free Book



Scan to Download

Concluding with a succinct summary, the chapter reiterates the key system calls, their algorithms, and how they interact within the file system, reinforcing the foundational knowledge crucial for managing files effectively.

5.20 Exercises

Finally, to solidify understanding, a series of exercises are provided, encouraging readers to apply the knowledge gained about file system calls and their operations within the UNIX environment.

More Free Book



Scan to Download

Chapter 6 Summary: The Structure of Processes

Chapter 6: The Structure of Processes

In this chapter, we explore the intricate architecture of processes within the UNIX operating system, highlighting their formal structure, state transitions, and management functionalities.

6.1 Process States and Transitions

The life cycle of a UNIX process is marked by various states that represent its current activity or status. These states include:

1. **User Mode Execution:** The process is actively running in user mode where it executes user-level instructions.
2. **Kernel Mode Execution:** The process operates in kernel mode, allowing it to execute privileged instructions and access system resources.
3. **Ready to Run:** The process is prepared to execute but is currently not running due to scheduling by the kernel.
4. **Sleeping in Main Memory:** The process is inactive, waiting for an event or resource.
5. **Ready to Run but Swapped Out:** The process is ready to run but has



been temporarily moved to disk to free up RAM.

6. **Sleeping in Secondary Storage:** Similar to sleeping in main memory, but the process is stored on disk.

7. **Preempted State:** The process was interrupted, typically by another high-priority process.

8. **Newly Created State:** The initial state immediately after the process is created.

9. **Zombie State:** The process has completed execution but still has an entry in the process table to handle its exit status.

Kernel algorithms facilitate transitions between these states, represented in a state transition diagram, allowing for efficient process management.

6.2 Layout of System Memory

The UNIX system organizes memory into logical segments: text, data, and stack. This chapter elaborates on memory management, detailing the concept of virtual memory and the kernel's role, which employs regions and page tables to manage address translation.

***Regions and Pages*:** Logical regions in memory enable resource sharing and protection among processes. The use of page tables allows the system to efficiently map virtual addresses to physical memory, optimizing memory



management.

6.3 The Context of a Process

A process's "context" encompasses its complete address space and the kernel structures associated with it. This context is divided into static components (like the process table entry and user area) and dynamic elements (including the kernel stack and context layers). During operations such as system calls, interrupts, and context switches, different context layers are engaged to maintain continuity.

6.4 Saving the Context of a Process

Whenever a process transitions between states—due to interrupts or system calls—the kernel must save its context. This meticulous saving process is crucial for ensuring that the process can accurately resume its execution. The chapter describes how the kernel efficiently manages context switching and preserves register states for this purpose.

6.5 Manipulation of the Process Address Space

More Free Book



Scan to Download

The chapter details the kernel's handling of the virtual address space through operations like attaching, growing, and loading regions. Additionally, it discusses different algorithms that allocate resources and manage memory in response to various system calls, ensuring smooth operation across processes.

6.6 Sleep

When processes require time to wait for necessary resources or events, they transition into a 'sleep' state. This section outlines the kernel's sleep and wakeup algorithms, which determine the conditions under which processes can be awakened or must remain in a sleeping state.

6.7 Summary

This chapter underscores the significance of process context and memory management within UNIX operating systems. It provides a detailed examination of state transitions, memory organization, and the corresponding algorithms that work together to facilitate efficient process execution and resource management.

6.8 Exercises

More Free Book



Scan to Download

Concluding the chapter, a variety of exercises are provided to reinforce understanding of process states, memory management algorithms, context switching, and the execution of system calls in the UNIX operating system. These activities serve to enhance comprehension and practical application of the concepts presented.

More Free Book



Scan to Download

Chapter 7 Summary: Process Control

Chapter 7: Process Control

Overview

This chapter focuses on the system calls that manage processes in the UNIX Operating System, outlining essential concepts such as process creation, termination, synchronization, memory management, and signal handling.

7.1 Process Creation

The process creation mechanism centers around the `fork` system call, which generates a new process known as the child from an existing process called the parent. This newly created child carries a copy of its parent's context but differs in several aspects, including receiving its own unique Process Identifier (PID). The implementation of `fork` involves multiple steps: allocating an entry in the process table, duplicating the parent's context, and preparing the child's execution state for its tasks.

7.2 Signals

Signals serve as notifications for processes concerning asynchronous events.



The ``kill`` system call is instrumental in sending signals to one or multiple processes. These signals can accomplish various tasks, such as terminating a process, handling exceptions, or signifying the completion of tasks. The UNIX kernel plays a critical role in signal management, determining how individual processes respond—whether by ignoring the signal or executing specific predefined functions.

7.3 Process Termination

Termination of processes is executed through the ``exit`` system call, which initiates cleanup tasks like deallocating resources and informing the parent process of the child's termination status. When a process finishes execution, it transitions into a zombie state, preserving its exit status until the parent processes it.

7.4 Awaiting Process Termination

To nab the exit status of terminated child processes, a parent process can use the ``wait`` system call. This functionality aids in process synchronization and addresses the issue of zombie processes, ensuring they are adequately managed and removed once acknowledged by their parent.

7.5 Invoking Other Programs

More Free Book



Scan to Download

The `exec` system call is pivotal for replacing the current execution context of a process with a new program. This section delves into how executable files are structured, the methodical transfer of parameters, and error handling during the program replacement process.

7.6 The User ID of a Process

Every process is assigned both a real user ID and an effective user ID, which govern its permissions for accessing files and managing processes. The `setuid` system call permits an alteration in the effective user ID, allowing processes to execute with escalated permissions, particularly useful for tasks requiring administrative access.

7.7 Changing the Size of a Process

The `brk` system call provides a means to modify a process's data segment size, facilitating both memory expansion and contraction. Discussion points include memory allocation protocols and the automatic extension of the stack when necessary.

7.8 The Shell

The shell acts as an interface between users and the operating system, utilizing various system calls to interpret and execute user commands,

More Free Book



Scan to Download

manage input/output redirection, and oversee asynchronous process functionality.

7.9 System Boot and the Init Process

The system boot process initializes the entire UNIX environment, culminating in the establishment of the first user-level process, ``init``. This foundational process is responsible for managing subsequent process creation in accordance with system configuration files.

7.10 Summary

This chapter encapsulates the fundamental system calls essential for process management in UNIX, emphasizing ``fork``, ``exec``, ``exit``, and signal management dynamics. Furthermore, it illustrates how the shell interfaces with these system calls to facilitate user command execution and process administration.

7.11 Exercises

A collection of exercises is included to deepen the reader's comprehension of process control concepts. These tasks encourage exploration of edge cases, the ramifications of signal handling, and the workings of various process management functionalities, providing hands-on experience in applying



theoretical knowledge.

More Free Book



Scan to Download

Chapter 8: Process Scheduling and Time

Chapter 8 delves into the intricacies of process scheduling and time management within the UNIX operating system, presenting a comprehensive examination of how processes are allocated CPU time and how the system keeps track of time.

8. PROCESS SCHEDULING AND TIME

UNIX utilizes a time-sharing system where the kernel efficiently allocates CPU time to processes through time slices. This preemptive scheduling method is influenced by process priority, allowing active processes to be paused and resumed based on their needs. The system harnesses a hardware clock to generate regular interrupts, known as clock ticks, which help manage time and track process execution through user commands.

8.1 PROCESS SCHEDULING

At the core of UNIX's scheduling is a highly efficient round-robin strategy fused with multilevel feedback. Each process is assigned a set time quantum for CPU usage. If a process exceeds this limit, it is preempted—meaning it is temporarily halted—and placed back in a priority queue. This dynamic ensures that processes waiting to run are prioritized based on their waiting time, and context switches facilitate the restoration of a process's execution



state when it's resumed.

8.1.1 Algorithm

The UNIX scheduler operates by evaluating processes in the run queue, selecting the one with the highest priority currently loaded in memory. If no eligible processes are found, the system remains idle until the next clock interrupt signals an opportunity for scheduling.

8.1.2 Scheduling Parameters

Each process carries a priority field which factors in recent CPU use and respective user and kernel priority classifications. User-level priorities are set when exiting kernel mode, while kernel-level priorities can fluctuate based on specific operational conditions.

8.1.3 Examples of Process Scheduling

In a fair share scenario, processes are represented as A, B, and C, competing for CPU time. Their priorities are recalibrated periodically—such as every second—depending on their CPU activity, ensuring that resource allocation remains equitable.

8.1.4 Controlling Process Priorities

More Free Book



Scan to Download

The `nice` system call empowers users to adjust their process priorities, fostering a fair allocation of CPU resources. This flexibility is vital for multi-user environments.

8.1.5 Fair Share Scheduler

In addition, UNIX incorporates a fair share scheduler which ensures that user groups receive a proportional amount of CPU time, effectively balancing resource distribution among users belonging to the same category.

8.1.6 Real-Time Processing

Real-time processing poses unique challenges, requiring immediate scheduling in response to specific events. Current UNIX systems lack built-in capabilities to guarantee real-time processing, necessitating additional modifications for such operational demands.

8.2 SYSTEM CALLS FOR TIME

UNIX provides a suite of time-related system calls—including `stime`, `time`, `times`, and `alarm`—which facilitate setting and retrieving time values as well as tracking CPU usage.

More Free Book



Scan to Download

8.3 CLOCK

The clock interrupt handler plays a pivotal role in executing scheduling, accounting, and profiling tasks, ensuring that system time is accurately maintained and process priorities are regularly adjusted.

8.3.1 Restarting the Clock

To maintain consistent operation, clock interrupts must be reprimed periodically.

8.3.2 Internal System Timeouts

Kernel operations may trigger processes based on real-time requirements, with specifics held in a callout table.

8.3.3 Profiling

Through kernel profiling, the system discerns between user and kernel execution times, providing insights for performance evaluations.

8.3.4 Accounting and Statistics

The kernel is tasked with maintaining execution time and memory usage



statistics for individual processes and their corresponding child processes, which aids in resource management.

8.3.5 Keeping Time

Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey





Positive feedback

Sara Scholz

...tes after each book summary
...erstanding but also make the
...and engaging. Bookey has
...ling for me.

Fantastic!!!



I'm amazed by the variety of books and languages Bookey supports. It's not just an app, it's a gateway to global knowledge. Plus, earning points for charity is a big plus!

Masood El Toure

Fi



Ab
bo
to
my

José Botín

...ding habit
...o's design
...ual growth

Love it!



Bookey offers me time to go through the important parts of a book. It also gives me enough idea whether or not I should purchase the whole book version or not! It is easy to use!

Wonnie Tappkx

Time saver!



Bookey is my go-to app for summaries are concise, ins curated. It's like having acc right at my fingertips!

Awesome app!



I love audiobooks but don't always have time to listen to the entire book! bookey allows me to get a summary of the highlights of the book I'm interested in!!! What a great concept !!!highly recommended!

Rahul Malviya

Beautiful App



This app is a lifesaver for book lovers with busy schedules. The summaries are spot on, and the mind maps help reinforce wh I've learned. Highly recommend!

Alex Walk

Free Trial with Bookey

Chapter 9 Summary: Memory Management Policies

Chapter 9: Memory Management Policies

In this chapter, the focus is on memory management policies within the UNIX operating system, highlighting the critical relationship between CPU scheduling and memory management. The primary memory is limited in size and cannot hold all active processes concurrently, thus necessitating an efficient management system to determine which processes remain in primary memory and how to handle those not currently in use.

9.1 Swapping

The chapter begins with an exploration of the swapping memory management policy, which involves moving entire processes between the primary memory and a designated swap device, effectively optimizing memory usage. Three key components are discussed:

- **Allocation of Swap Space:** The kernel allocates contiguous blocks of swap space, in contrast to standard file system allocation. This strategy mitigates fragmentation and ensures that the swap area is used efficiently.
- **Swapping Processes Out:** The kernel decides when processes should be



swapped out based on various triggers, such as memory constraints or the growth needs of processes.

- **Swapping Processes In:** The swapper, an essential kernel process, manages transferring processes back into memory, prioritizing tasks based on their state and how long they have been resident.

9.1.1 Allocation of Swap Space

Efficient memory allocation is vital, and the kernel uses a mapping system to keep track of free blocks on the swap device. A specific algorithm is employed to allocate memory quickly and effectively.

9.1.2 Swapping Processes Out

The kernel's mechanisms for swapping out processes are triggered by specific conditions, such as the demand for new child processes or when a process's stack needs to grow, revealing the dynamic nature of memory management.

9.1.3 Swapping Processes In

The swapper functions continuously, assessing which processes need to be swapped in. It prioritizes based on how long processes have been waiting

More Free Book



Scan to Download

and other factors, aiming to maintain system efficiency.

9.2 Demand Paging

Demand paging represents a significant advancement in memory management, allowing processes to execute even when their entire virtual address space is not loaded in memory. This section emphasizes critical features, such as:

- **Dynamic Loading:** Necessary memory pages are loaded on-the-fly, reducing initial memory demands.
- **Page Fault Management:** Essential for recovering from errors when accessing non-resident memory pages.

9.2.1 Data Structures for Demand Paging

The kernel relies on structured data forms, including page tables and disk block descriptors, to manage the state and usage of memory pages efficiently, linking them to physical memory locations.

9.2.2 The Page-Stealer Process

The page stealer serves as an automated handler, tasked with moving less



frequently accessed pages out of memory to reclaim resources. This process is engaged when available memory falls below a defined threshold.

9.2.3 Page Faults

The chapter defines two primary types of page faults: validity faults, which occur when an invalid memory access is attempted, and protection faults, which arise when a process attempts to access restricted memory. The kernel's response to these faults is pivotal for maintaining system stability.

9.3 A Hybrid System with Swapping and Demand Paging

To alleviate thrashing—a situation where excessive paging slows down the system significantly—the UNIX System V kernel employs a hybrid approach that combines traditional swapping and demand paging. When memory resources are low, entire processes may be swapped out rather than individual pages, enhancing overall system performance.

9.4 Summary

The chapter concludes by summarizing UNIX System V's memory management strategies. It emphasizes the dynamic allocation of resources, the effective handling of page faults, and the continuous efforts in system administration aimed at optimizing performance.

More Free Book



Scan to Download

9.5 Exercises

To reinforce the concepts discussed, exercises at the chapter's end present practical scenarios related to memory management. These tasks encourage readers to design algorithms, examine memory locking techniques, and evaluate the impact of different policies on overall system performance.

More Free Book



Scan to Download

Chapter 10 Summary: The I/O Subsystem

Chapter 10: The I/O Subsystem

Overview

The I/O subsystem in UNIX is a critical component that facilitates communication between processes and peripheral devices such as disks, terminals, and printers through device drivers. Each type of peripheral device typically has a unique driver, although a single driver can also manage multiple devices of the same type, enhancing system flexibility.

Driver Interfaces

In UNIX, devices are classified into two categories: block devices (like disks) that transfer data in blocks, and character devices (like terminals) that process data as streams of characters. Access to these devices occurs via special files within the UNIX file system hierarchy. System calls such as ``open``, ``close``, ``read``, ``write``, and ``ioctl`` are utilized to interact with these devices, although not every system call is necessarily implemented by every driver. Software devices, like memory, offer unique interfaces that do not correspond to specific physical devices.

More Free Book



Scan to Download

System Configuration

Device configuration can be handled in several ways: through hard-coded values, dynamic specifications, or automatic detection in the case of self-identifying devices. The kernel maintains crucial data structures that support interaction between the kernel and device drivers, enabling the execution of necessary function calls as dictated by system requests.

Interrupt Handlers

To efficiently manage hardware requests that occur independently from process actions, the kernel employs interrupt handlers associated with device drivers. This mechanism allows the operating system to respond promptly to hardware events, ensuring smooth operation.

Disk Drivers

Disk drivers are responsible for managing file system requests and translating logical addresses into physical locations on the disk. They play a vital role in optimizing data throughput while also managing areas of the disk to simplify file system management. This efficiency is crucial for maintaining reliable performance in data storage.

Terminal Drivers

More Free Book



Scan to Download

Terminal drivers serve as the user interface, facilitating interactions between users and the operating system. They provide two operational modes: raw mode for unprocessed input and canonical mode for structured input handling. Key elements of terminal management include line disciplines and clists (which handle the flow of data), essential for effective data transmission and superior user experience.

Streams

The innovation of Ritchie's Streams framework significantly enhances the modularity of device drivers. This framework allows for the dynamic stacking of processing layers (modules) over device drivers, thereby creating flexible protocols and processing models for I/O operations. This modular design fosters greater adaptability and customization in managing I/O tasks.

Summary

This chapter delves into the intricate mechanics of device drivers within the UNIX operating system, detailing their interactions with the kernel, the management of disk and terminal I/O, and the implementation of modular streams. Each category of device is governed by specific protocols designed to ensure efficient communication and effective handling of peripheral device requests.

More Free Book



Scan to Download

Exercises

The chapter concludes with a series of exercises aimed at deepening the reader's understanding of the discussed components. These exercises encourage practice with system call implementations and exploration of the implications related to device driver management in UNIX, reinforcing the practical knowledge gained throughout the chapter.

More Free Book



Scan to Download

Chapter 11 Summary: Interprocess Communication

Chapter 11: Interprocess Communication

Interprocess communication (IPC) plays a crucial role in enabling processes within an operating system to exchange data and synchronize their operations. While earlier IPC methods such as pipes, named pipes, and signals were discussed, they each possess limitations. Unnamed pipes restrict communication to descendant processes, named pipes generally do not support network communication or multiple connections, and signals only convey presence without detailed information. This chapter delves into more advanced IPC techniques, focusing on process tracing, the System V IPC package, network communications, and BSD sockets.

11.1 Process Tracing

Process tracing is a valuable debugging tool in UNIX systems, allowing one process to monitor and manipulate another. This is primarily achieved through the ``ptrace`` system call, which gives a debugger the power to set breakpoints and oversee the memory of a target process. The tracing process works by creating a child process that activates a trace bit before running the desired program. When the program executes, the kernel triggers a "trap" signal, allowing the debugger to intervene. While effective for debugging,

More Free Book



Scan to Download

process tracing can be inefficient due to the frequent context switches it requires and is solely applicable to child processes.

11.2 System V IPC

The System V IPC package enhances IPC through three core mechanisms: messages, shared memory, and semaphores, all united by a key-based creation and access method.

- **Messages** are processed through system calls like ``msgget``, ``msgsnd``, ``msgrcv``, and ``msgctl``. They allow the transfer of formatted data streams, organized in a linked list that facilitates prioritized access and control through permission settings.
- **Shared Memory** enables processes to directly share portions of their address space, with the associated system calls (``shmget``, ``shmat``, ``shmdt``, ``shmctl``) supporting efficient data handling. Once a shared region is attached, accessing shared data becomes straightforward, resembling standard memory operations.
- **Semaphores** serve as synchronization tools that manage access to shared resources. Utilizing operations to increment or decrement their values (``semop``), semaphores prevent simultaneous process access, thus avoiding deadlocks. They employ atomic operations and undo structures to



safeguard against risks associated with process terminations.

11.3 Network Communication

Network communication within UNIX has historically relied on makeshift methods that lacked a standardized approach, exposing compatibility issues across different network types. The ``ioctl`` system call helped specify control information but failed to offer uniformity, which could lead to challenges in application interoperability. To address these issues, a more organized framework leveraging various protocols through a streamlined streams mechanism has emerged, ensuring more effective network communication.

11.4 Sockets

BSD sockets provide a comprehensive solution for IPC, functioning both locally and across networks using a client-server model facilitated by socket pairs. These sockets can be categorized into different types—stream and datagram—and support diverse protocols. Key operations associated with sockets, such as ``socket``, ``bind``, ``listen``, ``accept``, ``send``, and ``recv``, are designed for seamless data transfer and simplified protocol management, marking a significant advancement in IPC methodologies.

11.5 Summary

More Free Book



Scan to Download

This chapter outlines various IPC methods, analyzing their strengths and weaknesses. While process tracing is limited and System V IPC offers performance advantages, it can lack flexibility. In contrast, the BSD sockets model presents a robust framework for network communications, fostering more consistent and effective interprocess interactions.

11.6 Exercises

The chapter concludes with a series of exercises that challenge readers to engage with IPC concepts, covering aspects like debugging techniques, semaphore operations, the intricacies of messaging implementations, and performance assessments of different IPC approaches. These practical exercises aid in solidifying understanding and highlight the real-world applicability of IPC mechanisms.

More Free Book



Scan to Download

Chapter 12: Multiprocessor Systems

Chapter 12, titled "Multiprocessor Systems," explores the adaptation of the UNIX operating system, originally designed for uniprocessor environments, to function efficiently in multiprocessor systems that utilize two or more CPUs. These systems allow for simultaneous process execution across multiple processors, but they introduce complexities such as data integrity challenges in the kernel due to potential concurrency issues.

12.1 Problem of Multiprocessor Systems

The inherent design of UNIX ensures data integrity by preventing context switches during kernel mode and disabling interrupts during critical code execution. However, in a multiprocessor scenario, the simultaneous execution of kernel code by multiple CPUs can compromise this integrity, leading to data corruption. To illustrate this point, the chapter provides examples of how concurrent operations can disrupt critical queue management. To address these risks, three primary strategies are proposed:

1. **Restricting critical activities to one CPU** to maintain control,
2. **Serializing access to critical code** through locking mechanisms,
3. **Redesigning algorithms** aimed at reducing contention among processes.



12.2 Solution with Master and Slave Processors

A significant solution discussed is Goble's master-slave architecture, where one CPU (the master) is designated to handle kernel-mode operations, while the slave processors manage user-mode processes. The master CPU is responsible for all system calls and interrupts, with slaves communicating back to the master about these calls. The implementation includes a processor ID within the process table to designate execution locations. However, potential conflicts may arise if multiple processors try to schedule the same process. Proposed solutions include assigning processes to specific CPUs and allowing only one processor to control the scheduling function at a time, thus maintaining order.

12.3 Solution with Semaphores

To ensure data integrity in a multiprocessor framework, semaphores are introduced as a mechanism to partition critical regions of the kernel. This ensures that only one processor may access these regions at any given time. Semaphore operations can be summarized as follows:

- **P (prober)**: Decreases the semaphore's count and blocks the process if the value is negative.
- **V (verifier)**: Increases the semaphore's count and activates any processes that were blocked waiting for the semaphore.



The chapter emphasizes the need for atomicity in semaphore operations to avoid race conditions, ensuring that checking and modifying lock statuses occurs as an indivisible operation. Various algorithms illustrate semaphore applications in practical scenarios, such as managing buffer allocation, processing parent waits to prevent orphaned processes, and coordinating access to driver interfaces.

12.4 The Tunis System

As an alternative to traditional semaphore use, the Tunis system is introduced. It maintains UNIX interface compatibility while employing kernel processes to enforce mutual exclusion through a monitor construct. Monitors only permit a single execution instance at a time, fostering a more modular approach compared to semaphores.

12.5 Performance Limitations

While the chapter outlines effective methods for implementing multiprocessor UNIX systems, it also acknowledges performance limitations. Throughput does not scale linearly with the addition of CPUs due to challenges such as memory contention and delays in semaphore usage. As more processors are engaged, performance may decline, highlighting constraints within the multiprocessor architecture.



12.6 Exercises

To cement understanding of the chapter's concepts, a series of practical exercises are provided. These engage readers in developing multiprocessor algorithms with semaphores while also addressing real-world challenges, such as preventing deadlocks and managing drivers during interrupt handling.

In conclusion, Chapter 12 presents a comprehensive overview of the challenges and solutions associated with adapting UNIX for multiprocessor systems, emphasizing the need for careful design to ensure system integrity and performance.

Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey






Read, Share, Empower

Finish Your Reading Challenge, Donate Books to African Children.

The Concept



This book donation activity is rolling out together with Books For Africa. We release this project because we share the same belief as BFA: For many children in Africa, the gift of books truly is a gift of hope.


The Rule



Earn 100 points



Redeem a book



Donate to Africa

Your learning not only brings knowledge but also allows you to earn points for charitable causes! For every 100 points you earn, a book will be donated to Africa.

Free Trial with Bookey

Chapter 13 Summary: Distributed Unix Systems

Chapter 13: Distributed UNIX Systems

In this chapter, the focus is on the concept of distributed UNIX systems, highlighting their differences from tightly coupled multiprocessor systems. Users often want to access files on much larger machines while still maintaining control over their personal computers, and this chapter illustrates how distributed systems fulfill that demand, allowing for seamless file access without the user needing to understand the complexities of network interactions.

13.0 Distributed System Architecture

Distributed systems stand out by enabling individual computers to retain their autonomy while sharing resources effectively. The chapter introduces various architectural frameworks:

- **Satellite Systems:** These configurations involve a central processor distributing tasks to a group of tightly clustered machines, enhancing computational efficiency through collaborative processing.
- **Newcastle Distributed Systems:** These systems enhance file



accessibility by modifying C library functions to recognize remote files. They operate independently of the kernel, thus streamlining file access.

- **Fully Transparent Distributed Systems** In these systems, users can access files across networks as if they were local, with the kernel invisibly handling network calls using standard path names as if they were accessing local resources.

13.1 Satellite Processors

This section elaborates on how satellite processor architectures function, where a central processor delegates specific tasks to satellite processors, thereby improving overall throughput. The communication process between the satellite and central processors is detailed, illustrating how messages are structured and exchanged during system calls.

13.2 The Newcastle Connection

The Newcastle system represents a significant advancement in remote file access by using altered path names to designate remote files, facilitating user access through a modified library. This introduces challenges such as ensuring consistent user ID mapping and maintaining security across different machines.



13.3 Transparent Distributed File Systems

In a transparent distributed architecture, users can interact with remote files seamlessly, unable to distinguish them from local files. The system's kernel operates requests behind the scenes but often requires additional network calls for permissions and file access, which can lead to performance reductions due to the overhead of messaging.

13.4 Transparent Distributed Model Without Stub Processes

This section examines the use of single-use server processes, an alternative to stub processes for managing remote calls. While this model simplifies certain operational mechanics, it also introduces complexities, particularly in handling system calls and signals in a distributed environment.

13.5 Summary

The chapter wraps up by contrasting the three primary distributed file access schemes—satellite processors, the Newcastle connection, and transparent distributed systems. Each approach has unique methodologies for managing remote file access, reflecting a spectrum of technological strategies for enhancing user experience in distributed environments.

13.6 Exercises

More Free Book



Scan to Download

To reinforce and encourage practical understanding, the chapter includes several exercises that invite readers to explore the discussed concepts further. These activities focus on implementing system calls within satellite systems, understanding signal management in distributed contexts, and tackling challenges linked with remote device access.

This chapter serves to underscore the evolution and innovative approaches in managing and accessing data across distributed UNIX systems, marking an essential aspect of contemporary computing environments.

More Free Book



Scan to Download

Chapter 14 Summary: Appendix — System Calls

Appendix: System Calls Overview

This appendix serves as a concise guide to essential UNIX system calls, defining their core functions and operations. Each call is accompanied by its purpose, highlighting their utility in managing process interactions, file operations, and overall system engagement. Understanding these calls is vital for programmers and users alike, providing foundational insights into the UNIX operating system's functionality.

System Call Descriptions

1. Access:

- **Function:** ``access(filename, mode)``

- **Purpose:** Verifies if the current process has the required permissions (read, write, execute) for a specified file.

2. Accounting:

- **Function:** ``acct(filename)``



- **Purpose:** Toggles system accounting based on the presence of a specified file.

3. Alarm:

- **Function:** ``alarm(seconds)``

- **Purpose:** Sets an alarm to signal after a defined number of seconds, returning the remaining time if already set.

4. Break and Segment Management:

- **Function:** ``brk(end_data_seg)``

- **Purpose:** Adjusts the data segment's upper limit for the calling process.

5. Change Directory:

- **Function:** ``chdir(filename)``

- **Purpose:** Updates the current working directory of the process to the specified path.



6. Change File Mode:

- **Function:** ``chmod(filename, mode)``

- **Purpose:** Modifies the permissions of a file according to the defined mode.

7. Change Owner:

- **Function:** ``chown(filename, owner, group)``

- **Purpose:** Alters the owner and user group of a specified file.

8. Change Root:

- **Function:** ``chroot(filename)``

- **Purpose:** Sets the new root directory for the process to a specified path.

9. Close File Descriptor:

More Free Book



Scan to Download

- **Function:** ``close(fildes)``

- **Purpose:** Closes the referenced file descriptor, releasing system resources.

10. Create File:

- **Function:** ``creat(filename, mode)``

- **Purpose:** Either creates a new file or truncates an existing one to zero length.

11. Duplicate File Descriptor:

- **Function:** ``dup(fildes)``

- **Purpose:** Clones the specified file descriptor, providing an identical handle to the file.

12. Execute File:

- **Function:** ``execve(filename, argv, envp)``

- **Purpose:** Replaces the current process image with a new program



defined by the file, passing arguments and environment.

13. **Exit Process:**

- **Function:** ``exit(status)``

- **Purpose:** Terminates the calling process, returning a specified status code.

14. **File Control:**

- **Function:** ``fcntl(fildes, cmd, arg)``

- **Purpose:** Executes multiple controls and operations on open files, such as altering lock status.

15. **Fork Process:**

- **Function:** ``fork()``

- **Purpose:** Spawns a new process, resulting in two processes: parent and child, returning the child's process ID (PID) to the parent.

16. **Get Process ID:**

More Free Book



Scan to Download

- **Function:** ``getpid()``

- **Purpose:** Retrieves the PID for the current process.

17. Get User ID:

- **Function:** ``getuid()``

- **Purpose:** Returns the real user ID of the process invoking the call.

18. Input/Output Control:

- **Function:** ``ioctl(fildes, cmd, arg)``

- **Purpose:** Executes device-specific control operations.

19. Kill Process:

- **Function:** ``kill(pid, sig)``

More Free Book



Scan to Download

- **Purpose:** Sends a specified signal to a process by its ID.

20. **Link Files:**

- **Function:** ``link(filename1, filename2)``

- **Purpose:** Creates an additional name (link) for a file.

21. **Lseek:**

- **Function:** ``lseek(fildes, offset, origin)``

- **Purpose:** Adjusts the read/write pointer for the file as indicated by the parameters.

22. **Make Node (File):**

- **Function:** ``mknod(filename, modes, dev)``

- **Purpose:** Creates special files (like devices) in the filesystem.



23. Mount File System:

- **Function:** ``mount(specialfile, dir, rwflag)``

- **Purpose:** Intertwines a file system with a directory structure at a specified mount point.

24. Message Queue Control:

- **Function:** ``msgctl(id, cmd, buf)``

- **Purpose:** Manages properties related to message queues.

25. Message Queue Creation:

- **Function:** ``msgget(key, flag)``

- **Purpose:** Generates an identifier for new message queues based on key parameters.

26. Send and Receive Messages:

- **Function:** ``msgsnd(id, msgp, size, flag)`` / ``msgrcv(id, msgp, size,`



type, flag)`

- **Purpose:** Handles sending and receiving messages through a specific message queue.

27. Nice Value Adjustment

- **Function:** `nice(increment)`

- **Purpose:** Modifies the scheduling priority of a process.

28. Open File:

- **Function:** `open(filename, flag, mode)`

- **Purpose:** Opens a specified file based on the given flags and access modes (like read/write).

29. Pause Process:

- **Function:** `pause()`

- **Purpose:** Halts execution of the process until a signal interrupts.



30. Pipe Creation:

- **Function:** ``pipe(fildes)``

- **Purpose:** Establishes a pair of file descriptors for communication between processes.

31. Profiling:

- **Function:** ``profil(buf, size, offset, scale)``

- **Purpose:** Requests profiling data useful for analyzing process execution.

32. Process Tracing

- **Function:** ``ptrace(cmd, pid, addr, data)``

- **Purpose:** Facilitates one process controlling or monitoring another during execution.

33. Read File:

More Free Book



Scan to Download

- **Function:** ``read(filides, buf, size)``

- **Purpose:** Retrieves data from an open file descriptor into a buffer.

34. Semaphore Control:

- **Function:** ``semctl(id, num, cmd, arg)``

- **Purpose:** Performs control operations on semaphore sets.

35. Semaphore Creation:

- **Function:** ``semget(key, nsems, flag)``

- **Purpose:** Creates a semaphore set, used for managing shared resources.

36. Semaphore Operations:

- **Function:** ``semop(id, ops, num)``

- **Purpose:** Applies defined operations to semaphores in a provided



set.

37. Set Process Group:

- **Function:** ``setpgrp()``

- **Purpose:** Assigns a process to a group to manage its behavior in relation to other processes.

38. Set User ID:

- **Function:** ``setuid(uid)`` / ``setgid(gid)``

- **Purpose:** Establishes the real and effective user or group IDs.

39. Shared Memory Control:

- **Function:** ``shmctl(id, cmd, buf)``

- **Purpose:** Modifies properties of shared memory segments.

40. Shared Memory Operations:

More Free Book



Scan to Download

- **Function:** ``shmget(key, size, flag)`` / ``shmat(id, addr, flag)`` / ``shmdt(addr)``
- **Purpose:** Facilitates allocation and management of shared memory areas.

41. Signal Handling:

- **Function:** ``signal(sig, function)``
- **Purpose:** Configures how processes respond to signals.

42. File Status:

- **Function:** ``stat(filename, statbuf)`` / ``fstat(fd, statbuf)``
- **Purpose:** Returns details about a file's attributes and status.

43. Synchronize File System:

- **Function:** ``sync()``



- **Purpose:** Ensures all buffered data is persisted to disk.

44. System Time:

- **Function:** ``time(tloc)``
- **Purpose:** Provides the current system time since the epoch.

45. Process Accounting:

- **Function:** ``times(tbuf)``
- **Purpose:** Delivers accounting information for process resource usage.

46. File Size Limits:

- **Function:** ``ulimit(cmd, limit)``
- **Purpose:** Sets or retrieves limitations on file sizes for the calling process.



47. File Mode Creation Mask:

- **Function:** ``umask(mask)``

- **Purpose:** Determines the default permissions for newly created files.

48. Unmount File System:

- **Function:** ``umount(specialfile)``

- **Purpose:** Detaches a mounted file system from the directory hierarchy.

49. System Information:

- **Function:** ``uname(name)``

- **Purpose:** Returns unique attributes about the operating system.

50. Unlink Files:

- **Function:** ``unlink(filename)``



- **Purpose:** Deletes the directory entry for a file, effectively removing it.

51. File System Statistics:

- **Function:** ``ustat(dev, ubuf)``

- **Purpose:** Provides analytical data about a specified filesystem.

52. Set Access and Modification Times:

- **Function:** ``utime(filename, times)``

- **Purpose:** Modifies the access and modification timestamps of a file.

53. Wait for Child Process:

- **Function:** ``wait(waitstat)``

- **Purpose:** Pauses execution until a child process alters its state.



54. Write File:

- **Function:** ``write(fd, buf, count)``

- **Purpose:** Sends data from a buffer to a specified file descriptor.

This summary encapsulates the critical system calls outlined in the appendix of "The Design of the Unix Operating System," providing a functional overview for quick reference and understanding of their applications in UNIX environments.

More Free Book



Scan to Download