



Processes

OPERATING SYSTEMS

Sercan Külcü | Operating Systems | 16.04.2023

Contents

Contents	1
1 Introduction	5
1.1 Definition of a process	5
1.2 The illusion of many cpus	6
1.3 Time sharing and space sharing	7
1.4 Importance of processes in an operating system	8
1.4.1 <i>Resource Allocation and Management</i>	9
1.4.2 <i>Multitasking and Concurrency</i>	9
1.4.3 <i>Security and Protection</i>	9
1.4.4 <i>Fault Isolation and Recovery</i>	10
1.5 Overview of the functions of processes in multi-tasking and concurrency	10
1.6 Process API	12
1.6.1 <i>Process creation</i>	13
1.6.2 <i>Process destroy</i>	14
1.6.3 <i>The fork() system call</i>	15
1.6.4 <i>The wait() system call</i>	17
1.6.5 <i>The exec() system call</i>	19
1.6.6 <i>The kill() system call</i>	22
1.6.7 <i>The pipe() system call</i>	25
1.7 Process hierarchies	27
1.8 Zombie and orphan process	28
2 Process States and Transitions	29
2.1 Process states:	29

2.1.1	<i>New State:</i>	30
2.1.2	<i>Ready State:</i>	30
2.1.3	<i>Running State:</i>	30
2.1.4	<i>Blocked State:</i>	30
2.1.5	<i>Terminated State:</i>	31
2.2	Transitions between process states	31
2.3	Importance of process states in multi-tasking and concurrency 32	
3	CPU virtualization.....	34
3.1	Limited Direct Execution.....	35
3.2	Restricted Operations	36
3.2.1	<i>User mode and kernel mode</i>	37
3.2.2	<i>System calls</i>	38
3.2.3	<i>System calls look like procedure calls</i>	39
3.2.4	<i>The use of trap and return-from-trap instructions</i>	40
4	Process Control Block (PCB)	41
4.1	Definition of a PCB.....	41
4.2	Contents of a PCB.....	42
4.3	Importance of PCB in process management	42
5	Process Scheduling.....	43
5.1	Definition of process scheduling.....	44
5.2	How to develop scheduling policy	45
5.3	Scheduling algorithms:	46
5.3.1	<i>First-Come-First-Serve (FCFS)</i>	47
5.3.2	<i>Shortest Job First (SJF)</i>	48
5.3.3	<i>Priority Scheduling</i>	50

5.3.4	<i>Round Robin (RR)</i>	51
5.3.5	<i>Multilevel Queue Scheduling (MLQS)</i>	53
5.3.6	<i>Multilevel Feedback Queue Scheduling (MLFQS)</i>	54
5.3.7	<i>Lottery Scheduling</i>	55
5.3.8	<i>Fair-Share Scheduling</i>	57
5.3.9	<i>Guaranteed Scheduling</i>	59
5.4	Importance of process scheduling in multi-tasking and concurrency.....	60
6	Interprocess Communication (IPC) and Synchronization.....	61
6.1	Definition of IPC and synchronization	62
6.2	Race conditions	63
6.3	Critical Regions.....	64
6.4	Mutual Exclusion with Busy Waiting.....	65
6.4.1	<i>Disabling interrupts</i>	66
6.4.2	<i>Lock variables</i>	66
6.4.3	<i>Strict alternation</i>	67
6.4.4	<i>Peterson's solution</i>	69
6.4.5	<i>The TSL instruction</i>	70
6.5	Sleep and Wakeup.....	71
6.6	Methods of IPC and synchronization:	72
6.6.1	<i>Semaphores</i>	73
6.6.2	<i>Mutexes</i>	73
6.6.3	<i>Shared Memory</i>	74
6.6.4	<i>Message Passing</i>	74
6.6.5	<i>Monitors</i>	75
6.6.6	<i>Barriers</i>	76

6.7	Importance of IPC and synchronization.....	77
7	Case Study: Process Management in Linux.....	78
7.1	Overview of Linux process management.....	78
7.2	Comparison with process management in other operating systems.....	79
7.2.1	<i>Windows Process Management:</i>	80
7.2.2	<i>macOS Process Management:</i>	80
7.2.3	<i>Linux Process Management:</i>	80
7.3	Impact on Linux Operating System's performance, reliability, and functionality.....	81
8	Conclusion.....	82

Chapter 3: Processes

1 Introduction

In this section, we will discuss the fundamental concepts of processes and their significance in an operating system. We will also provide an overview of the functions of processes in multi-tasking and concurrency.

A process is defined as a program in execution. It is a fundamental concept in an operating system that enables the system to execute multiple tasks concurrently. Processes are essential for the efficient use of system resources and play a crucial role in managing the system's overall performance.

In this section, we will explore the importance of processes in an operating system and provide a brief overview of the functions of processes in multi-tasking and concurrency. So, let's dive in and understand the critical role that processes play in an operating system!

1.1 Definition of a process

In an operating system, a process is defined as an instance of a program in execution. A process is a fundamental concept in operating systems, and it is responsible for executing instructions, allocating and releasing resources, and communicating with other processes. A process has its own memory space, execution context, and system resources that it uses to accomplish its tasks.

A process is created when a program is loaded into memory, and it is terminated when the program completes its execution or when it is terminated by the operating system. While a process is running, it may spawn child processes, communicate with other processes, and perform various operations on system resources.

Processes are essential components of an operating system, and they provide the foundation for the system's functionality. Without processes, an operating system would not be able to execute applications, manage system resources, or provide a platform for multi-tasking and concurrency.

In the following sections, we will explore the importance of processes in an operating system and provide an overview of their functions in multi-tasking and concurrency. We will also discuss the various attributes that define a process and the mechanisms that an operating system uses to manage processes.

1.2 The illusion of many cpus

In a world where there are many programs that need to run simultaneously, it would be ideal if there were enough physical CPUs to go around. Unfortunately, this is often not the case, as there may be only a few physical CPUs available. The good news is that there are ways to provide the illusion of many CPUs to programs, even when there are only a few physical ones available.

One of the primary techniques that the OS uses to provide the illusion of many CPUs is called time-sharing. This technique involves dividing the available CPU time into small chunks, typically on the order of a few milliseconds. The OS then assigns each program a slice of time, allowing it to execute for that slice before interrupting it and giving the CPU to another program. Because each program is given a slice of time, it

appears to the program as though it has its own CPU, even though it is actually sharing the physical CPU with other programs.

Another technique that the OS uses to provide the illusion of many CPUs is called multiprocessing. In this technique, multiple physical CPUs are used to execute programs simultaneously. The OS assigns each program to a specific CPU, allowing it to execute independently of the other programs. Because each program has its own CPU, it appears to the program as though it has its own CPU, even though it is actually sharing the physical CPU with other programs.

A third technique that the OS uses to provide the illusion of many CPUs is called multithreading. In this technique, each program is divided into multiple threads, each of which can be executed independently of the others. The OS then assigns each thread to a specific CPU or time slice, allowing it to execute independently of the other threads. Because each thread is executing independently, it appears to the program as though it has its own CPU, even though it is actually sharing the physical CPU with other threads.

Overall, there are many ways that the OS can provide the illusion of many CPUs to programs, even when there are only a few physical ones available. By using techniques like time-sharing, multiprocessing, and multithreading, the OS is able to make it appear as though each program has its own CPU, even when it is actually sharing the physical CPU with other programs or threads.

1.3 Time sharing and space sharing

Time sharing and space sharing are two fundamental techniques used by operating systems to manage resources and provide the illusion of multiple entities sharing those resources.

Time sharing refers to the allocation of a resource, such as the CPU or a network link, for a short period of time to one entity, and then to

another and so on. By doing this, many entities can share the resource and utilize it efficiently. In the context of the CPU, time sharing is commonly used to provide the illusion of multiple CPUs. In other words, while there may be only one physical CPU, the operating system can create the illusion of multiple virtual CPUs by time sharing the physical CPU between multiple processes.

On the other hand, space sharing involves the division of a resource among multiple entities. For example, disk space is a natural space-shared resource. Once a block of disk space is assigned to a file, it is typically not available for use by another file until the original file is deleted. In contrast to time sharing, space sharing does not involve time slicing or allocation of time slices. Instead, it involves dividing the resource into distinct pieces and allocating those pieces to different entities.

Both time sharing and space sharing are important techniques used by operating systems to manage resources and ensure efficient utilization of those resources. By using these techniques, the operating system can provide the illusion of multiple entities sharing a resource, even if there is only a single physical resource available. These techniques are crucial for ensuring that modern operating systems can effectively manage the large number of resources available on modern computer systems.

1.4 Importance of processes in an operating system

Processes are fundamental components of any operating system. They are essential for the efficient and effective execution of tasks, which makes them critical to the performance and functionality of an operating system. In this chapter, we will discuss the importance of processes in an operating system.

1.4.1 Resource Allocation and Management

Processes play a vital role in resource allocation and management. A process can request resources such as memory, CPU time, and I/O devices. The operating system is responsible for allocating these resources to the requesting process. Each process is allocated a specific amount of resources, which helps to ensure that all processes receive a fair share of resources. Proper resource allocation and management are essential to maintain the stability and reliability of the operating system.

1.4.2 Multitasking and Concurrency

The ability to run multiple processes simultaneously is known as multitasking. In a multitasking environment, processes share the CPU time, and the operating system manages the allocation of CPU time to each process. The operating system uses scheduling algorithms to ensure that all processes get a fair share of CPU time. This feature allows users to run multiple applications and perform multiple tasks simultaneously, enhancing the efficiency and productivity of the system.

Concurrency refers to the ability of a system to execute multiple tasks simultaneously. Processes enable concurrency by allowing multiple applications to run concurrently. This feature enables users to perform multiple tasks simultaneously, making the system more efficient and productive.

1.4.3 Security and Protection

Processes play a crucial role in maintaining the security and protection of the operating system. Each process runs in its address space, which prevents it from accessing the memory of other processes. This feature ensures that one process cannot interfere with the execution of another process. Additionally, the operating system assigns specific privileges and permissions to each process, which helps to ensure that processes only access the resources they are authorized to access.

1.4.4 Fault Isolation and Recovery

Processes provide fault isolation and recovery capabilities. Each process runs independently, which means that if one process fails, it does not affect the execution of other processes. The operating system can terminate a faulty process without affecting the other running processes. This feature helps to maintain the stability and reliability of the system.

In conclusion, processes are critical components of any operating system. They play a crucial role in resource allocation and management, multitasking and concurrency, security and protection, and fault isolation and recovery. The efficient execution of tasks is only possible because of the underlying processes that run in the background. Therefore, it is important to understand the importance of processes in an operating system to maintain the stability, reliability, and efficiency of the system.

1.5 Overview of the functions of processes in multi-tasking and concurrency

In a modern operating system, it is common to have multiple applications running concurrently. This means that the operating system needs to manage the execution of multiple processes and ensure that they can coexist and operate without interfering with each other. This is where multi-tasking and concurrency come into play.

A process is an executing program with its own memory space, set of resources, and state. Each process has a unique identifier and can interact with other processes through inter-process communication mechanisms. In a multi-tasking environment, the operating system can manage the execution of multiple processes simultaneously.

The primary function of processes in multi-tasking environments is to allow multiple applications to execute concurrently. This is achieved by

the operating system allocating time slices to each process, allowing them to execute for a specified period before suspending execution and allowing other processes to execute. This is known as time-sharing, and it enables the operating system to make the most efficient use of system resources.

Another function of processes in multi-tasking environments is to provide isolation and protection between applications. Each process has its own memory space, which means that it cannot access the memory of another process without explicit permission. This provides a level of security and protection between applications.

Concurrency is the ability of an operating system to manage multiple processes that execute simultaneously. This is achieved by the operating system dividing the system resources among the processes, allowing them to execute concurrently. The operating system provides mechanisms to ensure that the processes do not interfere with each other, and this is achieved through synchronization mechanisms such as semaphores and mutexes.

Processes can communicate with each other using inter-process communication mechanisms. These mechanisms allow processes to exchange data and synchronize their actions. This is an essential function of processes in multi-tasking and concurrent environments, as it enables different applications to work together and share resources.

In summary, processes are essential components of modern operating systems. They enable multi-tasking and concurrency, which allows multiple applications to execute concurrently, and they provide isolation and protection between applications. Inter-process communication mechanisms allow processes to communicate and synchronize their actions, enabling different applications to work together and share resources.

1.6 Process API

The Process API is a crucial component of any modern operating system. This interface provides the necessary functionality to create, manage, and terminate processes. Let's take a closer look at some of the essential functions that must be included in this API.

The first function is the creation of new processes. When a user initiates an action, such as typing a command into a shell or double-clicking on an application icon, the OS must create a new process to run the program. Therefore, the operating system must provide a way to create these new processes.

On the other hand, when a process is no longer needed, the operating system must provide a mechanism to destroy it forcefully. Some processes may end on their own, but when they don't, the user may need to terminate them. Therefore, a "Destroy" interface is necessary to stop these runaway processes.

Another important function provided by the Process API is the ability to wait for a process to complete its execution. For example, if a process is running in the background, the user may need to wait for it to finish before executing another command. Therefore, the OS must provide some kind of waiting interface.

The Process API also includes miscellaneous controls that allow the user to suspend a process (stop it from running temporarily) and then resume it (continue it running). This function can be useful when the user wants to free up resources that are being used by a process temporarily.

Finally, the Process API provides interfaces to get status information about a process, such as how long it has been running or what state it is currently in. This information can be useful in troubleshooting issues or monitoring the performance of the system.

In conclusion, the Process API is a vital part of any operating system. It provides a way to create and manage processes efficiently, allowing users to execute multiple programs simultaneously while sharing system resources fairly.

1.6.1 Process creation

When a user requests to run a program or an application, the operating system is responsible for creating a process to execute the program. The process creation process involves several steps, including memory allocation, file loading, and initialization.

The first step in the process creation process is memory allocation. The operating system must allocate memory for the new process to run. The amount of memory allocated depends on the requirements of the program being executed. The operating system must ensure that there is enough memory available to run the new process and that the memory is contiguous.

Once the memory has been allocated, the operating system loads the necessary files into the memory space. This includes the program code, any libraries that the program uses, and any other necessary resources. The operating system must ensure that the files are loaded in the correct order and that any dependencies are resolved.

Next, the operating system initializes the new process. This involves setting up the environment for the new process to run in. This includes setting the process priority, assigning any necessary resources (such as CPU time or I/O resources), and setting up the initial values of the process's registers and variables.

Finally, the operating system creates a process control block (PCB) for the new process. The PCB is a data structure that contains information about the process, such as its process ID, memory usage, CPU usage, and other important data. The operating system uses the PCB to manage the process and ensure that it is running correctly.

Overall, the process creation process is a complex task that requires the operating system to allocate memory, load files, initialize the process, and create a process control block. However, the process creation process is essential for the execution of programs and applications and is a fundamental part of any modern operating system.

1.6.2 Process destroy

Process destruction is the counterpart to process creation and involves terminating an existing process. The destruction of a process is typically initiated by a user, although it can also occur automatically when a program has completed execution.

When a user wishes to terminate a process, they may send a signal to the process asking it to shut down. The process will then handle the signal, typically by performing any cleanup tasks it needs to do before exiting. These cleanup tasks may include freeing up memory or releasing any resources that the process has been using.

If the process does not respond to the signal, or if it is in an unrecoverable state, the operating system may force the process to terminate. This is done by sending a signal that cannot be ignored or caught by the process, causing it to exit immediately.

The process of destroying a process may also involve the release of any resources that were being used by the process. This includes releasing memory and any other resources that were allocated to the process. In addition, the operating system may update any data structures that were being used to keep track of the process, such as process tables or queues.

One challenge in process destruction is ensuring that the process is terminated safely without causing any harm to the system or other processes. For example, if a process is holding a lock on a shared resource, terminating it abruptly could cause other processes to fail or become deadlocked. To avoid these kinds of issues, the operating system may use techniques such as graceful shutdown procedures or

synchronization mechanisms to ensure that all processes are aware of the impending termination and can take appropriate action.

In summary, process destruction is a critical aspect of the operating system's management of processes. It involves terminating an existing process and releasing any resources that were being used by the process, while ensuring that the system remains stable and other processes are not adversely affected.

1.6.3 The fork() system call

The fork() system call is one of the most important and fundamental functions provided by any operating system. It is the starting point for creating new processes, and is often used in combination with other system calls to build more complex programs.

In essence, the fork() system call creates a new process by duplicating the existing process. The new process, which is known as the child process, is an exact copy of the original process, which is known as the parent process. This means that the child process has the same memory, environment variables, and file descriptors as the parent process. However, the two processes have different process IDs (PIDs), and thus can be distinguished from each other.

The fork() system call takes no arguments and returns an integer value. In the parent process, the value returned by fork() is the PID of the newly created child process. In the child process, the value returned is zero. If fork() fails for any reason, a negative value is returned.

One of the most important things to keep in mind when using fork() is that the child process continues executing from the point where the fork() call was made. This means that both the parent and child processes are running simultaneously, and that the child process can begin executing its own code immediately after being created.

It's also important to understand that the child process is a separate entity from the parent process, and that changes made in one process do not affect the other. For example, if the parent process modifies a variable, the child process will not see that modification unless it explicitly reads the value of the variable.

One of the most common uses of the `fork()` system call is to create a child process to run a new program using the `exec()` system call. This allows a program to run another program in a separate process, without affecting the parent process.

Example: Here's a simple example code that demonstrates the `fork()` system call in C:

```
#include <stdio.h>

#include <unistd.h>

#include <sys/types.h>

int main() {
    pid_t pid;
    int x = 0;

    pid = fork();

    if (pid == -1) {
        printf("Fork failed!\n");
    } else if (pid == 0) {
        // Child process
        printf("Child process, x = %d\n", ++x);
    } else {
```

```

        // Parent process
        printf("Parent process, x = %d\n", x);
    }

    return 0;
}

```

In this example, we declare a variable `x` and initialize it to `0`. We then call `fork()`, which creates a new process (the child process) that is an exact copy of the parent process, except for its PID (process ID).

In the child process, we increment `x` and print out its value, while in the parent process, we simply print out the original value of `x`.

When we run this program, we will see that both the parent and child processes print out a different value of `x`:

```
Parent process, x = 0
```

```
Child process, x = 1
```

This is because the `fork()` system call creates a separate copy of the program's memory for the child process, so any modifications to `x` made by the child process do not affect the parent process.

1.6.4 The `wait()` system call

When a parent process creates a child process using the `fork()` system call, it often needs to wait for the child to complete its execution before it can continue with its own execution. The `wait()` system call provides a way for the parent process to suspend its execution and wait for the child process to terminate.

The `wait()` system call takes no arguments and returns the process ID of the terminated child process. If there are no child processes, the `wait()` system call returns immediately with a value of `-1`.

When a child process terminates, it remains in the system as a zombie process until its parent process performs a `wait()` system call to retrieve its exit status. Once the parent process calls `wait()` and retrieves the exit status, the child process is removed from the system.

Example: Here is an example of how the `wait()` system call can be used in a C program:

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/wait.h>

int main() {
    pid_t pid;
    int status;

    pid = fork();

    if (pid < 0) {
        printf("Fork failed.\n");
        exit(1);
    }
    else if (pid == 0) {
        printf("Child process.\n");
        exit(0);
    }
    else {
```

```
    printf("Parent process.\n");
    wait(&status);
    printf("Child process completed.\n");
}

return 0;
}
```

In this example, the parent process creates a child process using the `fork()` system call. The child process prints a message and exits, while the parent process waits for the child process to complete using the `wait()` system call. Once the child process completes, the parent process prints a message indicating that the child process has completed.

It is important to note that the `wait()` system call suspends the execution of the parent process until the child process completes. If the child process does not terminate, the parent process will be blocked indefinitely. To avoid this situation, it is common to use the `waitpid()` system call, which allows the parent process to specify which child process it is waiting for.

1.6.5 The `exec()` system call

The `exec()` system call is used to replace the current process image with a new process image. This new image can be a different program or script, with its own set of instructions, data, and memory space. This system call is often used in conjunction with the `fork()` system call to create a new process, and then replace its image with a new program or script.

The `exec()` system call comes in various forms, with different suffixes indicating different behaviors. For example, `execl()` takes a variable number of arguments, with the first argument being the name of the new program or script to execute. The remaining arguments are the

command-line arguments passed to the new program or script. Other forms of `exec()` allow for passing environment variables or an array of arguments.

One important thing to note is that when the `exec()` system call is used, the new process image completely replaces the current process image. This means that any code or data in the current process image is lost, along with any open file descriptors or other system resources. Therefore, it is important to make sure that any necessary setup or cleanup code is run before and after the `exec()` system call.

Example: Here's an example code snippet demonstrating the use of the `execl()` system call:

```
#include <unistd.h>

#include <stdio.h>

#include <stdlib.h>

int main() {
    printf("Before exec() call\n");

    // Fork a new process
    pid_t pid = fork();

    if (pid == 0) {
        // Child process
        printf("Child process executing...\n");

        // Replace process image with new program
        execl("/bin/ls", "ls", "-l", NULL);
    }
}
```

```

        // If we get here, the exec() call failed
        perror("execl");
        exit(EXIT_FAILURE);
    }
    else if (pid > 0) {
        // Parent process
        printf("Parent process waiting for child...\n");

        // Wait for child process to complete
        wait(NULL);

        printf("Child process completed\n");
    }
    else {
        // Fork failed
        perror("fork");
        exit(EXIT_FAILURE);
    }

    printf("After exec() call\n");
    return 0;
}

```

In this example, the `execl()` system call is used to replace the child process image with the `ls` program, which will list the files in the current directory. The parent process waits for the child process to complete

before continuing. If the `execl()` call fails, the child process will exit with a failure status.

1.6.6 The `kill()` system call

The `kill()` system call is used to send a signal to a process or a group of processes on the system. The signal is a software interrupt that can be used to convey various types of information to the target process or group of processes.

The signal can be used to request that the process terminate, to suspend or resume the process, or to notify the process of some event or condition. When a process receives a signal, it can either handle the signal itself or allow the operating system to handle it.

The `kill()` system call is a relatively simple system call that takes two arguments: the process ID of the target process and the signal to be sent. When the `kill()` system call is executed, it sends the specified signal to the specified process.

If the signal sent is `SIGKILL`, the target process is immediately terminated without any chance to handle the signal. If the signal sent is `SIGTERM`, the target process is given a chance to handle the signal before termination. Other signals may have different effects, depending on the signal and the process being targeted.

The `kill()` system call takes two arguments: the process ID of the target process and the signal to be sent. The process ID is a unique identifier that is assigned to each process on the system. The signal is an integer that represents the type of signal to be sent.

The `kill()` system call is used in a variety of situations, including:

- Termination of a process: When a process is no longer needed, the `kill()` system call can be used to terminate it.

- Error handling: If a process encounters an error condition, it may use the kill() system call to terminate itself or to request termination from another process.
- Process management: The kill() system call can be used to manage groups of processes. For example, a parent process may use the kill() system call to terminate all of its child processes.
- Debugging: The kill() system call can be used to send signals to a process for debugging purposes.

Example: Here's an example code in C that demonstrates the use of the kill() system call to send a signal to a process:

```
#include <stdio.h>
#include <signal.h>
#include <unistd.h>

int main() {
    pid_t pid;

    // Fork a child process
    pid = fork();

    if (pid == 0) {
        // Child process
        printf("Child process running...\n");

        // Wait for 10 seconds
        sleep(10);
    }
}
```



```

        // Print a message and terminate
        printf("Child process terminating...\n");
        return 0;
    }
else {
    // Parent process
    printf("Parent process running...\n");

    // Wait for 2 seconds
    sleep(2);

    // Send a signal to the child process
    printf("Sending signal to child process...\n");
    kill(pid, SIGTERM);

    // Print a message and terminate
    printf("Parent process terminating...\n");
    return 0;
}
}

```

In this example, the parent process forks a child process and then waits for 2 seconds before sending a SIGTERM signal to the child process using the kill() system call. The child process is programmed to terminate itself after 10 seconds. When the SIGTERM signal is sent, the child process receives it and terminates immediately.

Note that this is just a simple example, and in real-world scenarios, the `kill()` system call can be used to send a variety of signals to processes for different purposes.

1.6.7 The `pipe()` system call

The `pipe()` system call is a useful feature of operating systems that allows for inter-process communication. It creates a communication channel, or pipe, between two related processes, allowing them to exchange data. In this chapter, we will discuss the `pipe()` system call in detail, including its purpose, functionality, and how it is used in programming.

The `pipe()` system call is used to create a unidirectional data channel between two related processes, typically a parent and child process. This communication channel allows the processes to share data in a secure and efficient manner. The parent process can write data to the pipe, and the child process can read the data from the pipe. This is a useful feature for many applications, such as sending data from one process to another or for implementing interprocess synchronization.

The `pipe()` system call creates a pipe and returns two file descriptors representing the two ends of the pipe. One file descriptor represents the read end of the pipe, and the other file descriptor represents the write end of the pipe. The read end is used to read data from the pipe, while the write end is used to write data to the pipe. The data written to the write end of the pipe can be read from the read end of the pipe.

Example: Here is a sample code that demonstrates how to use the `pipe()` system call in programming. This code creates a pipe between a parent and child process, and the child process reads a message from the pipe that was written by the parent process:

```
#include <stdio.h>
#include <unistd.h>
```

```

int main() {
    int fd[2];
    pid_t pid;
    char buffer[20];

    // create pipe
    if (pipe(fd) == -1) {
        printf("Error creating pipe.\n");
        return 1;
    }

    // create child process
    pid = fork();

    if (pid < 0) {
        printf("Error creating child process.\n");
        return 1;
    }

    if (pid > 0) {
        // parent process
        close(fd[0]); // close read end of pipe
        write(fd[1], "Hello, child!", 13); // write message to pipe
        close(fd[1]); // close write end of pipe
    } else {

```

```
        // child process
        close(fd[1]); // close write end of pipe
        read(fd[0], buffer, 13); // read message from pipe
        printf("Received message: %s\n", buffer);
        close(fd[0]); // close read end of pipe
    }

    return 0;
}
```

In this code, the `pipe()` system call creates a pipe, which is represented by the `fd` array. The parent process writes the message "Hello, child!" to the write end of the pipe using the `write()` system call. The child process then reads the message from the read end of the pipe using the `read()` system call and prints it to the console.

1.7 Process hierarchies

Process hierarchies are an important aspect of modern operating systems. When a process creates another process, the parent-child relationship is established. The parent process creates the child process, and the child process inherits many of its parent's attributes, such as its environment variables and file descriptors.

In some systems, this parent-child relationship is maintained even after the child process has been created. This means that a child process can itself create more processes, forming a process hierarchy. In other words, the child process becomes a parent process and can create its own children.

This process hierarchy can be visualized as a tree-like structure, with the original parent process at the root of the tree and all subsequent

processes branching off from it. Each process in the hierarchy has a unique process ID (PID) assigned to it, which allows the operating system to track it and manage its resources.

One of the advantages of process hierarchies is that they allow for the efficient management of resources. For example, if a parent process creates a child process to perform a particular task, the child process can be terminated when the task is completed. This frees up resources that would have otherwise been tied up by the child process, allowing them to be used by other processes.

Another advantage of process hierarchies is that they make it easy to manage the relationship between different processes. For example, a parent process may need to communicate with its child process or monitor its activity. By maintaining the parent-child relationship, the operating system can provide a mechanism for this communication and monitoring to take place.

1.8 Zombie and orphan process

A zombie process is a process that has terminated but whose exit status has not yet been collected by its parent process. This means that the process entry still exists in the process table, but the process is no longer executing. In other words, a zombie process is a process that has completed execution but still has an entry in the process table.

An orphan process, on the other hand, is a child process that has terminated, but whose parent process has not yet collected its exit status. When the parent process terminates before collecting the exit status of its child process, the orphan process becomes a child of the init process (with process ID 1), which is responsible for collecting the exit status of all orphan processes.

When a process terminates, the operating system sets a flag indicating that the process is ready to be removed from the system. However, if the

parent process is still running and has not yet collected the exit status of its child process, the process table entry for the child process cannot be removed. Therefore, the operating system leaves the entry in the process table and marks the process as a zombie.

The parent process can collect the exit status of its child process by calling the `wait()` or `waitpid()` system call. When the parent process calls one of these system calls, it is suspended until one of its child processes terminates. At that point, the parent process is awakened and can collect the exit status of its child process.

In the case of an orphan process, the `init` process will automatically collect the exit status of the child process. This ensures that no process becomes a zombie and the system resources are effectively managed.

2 Process States and Transitions

A process is a fundamental concept in an operating system, and it is essential to understand the different states a process can be in and how it transitions between these states. The concept of process states is crucial in multi-tasking and concurrency because it allows the operating system to manage and prioritize the execution of processes.

In this chapter, we will cover the different process states, including new, ready, running, blocked, and terminated. We will also discuss how a process transitions between these states and the importance of understanding process states in a multi-tasking and concurrent environment.

2.1 Process states:

Processes are the core of an operating system, and they have various states throughout their execution. In this chapter, we will discuss the

different process states, which are new, ready, running, blocked, and terminated.

2.1.1 New State:

When a process is created, it is in the new state. At this point, the process is just an idea or a request, and the operating system has not yet allocated resources to it. Once the operating system assigns resources to the process, it moves to the next state.

2.1.2 Ready State:

When a process has been assigned resources, it moves to the ready state. In this state, the process is waiting to be allocated a processor by the operating system. The process is ready to execute, but the operating system has not yet assigned a processor to it.

2.1.3 Running State:

When the operating system assigns a processor to a process, it moves to the running state. In this state, the process is executing its instructions on the assigned processor.

2.1.4 Blocked State:

When a process is waiting for an event, such as I/O or a system resource, it moves to the blocked state. In this state, the process cannot execute until the event it is waiting for occurs. Once the event occurs, the process moves back to the ready state.

2.1.5 Terminated State:

When a process completes its execution, it moves to the terminated state. In this state, the process is no longer executing, and its resources have been deallocated by the operating system.

Understanding the different process states is crucial to the efficient operation of an operating system. The ability to manage and manipulate process states is a critical component of any operating system, allowing the operating system to allocate resources effectively and provide a seamless user experience.

2.2 Transitions between process states

Transitions between process states are crucial to the functioning of a modern operating system. A process can move between different states during its lifetime, depending on the type of operation it is currently performing. Understanding these transitions is key to understanding how a multi-tasking and concurrent system works.

When a process is first created, it enters the "new" state. In this state, the operating system has allocated resources such as memory and process control blocks to the process, but it is not yet ready to execute. From here, the process may transition to the "ready" state, where it is waiting for the CPU to be assigned to it.

Once the CPU is assigned to the process, it enters the "running" state. In this state, the process is actively executing its instructions. However, at any point, the process may be interrupted and transition back to the "ready" state. This can happen, for example, if the operating system needs to switch to another process that has become ready to execute.

A process can also transition to the "blocked" state, which occurs when the process is waiting for some external event, such as user input or data from a file. In this state, the process is not executing any instructions and is temporarily suspended.

Finally, a process may transition to the "terminated" state when it has completed its execution. In this state, the operating system releases any resources that were allocated to the process.

Understanding the transitions between process states is important for building a robust and efficient operating system. It enables the operating system to prioritize processes based on their state, ensuring that processes that are ready to execute are given access to the CPU. It also allows the operating system to manage resources effectively, by releasing resources when they are no longer needed.

2.3 Importance of process states in multi-tasking and concurrency

In a modern operating system, the ability to manage multiple processes simultaneously is a critical feature. This feature enables the system to be more efficient and responsive, as well as providing the ability to execute multiple programs simultaneously.

The management of multiple processes is a complex task, requiring careful attention to detail and the use of sophisticated algorithms. One of the key components of process management is the concept of process states. The state of a process is a reflection of its current activity level and can be used to control its behavior.

There are five process states: new, ready, running, blocked, and terminated. Each of these states plays an important role in the management of processes and is used to control how processes are scheduled for execution.

When a process is first created, it is in the new state. In this state, the process has been created but has not yet been assigned any resources or executed by the system. Once the process has been assigned the necessary resources and is ready to be executed, it enters the ready state.

In the ready state, the process is waiting for its turn to be executed by the system. The operating system uses a scheduling algorithm to determine which process to execute next from the pool of ready processes.

Once a process has been selected for execution, it enters the running state. In this state, the process is actively executing its code and using system resources.

The blocked state is used to indicate that a process is waiting for some external event to occur before it can continue executing. For example, a process may be blocked while waiting for user input or while waiting for a file to be loaded from disk.

Finally, when a process has completed its execution or has been terminated by the system, it enters the terminated state. In this state, the process is no longer executing and its resources have been freed by the system.

The importance of process states in multi-tasking and concurrency cannot be overstated. These states provide a mechanism for the operating system to control how processes are executed and to ensure that system resources are used efficiently.

By carefully managing the state of each process, the system can ensure that all processes are executed fairly and that no single process monopolizes the system's resources. Additionally, the use of process states makes it possible to manage complex multi-tasking and concurrency scenarios, allowing the system to execute multiple processes simultaneously without conflicts or other issues.

In conclusion, the management of process states is a critical aspect of modern operating systems. By providing a mechanism for controlling the behavior of processes and managing system resources, process states make it possible to execute multiple processes simultaneously and ensure that the system is both efficient and responsive.

3 CPU virtualization

In order to allow multiple processes to run seemingly at the same time on a single CPU, the operating system needs to virtualize the CPU. The basic idea behind CPU virtualization is simple: run one process for a short period of time, then switch to another process, and repeat. By time-sharing the CPU in this way, virtualization is achieved.

However, building the machinery required for CPU virtualization is not without its challenges. The first and foremost challenge is that of performance. The overhead introduced by the virtualization process should be minimal so that the system can operate as efficiently as possible. Every time a context switch occurs, there is a cost in terms of time and resources. Thus, the virtualization machinery must be designed to minimize this overhead while still allowing multiple processes to run on the same CPU.

The second challenge in building CPU virtualization machinery is control. The operating system must retain control over the CPU and the resources it manages. Without control, a rogue process could take over the machine, run forever, or access information it shouldn't have access to. Thus, the OS must be designed in such a way that it can maintain control over the CPU and the resources it manages, while still allowing processes to run efficiently.

Obtaining high performance while maintaining control is one of the central challenges in building an operating system. Many modern operating systems use a variety of techniques to achieve this, such as

preemptive multitasking, priority-based scheduling, and hardware support for virtualization. By carefully designing the virtualization machinery, the operating system can achieve high performance while maintaining control over the CPU and resources.

3.1 Limited Direct Execution

When it comes to running programs on an operating system, the simplest approach is to just run the program directly on the CPU. However, this approach poses a few challenges when it comes to virtualizing the CPU and implementing time sharing.

To begin with, if we just run a program, there is no way for the operating system to ensure that the program doesn't perform any actions that we don't want it to do. In other words, there is no control over the program's behavior, which can be dangerous in a multi-user system. The OS needs to ensure that it maintains control over the resources and activities of each program.

Secondly, when a process is running, the OS needs to be able to stop it and switch to another process, thus enabling time sharing and virtualization of the CPU. This is done through a mechanism called process scheduling, where the OS decides which process should be run next and for how long.

To overcome these challenges, the operating system needs to implement mechanisms for process management, memory management, and resource allocation. When the OS starts a program, it creates a process entry for it in a process list, allocates memory for it, loads the program code into memory from the disk, locates the program's entry point (usually the `main()` routine), and starts running the user's code.

To maintain control over the program's behavior, the OS employs various mechanisms such as system calls, which allow the program to

interact with the OS in a controlled manner. The OS also employs techniques such as memory protection, where each process is allocated a separate address space to prevent it from accessing the memory of other processes.

To implement time sharing and process scheduling, the OS uses a scheduler that decides which process should be run next and for how long. There are different types of schedulers such as round-robin, priority-based, and multi-level feedback queues.

In summary, although the "direct execution" approach of running programs on the CPU is simple, it poses several challenges when it comes to virtualizing the CPU and implementing time sharing. The operating system needs to employ various mechanisms and techniques to ensure control over the program's behavior and efficient process scheduling.

3.2 Restricted Operations

Direct execution, as we discussed earlier, has the benefit of being fast and efficient, but it also poses some significant challenges. One of the main issues is how to handle restricted operations that a process may try to perform. When a process runs directly on the CPU, it has access to all the hardware resources, including the disk, CPU, and memory.

However, this unrestricted access can be dangerous, especially when it comes to system security and stability. For example, if a process is allowed to access system memory without permission, it can cause a crash or even compromise the entire system's security.

To address this problem, operating systems implement a mechanism known as system calls. System calls provide a safe and controlled way for processes to interact with system resources, including I/O devices, memory, and other hardware.

When a process wants to perform a restricted operation, it must first make a system call to the operating system, requesting permission to access the resource. The operating system then performs the requested operation on behalf of the process and returns the result to the process.

For instance, when a process needs to access a file on the disk, it must first make a system call to the operating system to open the file. The operating system will then check whether the process has permission to access the file and perform the operation on the process's behalf.

3.2.1 User mode and kernel mode

When it comes to running processes on a computer, it is essential to have some way of controlling what a process can and cannot do. Allowing processes to perform unrestricted operations, such as issuing I/O requests, could result in the system being compromised or even taken over entirely. This is where the concept of user mode and kernel mode comes in.

User mode is a restricted mode that allows processes to run on the CPU while limiting what they can do. This mode is designed to prevent processes from accessing resources they shouldn't or performing operations they're not authorized to perform. For example, in user mode, a process cannot issue I/O requests. If it does, the CPU will raise an exception, and the operating system will kill the process to prevent any damage.

Kernel mode, on the other hand, is a privileged mode that only the operating system or kernel can run in. Code running in kernel mode can perform any operation, including restricted instructions and issuing I/O requests. This mode is designed to give the operating system full access to all system resources and to perform privileged operations.

The idea of user mode and kernel mode allows the operating system to have control over what processes can do on the system. By limiting the actions a process can perform in user mode, the operating system can

ensure the integrity and security of the system. At the same time, by allowing the operating system to run in kernel mode, the system can perform critical tasks that would otherwise be impossible.

3.2.2 System calls

One of the fundamental challenges in building an operating system is allowing user processes to perform privileged operations, such as accessing the file system or communicating with other processes. To enable this, modern hardware provides a mechanism known as system calls. System calls allow user programs to request services from the kernel, which is running in kernel mode and has the authority to perform these privileged operations.

The idea of system calls dates back to early machines such as the Atlas, where they were used to expose certain key pieces of functionality to user programs. Today, most operating systems provide a few hundred calls, each with a specific purpose. These include accessing the file system, creating and destroying processes, communicating with other processes, and allocating more memory. The POSIX standard provides a comprehensive list of system calls that are available on most modern operating systems.

When a user process wishes to perform a system call, it must first switch from user mode to kernel mode. This is typically accomplished through a software interrupt, which causes the processor to transition to kernel mode and begin executing the corresponding kernel code. Once the system call has completed, the kernel returns control to the user process and switches back to user mode.

System calls provide an essential interface between user programs and the kernel, allowing programs to perform privileged operations without compromising the security or stability of the system. However, the implementation of system calls can be complex, and care must be taken to ensure that they are designed and implemented correctly. Furthermore, the performance of system calls can have a significant

impact on the overall performance of the system, so optimization is often a critical concern for operating system developers.

3.2.3 System calls look like procedure calls

Have you ever wondered how a system call, such as `open()` or `read()`, looks like a regular procedure call in C? How does the system know it's a system call and perform all the necessary actions? In this chapter, we will explore why system calls look like procedure calls.

The answer is simple: a call to a system call is, in fact, a procedure call that hides a trap instruction. When a user program calls `open()`, for example, it executes a procedure call to the C library, which has an agreed-upon calling convention with the kernel. The library places the arguments for `open()` in predetermined locations (such as the stack or specific registers), sets the system call number in a well-known location (such as the stack or a register), and then executes the trap instruction. The code in the library then unpacks the return values and returns control to the program that made the system call.

As a result, the parts of the C library that make system calls are hand-coded in assembly language. They must carefully follow conventions to handle arguments and return values correctly and execute the hardware-specific trap instruction. Thus, you do not need to write assembly code to trap into an operating system because someone else has already written that assembly code for you.

In conclusion, system calls look like procedure calls because they are procedure calls that contain a trap instruction that triggers the kernel to perform the necessary operations. Understanding how system calls work is critical for anyone interested in operating systems and computer architecture.

3.2.4 The use of trap and return-from-trap instructions

When a program needs to perform a privileged operation, such as accessing the file system or communicating with other processes, it needs to execute a system call. As we discussed earlier, a system call is just like a regular procedure call, but it includes a special trap instruction that raises the privilege level to kernel mode and transfers control to the operating system. Once in kernel mode, the operating system can perform the requested operation on behalf of the calling program.

But what happens when the operating system is finished executing the system call? How does control return to the calling program? The answer lies in a special return-from-trap instruction, which is called by the operating system when it has finished executing the system call. This instruction returns control to the calling program while simultaneously reducing the privilege level back to user mode.

The use of trap and return-from-trap instructions is a fundamental technique used by modern operating systems to provide a safe and controlled environment for executing user programs. It allows programs to perform privileged operations without compromising the integrity and security of the system as a whole. By using these instructions, the operating system can ensure that user programs can only perform authorized operations, and that they do not interfere with other programs or the system as a whole.

In summary, the use of trap and return-from-trap instructions is an essential aspect of modern operating systems. They allow programs to execute system calls and perform privileged operations while maintaining the security and integrity of the system as a whole.

4 Process Control Block (PCB)

Welcome to the chapter on Process Control Block (PCB). In an operating system, managing processes is a critical task, and one of the key components used for this purpose is the Process Control Block. The Process Control Block is a data structure that contains essential information about a process, and it serves as a central point for the operating system to manage and control the process. In this chapter, we will cover the definition of a PCB, its contents, and its importance in process management. Understanding PCB is essential for anyone studying operating systems and its processes. So, let's dive into the world of PCBs and learn how they help in managing processes in an operating system.

In modern operating systems, process management is an essential component to ensure the efficient and effective execution of processes. One of the key structures used in process management is the Process Control Block (PCB). In this chapter, we will discuss what a PCB is, its functions, and how it is used in process management.

4.1 Definition of a PCB

A Process Control Block (PCB) is a data structure used by the operating system to manage information about a process. It is also known as a task control block or a process descriptor. A PCB is created by the operating system when a process is created and is responsible for keeping track of important information about the process, such as its current state, register values, memory allocation, and other attributes that are necessary for process execution.

4.2 Contents of a PCB

The contents of a PCB may vary depending on the operating system, but it generally contains the following information:

Process state: The current state of the process, which can be running, ready, blocked, or terminated.

Process ID: A unique identifier assigned by the operating system to each process.

Program counter: A register that holds the address of the next instruction to be executed.

CPU registers: The values of the CPU registers that are associated with the process, such as the accumulator, stack pointer, and index register.

Memory management information: Information about the memory allocated to the process, including its base address, limit, and page table information.

Priority: A value that determines the relative importance of the process compared to other processes.

I/O status information: Information about any I/O devices that are associated with the process, including the device status and any pending I/O operations.

4.3 Importance of PCB in process management

The PCB is a critical data structure used by the operating system to manage processes effectively. The operating system uses the information stored in the PCB to make decisions about how to allocate resources to the process, such as CPU time, memory, and I/O devices. Without the PCB, it would be difficult for the operating system to manage multiple processes concurrently and efficiently. The PCB is also

used by the operating system to switch between processes quickly, as it contains all the necessary information required to save the state of a process and restore it later.

In conclusion, the Process Control Block (PCB) is a vital component of process management in modern operating systems. It is responsible for storing and managing critical information about a process, including its current state, register values, memory allocation, and other attributes. The PCB allows the operating system to manage multiple processes efficiently and switch between them quickly, making it an essential part of any modern operating system.

5 Process Scheduling

In a multi-tasking and concurrent operating system, several processes compete for the same resources, such as CPU time and memory. The scheduler is responsible for assigning these resources to the processes efficiently and fairly. This is where process scheduling comes into play. In this chapter, we will discuss the definition of process scheduling, popular scheduling algorithms such as First-Come-First-Serve (FCFS), Shortest Job First (SJF), Round Robin (RR), Priority Scheduling, and Multilevel Feedback Queue (MLFQ), and the importance of process scheduling in multi-tasking and concurrency.

Process scheduling is a fundamental task in the operating system, and its primary purpose is to allocate CPU time to multiple processes in an efficient and effective manner. The CPU is a valuable resource that must be allocated carefully, as it can significantly impact the performance of the entire system. The scheduler is responsible for selecting the next process to run based on a set of predefined criteria. These criteria can vary depending on the scheduling algorithm used.

We will explore the most popular scheduling algorithms in detail in this chapter, including their strengths and weaknesses. We will also discuss the contents of a process control block (PCB), which is an essential data structure used by the scheduler to store information about each process. Finally, we will discuss the importance of process scheduling in multi-tasking and concurrency.

5.1 Definition of process scheduling

Process scheduling is a critical aspect of operating systems that allows for efficient and effective utilization of system resources. It involves determining which process will be executed by the CPU and when. The objective of process scheduling is to optimize system performance by reducing the turnaround time, waiting time, and response time of processes. In this chapter, we will explore the concept of process scheduling and its importance in modern operating systems.

Process scheduling is the mechanism used by operating systems to determine which process will be executed by the CPU next. The process scheduler is responsible for managing the queue of processes waiting to be executed and allocating system resources to them. The scheduler must balance competing demands for resources while ensuring that each process receives a fair share of CPU time. Process scheduling can be preemptive or non-preemptive. In preemptive scheduling, the CPU can be taken away from a process at any time, while in non-preemptive scheduling, a process holds the CPU until it voluntarily relinquishes it.

Process scheduling is a key component of any operating system, and different scheduling algorithms have been developed to manage system resources efficiently. These algorithms are designed to prioritize certain processes over others based on various criteria, such as the length of time a process has been waiting, its priority level, or the amount of CPU time it has already consumed.

5.2 How to develop scheduling policy

Scheduling policies are a fundamental aspect of operating systems. The scheduler is responsible for determining which processes to run, when to run them, and for how long. The goal of scheduling policies is to efficiently utilize system resources, while also providing good performance and fairness to all running processes. In this chapter, we will discuss the key considerations when developing scheduling policies.

The first step in developing a scheduling policy is to make some key assumptions about the system. The most important assumption is that the system has limited resources, including CPU time, memory, and I/O devices. Additionally, we assume that there are multiple processes that require access to these resources. These processes can have different priorities, requirements, and characteristics, such as I/O-bound or CPU-bound.

The next step is to determine the metrics that are important for evaluating the performance of the scheduling policy. The two most common metrics are throughput and turnaround time. Throughput measures the number of processes that are completed per unit of time, while turnaround time measures the time it takes for a process to complete from start to finish. Other important metrics include response time, which measures the time it takes for a process to start executing, and fairness, which measures how equitably resources are distributed among processes.

There are many different approaches to scheduling policies. One of the earliest approaches was the first-come, first-served (FCFS) policy, which simply scheduled processes in the order that they arrived in the system. However, this policy can result in long turnaround times for processes that arrive later and have to wait for earlier processes to complete. Another approach is the shortest job first (SJF) policy, which schedules the process with the shortest expected running time. This policy can

improve turnaround times, but requires knowledge of the expected running time of each process, which may not be available.

Other popular scheduling policies include round-robin, priority-based, and multilevel feedback queue (MLFQ) policies. The round-robin policy allocates a fixed amount of CPU time to each process in turn, while the priority-based policy assigns a priority level to each process and schedules them in order of priority. The MLFQ policy divides processes into different queues based on their characteristics, and applies different scheduling policies to each queue.

Developing an effective scheduling policy is a complex task that requires careful consideration of the system's resources, metrics, and available approaches. By understanding these key considerations, operating system developers can design scheduling policies that efficiently allocate resources, provide good performance and fairness, and meet the needs of the system and its users.

5.3 Scheduling algorithms:

As we have seen earlier, the operating system (OS) uses various mechanisms to manage and share resources efficiently. However, these mechanisms alone are not enough to ensure optimal performance and utilization of resources. This is where the policies come into play. Policies are algorithms that help the OS make decisions about resource allocation and scheduling.

One of the most common policies in the OS is the scheduling policy. The scheduling policy determines which program or process should be executed next on the CPU. The scheduling algorithm takes into account various factors such as the priority of the process, the amount of time it has already spent executing, and the type of workload the system is experiencing. The ultimate goal of the scheduling policy is to maximize

the overall throughput of the system, while also ensuring that interactive processes are given a fair share of the CPU time.

Another example of a policy is the memory allocation policy. The OS needs to decide which process gets to use which part of the memory. The memory allocation policy takes into account the size of the process, the amount of memory currently available, and the amount of memory requested by other processes. The goal of the memory allocation policy is to ensure that all processes get their required amount of memory, while also minimizing the fragmentation of memory.

In addition to the scheduling and memory allocation policies, there are many other policies implemented in the OS. For example, there are policies for disk scheduling, network bandwidth allocation, and power management. Each of these policies has its own unique algorithm that takes into account the specific characteristics of the resource being managed.

Policies are an essential part of the OS because they allow the system to adapt to changing workloads and usage patterns. Without policies, the OS would be unable to make intelligent decisions about how to allocate resources and manage competing demands. By combining policies with the mechanisms we discussed earlier, the OS is able to provide a high level of performance, reliability, and ease of use to its users.

5.3.1 First-Come-First-Serve (FCFS)

First-Come-First-Serve (FCFS) is one of the simplest CPU scheduling algorithms used in operating systems. In this algorithm, the process that arrives first is allocated the CPU first. The FCFS algorithm is also known as the First-In-First-Out (FIFO) scheduling algorithm because the process that comes first into the ready queue will be the first one to be executed.

When a process enters the ready queue, it is assigned the CPU based on the order in which it entered the queue. The process continues to use the CPU until it finishes its execution, blocks for I/O, or terminates.

The FCFS algorithm is simple to implement and understand, but it can cause long waiting times for the processes that arrive later. This is because a process with a longer burst time can hold the CPU for a long time, causing other processes to wait in the ready queue.

Moreover, the FCFS algorithm can lead to a phenomenon known as convoy effect. The convoy effect occurs when a long process holds the CPU, causing other short processes to wait behind it. This can lead to poor resource utilization and long average waiting times.

Despite its drawbacks, the FCFS algorithm is still widely used in batch processing systems, where it is essential to execute processes in the order they were submitted. However, it is not suitable for interactive systems, where users expect a quick response time.

In summary, the FCFS algorithm is a simple and straightforward scheduling algorithm that can be used in batch processing systems. However, it can cause long waiting times and poor resource utilization in interactive systems.

5.3.2 Shortest Job First (SJF)

In process scheduling, the goal is to maximize the throughput, minimize response time, and minimize turnaround time. The Shortest Job First (SJF) scheduling algorithm is a non-preemptive algorithm that prioritizes the process with the shortest CPU burst time.

The SJF scheduling algorithm is a non-preemptive algorithm that selects the process with the smallest CPU burst time. The CPU burst time is the amount of time required by a process to complete its execution on the CPU. In this algorithm, the ready queue is maintained in order of the burst time of the processes.

The SJF scheduling algorithm can be either preemptive or non-preemptive. In preemptive SJF scheduling, if a new process with a shorter burst time enters the ready queue, the currently executing process is interrupted, and the new process is scheduled to execute. In non-preemptive SJF scheduling, once a process starts executing, it continues until completion or until it enters the blocked state.

The SJF scheduling algorithm has the following advantages:

- It results in the shortest average waiting time for the processes in the ready queue.
- It is optimal in the sense that it minimizes the average waiting time for the processes, assuming that the CPU burst times are known in advance.
- It prioritizes short processes, which leads to a faster turnaround time.

The SJF scheduling algorithm has the following disadvantages:

- It requires knowledge of the CPU burst time of each process, which is not always available.
- It can lead to starvation for long processes if there are a large number of short processes in the system.

The SJF scheduling algorithm is a non-preemptive algorithm that selects the process with the shortest CPU burst time. It prioritizes short processes, resulting in a faster turnaround time and shorter average waiting time for processes in the ready queue. However, it requires knowledge of the CPU burst time of each process, which may not always be available, and it can lead to starvation for long processes if there are a large number of short processes in the system.

5.3.3 Priority Scheduling

Priority Scheduling is one of the most widely used scheduling algorithms in modern operating systems. As the name suggests, this algorithm assigns priorities to each process based on their importance and then schedules them according to their priority. This approach ensures that the high-priority tasks get executed first, thus improving the overall performance and responsiveness of the system.

The basic idea behind Priority Scheduling is to assign a priority level to each process based on its importance. This priority level can be determined based on various factors, such as the amount of CPU time required by the process, the amount of memory it needs, or the urgency of the task. Once the priority levels are assigned, the scheduler can then schedule the processes based on their priority, giving the highest-priority process the CPU time first.

Priority levels can be assigned either statically or dynamically. In static priority scheduling, the priority level of a process is fixed at the time of creation and remains unchanged throughout the life of the process. In dynamic priority scheduling, on the other hand, the priority level of a process can be adjusted dynamically based on its behavior.

Priority Scheduling can be further divided into two different scheduling policies: Preemptive and Non-preemptive. In Preemptive Priority Scheduling, a higher-priority process can interrupt a lower-priority process while it is running. This allows the higher-priority process to start executing immediately, thus ensuring that the most important tasks get executed first. In Non-preemptive Priority Scheduling, however, a running process cannot be interrupted by a higher-priority process. In this case, the scheduler must wait for the currently running process to finish before scheduling the higher-priority process.

One of the key advantages of Priority Scheduling is that it allows the most important tasks to be executed first. This improves the overall performance and responsiveness of the system, especially in real-time

systems where timely execution of critical tasks is essential. Priority Scheduling also allows for efficient utilization of system resources, as it ensures that the most important tasks get executed first, thus reducing wastage of CPU time.

One major disadvantage of Priority Scheduling is that it can lead to starvation. If a low-priority process is constantly preempted by higher-priority processes, it may never get a chance to execute, thus leading to starvation. Another disadvantage is that it can lead to low-priority processes getting neglected. If there are too many high-priority processes, the low-priority processes may never get a chance to execute, leading to poor performance and responsiveness.

Priority Scheduling is widely used in modern operating systems, such as Linux, Windows, and macOS. In these systems, priority levels are assigned based on various factors, such as the type of process, its behavior, and its importance. The scheduler then uses these priorities to schedule the processes, giving the highest-priority process the CPU time first.

5.3.4 Round Robin (RR)

Round Robin (RR) is a widely used CPU scheduling algorithm in operating systems. This algorithm is designed to schedule multiple processes concurrently in a fair and efficient manner. In this chapter, we will discuss the details of the Round Robin scheduling algorithm.

The Round Robin (RR) scheduling algorithm is a preemptive scheduling algorithm. In this algorithm, each process is assigned a fixed time interval, called a time quantum or time slice, to use the CPU. The time quantum is usually a small unit of time, typically between 10 to 100 milliseconds.

When a process is given the CPU, it is allowed to run for a time quantum. If the process completes its task before the end of the time quantum, it releases the CPU voluntarily. If the process does not complete its task

before the end of the time quantum, it is preempted and moved to the back of the ready queue. The next process in the ready queue is then given the CPU to execute.

If a process arrives while the CPU is busy, it is placed at the end of the ready queue. This ensures that processes are executed in the order in which they arrived.

Round Robin (RR) scheduling algorithm offers the following advantages:

- **Fairness:** The Round Robin scheduling algorithm provides fairness to all processes by giving each process an equal opportunity to use the CPU.
- **Time-sharing:** The Round Robin scheduling algorithm is ideal for time-sharing systems where multiple users access the system simultaneously.
- **Low response time:** The Round Robin scheduling algorithm ensures that each process gets a turn to execute quickly, resulting in a low response time.

The Round Robin scheduling algorithm has the following disadvantages:

- **Inefficiency:** The Round Robin scheduling algorithm can be inefficient when the time quantum is too long or too short.
- **Overhead:** The Round Robin scheduling algorithm has a higher overhead compared to other scheduling algorithms, as it requires the scheduler to keep track of the time quantum for each process.

The Round Robin (RR) scheduling algorithm is a widely used scheduling algorithm in operating systems. It provides fairness to all processes and

is ideal for time-sharing systems. However, it can be inefficient and has a higher overhead compared to other scheduling algorithms.

5.3.5 Multilevel Queue Scheduling (MLQS)

In the previous chapters, we discussed various CPU scheduling algorithms, such as FCFS, SJF, Priority Scheduling, and Round Robin. In this chapter, we will discuss another important algorithm called Multilevel Queue Scheduling (MLQS). The MLQS algorithm is widely used in modern operating systems, especially in systems that need to manage multiple types of processes with varying priority levels.

Multilevel Queue Scheduling is a scheduling algorithm that divides the ready queue into several separate queues, each with its own scheduling algorithm. Each queue has its own priority level, and the scheduling algorithm is applied to each queue based on the priority level. This approach allows the operating system to prioritize different types of processes based on their needs, without compromising the responsiveness of the system.

In Multilevel Queue Scheduling, the ready queue is divided into multiple queues, each with its own priority level. The operating system assigns each process to a queue based on its priority level. Each queue can have its own scheduling algorithm, such as FCFS, SJF, Priority Scheduling, or Round Robin.

In most implementations of MLQS, the highest priority queue is served first, followed by the next highest priority queue, and so on. Within each queue, the scheduling algorithm is applied to determine the order in which processes are executed. The processes in each queue are typically scheduled in a non-preemptive manner, meaning that a process must yield the CPU voluntarily before another process can be executed.

Multilevel Queue Scheduling is an important scheduling algorithm in modern operating systems. It allows the operating system to prioritize different types of processes based on their priority level and the needs

of the system. This is important in systems that need to manage multiple types of processes with varying priority levels, such as real-time systems and systems that run both user-level and system-level processes.

The MLQS algorithm can also help improve the responsiveness of the system by ensuring that high-priority processes are given priority access to the CPU. By dividing the ready queue into multiple queues, the operating system can ensure that each type of process is given the appropriate amount of CPU time, without compromising the performance of the system.

Multilevel Queue Scheduling is an important scheduling algorithm in modern operating systems. It allows the operating system to prioritize different types of processes based on their priority level and the needs of the system. The MLQS algorithm can help improve the responsiveness of the system and ensure that high-priority processes are given priority access to the CPU.

5.3.6 Multilevel Feedback Queue Scheduling (MLFQS)

Multilevel Feedback Queue Scheduling (MLFQS) is a complex scheduling algorithm that dynamically adjusts the priority of processes based on their behavior over time. It is an extension of the Multilevel Queue Scheduling (MLQS) algorithm and is widely used in modern operating systems.

The MLFQS algorithm works by dividing the ready queue into multiple priority queues. Each queue is assigned a different priority level, with the highest priority queue being reserved for the most important processes, such as system processes or real-time processes. Each queue also has its own scheduling algorithm, with different algorithms being used for different priority levels.

When a process is created, it is placed in the highest priority queue. The process is then given a certain amount of time to execute, known as a time slice or quantum. If the process completes its execution before the

time slice expires, it is removed from the queue. If the time slice expires before the process completes its execution, the process is preempted and moved to a lower priority queue.

The priority of a process in MLFQS is determined dynamically based on its behavior over time. Processes that use a lot of CPU time are given lower priority to prevent them from monopolizing the CPU. Conversely, processes that use less CPU time are given higher priority to ensure that they are executed quickly.

MLFQS also includes a mechanism for promoting and demoting processes between priority levels. When a process is blocked, it is moved to a lower priority queue. When it becomes unblocked, it is moved back to its original priority level. This ensures that long-running processes do not hold up the system and that important processes are executed as quickly as possible.

In summary, Multilevel Feedback Queue Scheduling (MLFQS) is a powerful scheduling algorithm that dynamically adjusts the priority of processes based on their behavior over time. It is designed to ensure that important processes are executed quickly while preventing long-running processes from monopolizing the CPU.

5.3.7 Lottery Scheduling

In operating systems, the lottery scheduling algorithm is a probabilistic scheduling algorithm used to allocate resources to processes. It is a unique scheduling algorithm because it provides equal opportunities for all processes to win a "lottery ticket" and acquire CPU time. The algorithm uses a lottery ticket metaphor to determine which process gets the CPU next, making the process selection entirely random.

The lottery scheduling algorithm assigns each process a set of lottery tickets. A lottery ticket represents the process's share of the CPU time. The more lottery tickets a process has, the higher the chances of it winning the lottery and acquiring the CPU.

To allocate CPU time, the lottery scheduling algorithm randomly selects a ticket from the ticket pool. The process that owns the ticket wins the lottery and gets to run on the CPU for a set time quantum. After the time quantum expires, the process returns the CPU, and the lottery begins again.

The lottery scheduling algorithm uses a number of data structures to maintain the ticket pool and manage processes' ticket allocation. One such data structure is a list of processes, each with a number of lottery tickets associated with it. The algorithm also maintains a list of unused tickets and a counter that tracks the number of tickets in the pool.

The lottery scheduling algorithm has several advantages over other scheduling algorithms. One significant advantage is its ability to provide fairness to all processes. Since each process has an equal chance of winning the lottery, the scheduling algorithm ensures that no process is left behind or unfairly treated.

Another advantage of the lottery scheduling algorithm is its simplicity. The algorithm is easy to implement and does not require complex data structures or sophisticated algorithms. This simplicity translates to low overhead costs and makes it a good choice for systems with limited resources.

However, the lottery scheduling algorithm has some disadvantages. One significant disadvantage is that the algorithm is entirely random. There is no guarantee that a process with a large number of tickets will win the lottery or that a process with few tickets will not win the lottery several times in a row. This randomness can lead to inefficiencies in the system's performance and unpredictability in the scheduling results.

Additionally, the lottery scheduling algorithm may not be suitable for systems with strict scheduling requirements. The randomness of the algorithm may lead to scheduling delays and missed deadlines, which can be detrimental in real-time systems.

The lottery scheduling algorithm is a unique scheduling algorithm that uses a lottery ticket metaphor to allocate CPU time to processes. The algorithm provides fairness to all processes and is easy to implement, making it a good choice for systems with limited resources.

However, the algorithm's randomness can lead to inefficiencies and unpredictability in the scheduling results, making it unsuitable for systems with strict scheduling requirements. In such cases, other scheduling algorithms, such as Round Robin or Priority Scheduling, may be more appropriate.

5.3.8 Fair-Share Scheduling

In a multi-user system, it is important to ensure fairness and prevent any single user from monopolizing system resources. Fair-Share Scheduling is a scheduling algorithm that addresses this issue by allocating system resources fairly among all users. In this chapter, we will discuss the concept of Fair-Share Scheduling and how it works in an operating system.

Fair-Share Scheduling is a scheduling algorithm that dynamically allocates system resources based on the proportion of resources each user is entitled to. Each user is assigned a "share" of the system resources, and the scheduler ensures that each user gets their fair share. The concept of fair sharing can be applied to various system resources, including CPU time, memory, and I/O devices.

The Fair-Share Scheduling algorithm works by maintaining a record of each user's resource usage over time. This record is used to calculate each user's share of the resources based on a predetermined policy. The policy may be based on factors such as the number of active processes, the amount of CPU time used, or a combination of factors.

When a user requests a resource, such as CPU time or memory, the scheduler checks their entitlement to that resource based on the user's share. If the user's share has been used up, they may be placed in a

waiting queue until their share becomes available again. This prevents any user from monopolizing system resources and ensures that all users are treated fairly.

The main advantage of Fair-Share Scheduling is that it ensures fairness in the allocation of system resources. This is particularly important in multi-user systems where resources are shared among multiple users. By dynamically adjusting each user's share based on their resource usage, the scheduler ensures that no user can monopolize the system resources.

Another advantage of Fair-Share Scheduling is that it allows administrators to set policies that reflect the organization's priorities. For example, a policy can be set to give higher priority to certain users or groups of users, based on their role within the organization.

One potential disadvantage of Fair-Share Scheduling is that it can be complex to implement and maintain. The scheduler needs to keep track of each user's resource usage over time, which requires additional system overhead. Additionally, the policies used to calculate each user's share can be complex and difficult to configure.

Another potential disadvantage of Fair-Share Scheduling is that it may not be suitable for all types of systems. For example, in a system where users are performing real-time tasks, such as video streaming or audio processing, Fair-Share Scheduling may not provide the necessary responsiveness.

Fair-Share Scheduling is a scheduling algorithm that ensures fairness in the allocation of system resources. It works by dynamically adjusting each user's share based on their resource usage, preventing any single user from monopolizing system resources. While Fair-Share Scheduling has its advantages, such as allowing administrators to set policies that reflect the organization's priorities, it also has potential disadvantages, such as increased complexity and possible lack of responsiveness in real-time systems.

5.3.9 Guaranteed Scheduling

In the context of Operating Systems, a Guaranteed Scheduling Algorithm is a type of scheduling algorithm that ensures that certain processes are guaranteed a certain amount of CPU time, regardless of the presence of other processes in the system. This is particularly useful in situations where there are real-time processes that require a certain level of responsiveness from the system.

In a Guaranteed Scheduling Algorithm, the system sets aside a certain amount of CPU time for specific processes, known as guaranteed processes. These guaranteed processes are given a certain priority level, which determines the amount of CPU time they are allocated. Once a guaranteed process is scheduled to run, it is given the CPU until it either completes or reaches its maximum allocated time slice.

If there are no guaranteed processes ready to run, the system switches to a different scheduling algorithm to assign CPU time to non-guaranteed processes. This helps to ensure that non-guaranteed processes don't starve for CPU time.

One of the main advantages of a Guaranteed Scheduling Algorithm is that it ensures that certain processes receive the CPU time they need to function properly. This is particularly important in real-time systems, where a delay in processing a critical task could have serious consequences.

Another advantage of Guaranteed Scheduling Algorithm is that it allows for more precise control over system resources. By allocating specific amounts of CPU time to specific processes, system administrators can ensure that all processes receive an appropriate amount of resources, while also preventing any single process from monopolizing the CPU.

The main disadvantage of a Guaranteed Scheduling Algorithm is that it can be difficult to implement in a way that balances the needs of all processes in the system. For example, if too much CPU time is allocated

to guaranteed processes, non-guaranteed processes may experience unacceptable levels of latency or may be starved for CPU time.

Additionally, a Guaranteed Scheduling Algorithm can be complex to implement, requiring careful tuning of the system parameters and a thorough understanding of the needs of each process in the system.

5.4 Importance of process scheduling in multi-tasking and concurrency

In multi-tasking and concurrent environments, it is essential to have a proper process scheduling mechanism in place. The scheduling algorithm plays a vital role in deciding which process gets to execute on the CPU and for how long. The process scheduling algorithm decides the efficiency of an operating system in handling multiple tasks simultaneously.

The primary goal of process scheduling is to improve system performance by reducing the CPU idle time and improving the response time of the system. The following are the key reasons why process scheduling is essential in a multi-tasking and concurrent environment:

Resource Utilization: In a multi-tasking environment, several processes compete for resources such as CPU, memory, and I/O devices. Process scheduling ensures that these resources are allocated efficiently and utilized to their maximum capacity.

Throughput: Process scheduling influences the throughput of the system. The throughput refers to the number of processes that the system can execute in a given period. A good process scheduling algorithm can significantly improve the throughput of the system.

Response Time: The response time of a system refers to the time taken by the system to respond to a user's input. A good process scheduling

algorithm can ensure that the system responds to the user's input promptly.

Fairness: Process scheduling ensures that every process gets a fair share of the system resources. A fair process scheduling algorithm can ensure that no process is starved of system resources.

Prioritization: Process scheduling can prioritize processes based on their importance. The priority of a process determines its access to the system resources. A good process scheduling algorithm can ensure that critical processes get higher priority, ensuring that the system operates efficiently.

In conclusion, process scheduling is an essential aspect of any operating system, particularly in a multi-tasking and concurrent environment. The scheduling algorithm used by an operating system can significantly impact its performance, response time, throughput, fairness, and prioritization. Operating system designers need to consider these factors while designing the process scheduling algorithm to ensure that the operating system is efficient and responsive.

6 Interprocess Communication (IPC) and Synchronization

In this chapter, we will explore the various methods of IPC and synchronization, such as shared memory, message passing, semaphores, and monitors. These methods provide a means for different processes to exchange information and coordinate their activities. We will also discuss the importance of IPC and synchronization in multi-tasking and concurrency, and how they help in preventing race conditions, deadlocks, and other synchronization problems that may arise when multiple processes access shared resources simultaneously. Overall, this

chapter will provide you with an understanding of how IPC and synchronization play a crucial role in the effective management of processes in operating systems.

6.1 Definition of IPC and synchronization

Interprocess Communication (IPC) and synchronization are two important concepts in operating systems that are necessary for effective multi-tasking and concurrency.

IPC refers to the mechanisms and techniques used by different processes to communicate with each other and share resources. In a multi-tasking environment, it is essential for different processes to communicate with each other to coordinate their activities, share data, and perform tasks collaboratively.

Synchronization, on the other hand, refers to the process of coordinating access to shared resources among different processes. In a multi-tasking environment, multiple processes may require access to the same resources simultaneously, and synchronization ensures that they do not interfere with each other.

IPC and synchronization are closely related concepts, as synchronization is necessary for proper communication and resource sharing between processes. There are various methods available for IPC and synchronization, each with its own advantages and disadvantages.

In the following chapters, we will explore different methods of IPC and synchronization in detail, along with their advantages, disadvantages, and real-world applications.

6.2 Race conditions

Race conditions are a common issue in operating systems, particularly in systems that allow multiple processes to access shared storage. In a race condition, two or more processes or threads access shared storage concurrently, and the order in which they access the storage affects the outcome of the program. This can lead to unpredictable behavior and errors in the system.

For example, in a print spooler, two or more processes might try to add a file to the spooler directory at the same time. If the directory is not protected against simultaneous access, one process might overwrite the file added by another process. Alternatively, the printer daemon might try to print a file that has not been completely spooled, leading to printing errors.

To avoid race conditions, operating systems use various synchronization mechanisms, such as semaphores, mutexes, and monitors. These mechanisms ensure that only one process or thread can access shared storage at a time, preventing conflicts and ensuring the correct operation of the system.

It's important to note that while synchronization mechanisms can prevent race conditions, they can also introduce new problems, such as deadlocks and livelocks. Deadlocks occur when two or more processes or threads are blocked, waiting for each other to release a resource. Livelocks occur when two or more processes or threads are blocked, but continue to perform nonproductive actions, preventing progress in the system.

In summary, race conditions are a common issue in operating systems that allow multiple processes or threads to access shared storage. To avoid race conditions, operating systems use synchronization mechanisms such as semaphores, mutexes, and monitors. However, these mechanisms can introduce new problems such as deadlocks and

livelocks. Therefore, careful design and implementation are necessary to ensure the correct and efficient operation of the system.

6.3 Critical Regions

When multiple processes access shared data, they may run into race conditions, causing unpredictable behavior and potential data corruption. To avoid these issues, we need to ensure that only one process accesses the shared data at any given time. This is achieved by creating critical regions or mutual exclusion.

A critical region is a section of code in a process that accesses shared resources. To prevent race conditions, only one process should execute the critical region at any given time. This is accomplished by using synchronization mechanisms that ensure exclusive access to the critical region.

The two most commonly used synchronization mechanisms for implementing critical regions are semaphores and monitors. Semaphores are integer variables that are used to control access to shared resources. Monitors are a higher-level synchronization construct that includes shared variables and procedures for accessing them.

Semaphores and monitors work by allowing only one process to access the shared resource at a time. If a process tries to access the resource while it is being used by another process, it is blocked until the resource is released. This ensures that only one process executes the critical region at any given time.

One important consideration when implementing critical regions is deadlock, a situation where two or more processes are blocked, waiting for resources that are held by other blocked processes. Deadlock can occur when synchronization mechanisms are not properly implemented, leading to a situation where no progress can be made.

To prevent deadlock, it is essential to ensure that synchronization mechanisms are used consistently throughout the system. This includes ensuring that all critical regions are properly protected and that processes release shared resources when they are no longer needed.

In summary, critical regions are essential for ensuring that shared resources are accessed safely and effectively in a multi-process system. By implementing synchronization mechanisms such as semaphores and monitors, we can prevent race conditions and ensure that processes access shared resources in a mutually exclusive manner. However, care must be taken to prevent deadlock by ensuring consistent use of synchronization mechanisms throughout the system.

6.4 Mutual Exclusion with Busy Waiting

One of the most straightforward ways to achieve mutual exclusion is through busy waiting. In this approach, a process repeatedly tests a shared variable in a loop, waiting until the variable is available for use. The process then enters its critical region, performs its task, and then releases the variable so that another process can access it.

While this method is simple to understand and implement, it has some significant drawbacks. The main issue with busy waiting is that it can waste a lot of CPU time. Since the process is continuously looping while waiting for the variable to be available, it is using CPU resources that could be better used by other processes. In addition, if a process forgets to release the variable after it is done, other processes will be locked out of the critical region indefinitely.

Despite these issues, busy waiting is still used in some situations. For example, it is commonly used in low-level synchronization primitives such as spinlocks, which are used in kernel-level code to protect shared data structures.

To mitigate the downsides of busy waiting, operating systems also provide alternative methods for mutual exclusion, such as sleep-and-wakeup and semaphores, which we will explore in the following sections.

6.4.1 Disabling interrupts

Disabling interrupts is a technique for achieving mutual exclusion in a single-processor system. When a process enters its critical region, it disables all interrupts to prevent other processes from interrupting it and accessing the shared memory. Once the process has finished updating the shared memory, it re-enables interrupts before leaving the critical region.

While disabling interrupts may seem like a simple and effective solution, it has some drawbacks. For one, it can cause the system to become unresponsive if a process forgets to re-enable interrupts, as the clock interrupt will not occur and the system will hang. Additionally, disabling interrupts can also lead to problems with I/O operations, such as disk accesses, which require interrupts to signal when an operation is complete.

Despite these drawbacks, disabling interrupts can still be a useful technique in certain situations, particularly in embedded systems where the hardware is tightly integrated with the software and the system is designed specifically to handle this type of mutual exclusion. However, in most general-purpose operating systems, disabling interrupts is not the preferred method for achieving mutual exclusion, as there are other techniques that are safer and more efficient, such as semaphores and mutexes.

6.4.2 Lock variables

Lock variables are a commonly used software solution to achieve mutual exclusion. The idea is simple: have a single, shared variable that acts as a lock. When a process wants to enter its critical region, it first tests the

lock. If the lock is 0, the process sets it to 1 and enters the critical region. If the lock is already 1, the process just waits until it becomes 0 again.

The use of lock variables is a simple and effective way of ensuring mutual exclusion. However, there are some potential issues that need to be considered. For example, suppose two processes are waiting to enter their critical regions, and the lock variable is currently 0. Both processes may simultaneously try to set the lock variable to 1, which could result in a race condition. In order to avoid such problems, some form of atomic operation is required to set the lock variable to 1.

Another issue to consider is the possibility of deadlock. Deadlock can occur when two or more processes are waiting for each other to release locks that they currently hold. For example, suppose two processes each hold a lock variable, and they are both waiting to acquire the other's lock. This situation can be avoided by ensuring that lock variables are always acquired in a predetermined order.

Lock variables can also be used to implement more sophisticated synchronization mechanisms. For example, a semaphore is a synchronization primitive that uses a lock variable and a counter. The counter is used to keep track of the number of available resources, and the lock variable is used to ensure that only one process at a time can access the counter.

Overall, lock variables are a simple and effective way of ensuring mutual exclusion in a shared memory system. However, care must be taken to avoid race conditions and deadlocks, and more sophisticated synchronization mechanisms may be required in some situations.

6.4.3 Strict alternation

Strict alternation is a technique for achieving mutual exclusion that uses a shared boolean variable to enforce strict alternation between two processes. The idea behind strict alternation is that two processes, P_0 and P_1 , take turns accessing a shared resource. When P_0 wants to enter

its critical section, it first sets the shared boolean variable to true, indicating that it is its turn to enter the critical section. It then enters its critical section and performs its task. When P₀ is finished, it sets the shared boolean variable to false, indicating that it is P₁'s turn to enter the critical section. P₁ can then proceed in the same way, and so on.

Example: Here is an example implementation of strict alternation:

```
bool turn = false; // shared boolean variable

// Process P0
while (true) {
    while (turn == true); // wait for P1 to finish
    // critical section for P0
    turn = true; // pass the turn to P1
}

// Process P1
while (true) {
    while (turn == false); // wait for P0 to finish
    // critical section for P1
    turn = false; // pass the turn to P0
}
```

In this implementation, the two processes P₀ and P₁ take turns entering their critical sections, based on the value of the shared boolean variable `turn`. The while loops ensure that a process waits until it is its turn to enter the critical section, and the setting of `turn` at the end of each critical section ensures that the other process will eventually have its turn.

While strict alternation is simple and easy to implement, it has several limitations. The most significant limitation is that it does not scale well to more than two processes. With more than two processes, the processes would need to wait for each other in a specific order, which can become complicated and error-prone. Additionally, strict alternation can lead to starvation, where a process is unable to enter its critical section because it is repeatedly passed over in favor of the other process. To avoid starvation, it may be necessary to introduce additional rules or mechanisms, such as a timeout or a fairness policy.

6.4.4 Peterson's solution

Peterson's solution is a software-based algorithm designed to solve the mutual exclusion problem in a concurrent system without using busy waiting. The algorithm was proposed by Gary L. Peterson in 1981, and it builds on the idea of taking turns between two processes in a critical section.

Peterson's solution uses two shared variables, turn and flag, to control the entry of two processes into their critical regions. The turn variable is used to indicate whose turn it is to enter the critical section, and the flag variable is used to indicate whether a process wants to enter its critical region. The two processes alternate turns, and when one process is in its critical section, the other process waits until it is its turn.

The algorithm works as follows:

- Initialize the turn variable to either 0 or 1, depending on which process should go first.
- Initialize both flag variables to false.
- Process P₀ sets its flag to true and sets turn to 1.
- Process P₀ enters the critical section if and only if P₁'s flag is false or turn is 0.
- Process P₀ exits the critical section and sets its flag to false.
- Process P₁ sets its flag to true and sets turn to 0.

- Process P₁ enters the critical section if and only if P₀'s flag is false or turn is 1.
- Process P₁ exits the critical section and sets its flag to false.

The key idea behind Peterson's solution is that a process checks the turn variable to see if it is its turn to enter the critical section. If it is not its turn, the process waits until it is its turn. Additionally, a process sets its flag to indicate that it wants to enter the critical section. If the other process's flag is already set, then the process waits until the other process has finished.

Peterson's solution is an improvement over strict alternation and lock variables because it avoids busy waiting. The algorithm is also simple and easy to implement, although it only works for two processes.

However, Peterson's solution suffers from several drawbacks. The algorithm assumes that processes take turns, which may not be true in all cases. Moreover, the algorithm does not work for more than two processes, and it is vulnerable to priority inversion, a situation where a low-priority process holds the lock and prevents a high-priority process from entering its critical section.

Overall, Peterson's solution is a useful algorithm for solving the mutual exclusion problem in a concurrent system, but it has limitations that must be taken into account when designing a real-time operating system.

6.4.5 The TSL instruction

In the previous sections, we have examined various software solutions to the mutual exclusion problem. In this section, we will take a look at a hardware solution that relies on a special instruction available on some computers called Test and Set Lock (TSL).

TSL is a single instruction that performs two operations atomically: it reads the value at a given memory location into a register and sets the

value of the memory location to a non-zero value. This means that the read and write operations happen without any interruption from other processors, ensuring mutual exclusion.

The TSL instruction is especially useful in systems with multiple processors where software solutions like lock variables and Peterson's solution may not work effectively. By relying on a hardware instruction, TSL avoids the need for busy waiting and reduces the overhead involved in mutual exclusion.

However, the use of TSL also requires careful consideration. If multiple processes are contending for the same lock, the order in which they are granted access can affect system performance. Also, in systems without TSL, software solutions like Peterson's algorithm may be more portable and easier to implement.

In conclusion, the TSL instruction provides a powerful mechanism for achieving mutual exclusion in hardware, but its use requires careful consideration of its advantages and limitations. It is up to the operating system designer to weigh the trade-offs involved and choose the most appropriate mechanism for their system.

While the solutions we have discussed so far, such as Peterson's solution and the ones using TSL or XCHG instructions, are correct in ensuring mutual exclusion, they have the drawback of requiring busy waiting. This means that when a process is unable to enter its critical region, it repeatedly checks for availability in a tight loop, wasting CPU time and potentially causing unexpected effects.

6.5 Sleep and Wakeup

To overcome this drawback, we need to find a solution that allows the waiting process to give up the CPU and only resume its attempt to enter the critical region when it receives some sort of notification that the

resource is available. This notification can be in the form of an interrupt, a signal, or a message.

One way to implement this solution is to use a sleep/wakeup mechanism. When a process finds that the critical region is currently in use, it calls a sleep function that blocks the process until the resource is available. The process is then placed in a waiting queue, and the CPU is made available for other processes to run.

When the resource becomes available, the process holding the resource calls a wakeup function that signals the waiting processes that the resource is now available. The processes in the waiting queue are then made ready to run, and the operating system scheduler decides which process should be granted the CPU next.

This approach not only saves CPU time by avoiding busy waiting but also allows for better use of resources, as other processes can use the CPU while a process is waiting for a critical resource.

In conclusion, while solutions like Peterson's solution and those using TSL or XCHG instructions are effective in ensuring mutual exclusion, they have the drawback of requiring busy waiting. A sleep/wakeup mechanism can be used to overcome this drawback by allowing waiting processes to give up the CPU and only resume their attempt to enter the critical region when the resource is available, thus saving CPU time and enabling better resource utilization.

6.6 Methods of IPC and synchronization:

In a multi-tasking and concurrent environment, processes often need to communicate and synchronize with each other to achieve their intended goals. This is where Interprocess Communication (IPC) and synchronization techniques come into play. IPC and synchronization allow processes to exchange information and coordinate their activities, ensuring that the system operates correctly and efficiently.

There are several methods of IPC and synchronization, each with its strengths and weaknesses. In this chapter, we will explore four of the most commonly used methods: shared memory, message passing, semaphores, and monitors.

6.6.1 Semaphores

Semaphores are a synchronization technique that allows processes to coordinate their activities by controlling access to shared resources. A semaphore is essentially a counter that can be incremented and decremented by processes. When the counter reaches zero, the semaphore is considered to be locked, and any process attempting to access the shared resource must wait until the semaphore is unlocked.

Semaphores can be used to implement critical sections, where only one process at a time is allowed to access a shared resource. They can also be used to implement synchronization between processes, ensuring that one process completes its task before another process begins.

6.6.2 Mutexes

Mutexes, short for mutual exclusion objects, are a synchronization primitive that is widely used in operating systems to manage concurrent access to shared resources or pieces of code. They work in a similar way to semaphores but are simpler and more efficient to implement, which makes them popular in user-space thread packages.

A mutex is a shared variable that can be either in an unlocked or locked state. This state is indicated by a single bit or an integer, with zero indicating that the mutex is unlocked and any other value indicating that the mutex is locked. Two procedures are associated with mutexes: lock and unlock.

When a thread or process needs access to a critical region, it calls the lock procedure. If the mutex is currently unlocked (i.e., the critical region is available), the lock call succeeds, and the thread is allowed to

enter the critical region. If the mutex is already locked (i.e., the critical region is currently in use), the lock call blocks the thread until the mutex becomes unlocked again. This blocking mechanism prevents busy waiting, which can waste CPU time and degrade system performance.

Once a thread finishes executing in the critical region, it must call the unlock procedure to release the mutex and make the critical region available to other threads. The unlock procedure simply sets the mutex to an unlocked state, indicating that the critical region is available again.

Mutexes provide a simple and effective way of managing mutual exclusion and preventing race conditions in a concurrent environment. They are widely used in operating systems and are an essential component of many synchronization mechanisms.

6.6.3 Shared Memory

Shared memory is a technique that allows multiple processes to access the same region of memory. This region of memory is called a shared memory segment and is typically created by one process and then shared with other processes. Once a process has access to the shared memory segment, it can read from and write to it just like any other region of memory.

Shared memory is a fast and efficient way for processes to exchange large amounts of data because there is no need to copy the data between processes. However, it requires careful management to ensure that processes do not overwrite each other's data or access the shared memory segment at the same time.

6.6.4 Message Passing

Message passing is a technique that involves sending messages between processes. In this method, one process sends a message to another process, which receives and processes the message. Messages can be sent using either synchronous or asynchronous communication.

Synchronous communication means that the sender and receiver must synchronize their actions, such that the sender will not send another message until the receiver has processed the first message. Asynchronous communication, on the other hand, allows the sender to continue processing without waiting for the receiver to respond.

Message passing is a flexible and reliable method of IPC, but it can be slower and less efficient than shared memory, especially when large amounts of data need to be exchanged.

6.6.5 Monitors

Monitors are a synchronization technique that provides a higher-level abstraction than semaphores. A monitor is a module that encapsulates shared data and the procedures that operate on that data. Only one process can access a monitor at a time, and any other process that attempts to access the monitor is blocked until the first process completes its work.

Monitors are a powerful tool for IPC and synchronization because they simplify the development of concurrent programs. They provide a natural way to encapsulate shared data and procedures and ensure that processes do not interfere with each other.

IPC and synchronization are essential concepts in the field of operating systems. The methods we have discussed in this chapter provide ways for processes to communicate and coordinate their activities effectively, ensuring that the system operates correctly and efficiently. Shared memory, message passing, semaphores, and monitors all have their strengths and weaknesses, and the choice of which method to use depends on the specific requirements of the system. As such, it is important for operating system developers to have a solid understanding of these concepts and their implementation.

6.6.6 Barriers

Barriers are synchronization mechanisms designed for groups of processes rather than just two processes. They are commonly used in applications that are divided into phases, where no process is allowed to proceed to the next phase until all processes have completed the current one. Barriers are particularly useful for parallel computing, where multiple processes work together to solve a problem.

The basic idea behind barriers is to place a synchronization point at the end of each phase, which ensures that all processes have completed their work before moving on to the next phase. When a process reaches the barrier, it is blocked until all other processes have also reached the barrier. Once all processes have reached the barrier, they are released and can proceed to the next phase.

One common implementation of barriers is the "counting barrier", which works as follows: when a process reaches the barrier, it decrements a counter that keeps track of the number of processes that have reached the barrier. If the counter reaches zero, all processes have reached the barrier, and they are released. If the counter is not zero, the process is blocked until all other processes have reached the barrier.

Another type of barrier is the "sense-reversing barrier", which requires each process to maintain a "sense" variable that is toggled between true and false at each barrier. When a process reaches the barrier, it checks the sense variable of another process. If the sense variable is different from its own, the process is blocked until the other process has also reached the barrier. Once all processes have reached the barrier, they toggle their sense variables and release the waiting processes.

Barriers can be implemented using any synchronization mechanism, such as semaphores or mutexes, but they are often implemented using specialized barrier objects provided by the operating system or programming language. These objects typically provide higher-level

interfaces for creating and using barriers, making it easier for programmers to synchronize groups of processes.

In summary, barriers are a powerful synchronization mechanism that can be used to ensure that groups of processes complete their work in a coordinated and efficient manner. They are particularly useful in parallel computing and other applications that are divided into phases.

6.7 Importance of IPC and synchronization

In a multi-tasking and concurrent operating system, several processes run simultaneously, accessing and manipulating shared resources, such as memory, files, or hardware devices. To ensure correct and safe operation, these processes must communicate and synchronize with each other through Interprocess Communication (IPC) and synchronization mechanisms.

IPC allows processes to exchange information and coordinate their activities. This is crucial for tasks such as data sharing, interlocking, or coordination, which can be accomplished using different methods, such as shared memory, message passing, semaphores, or monitors.

Synchronization mechanisms, on the other hand, ensure that processes access shared resources in an orderly and safe manner. For instance, mutual exclusion mechanisms, like semaphores or monitors, prevent two processes from accessing the same resource simultaneously, thus avoiding race conditions and data inconsistencies. Similarly, synchronization mechanisms like barriers, locks, or condition variables, ensure that processes wait for each other until a specific condition is met, allowing them to coordinate their activities.

In conclusion, IPC and synchronization are critical components of multi-tasking and concurrent operating systems, as they ensure that processes can communicate and coordinate with each other in a safe and efficient manner. By using these mechanisms, processes can access

and manipulate shared resources while avoiding data inconsistencies, race conditions, or deadlocks.

7 Case Study: Process Management in Linux

In this chapter, we will explore the process management in Linux, one of the most popular operating systems in the world. We will begin by giving an overview of the Linux process management, including its design principles and features.

Next, we will compare Linux process management with other operating systems, such as Windows and MacOS. This will help us to understand the strengths and weaknesses of Linux process management and how it differs from other systems.

Finally, we will discuss the impact of process management on the performance, reliability, and functionality of the Linux Operating System. We will analyze the various techniques and strategies employed by the Linux process management system to achieve these goals.

7.1 Overview of Linux process management

Linux is one of the most widely used operating systems in the world, powering everything from servers to mobile devices. One of the key reasons for its success is the robust process management capabilities it provides. In this chapter, we will take a closer look at the process management features of Linux.

Processes in Linux are managed using the process table, which is a data structure that holds information about all currently running processes. Each process is identified by a unique process ID (PID) and is associated with other data, including its state, priority, parent process, and resource usage.

Linux provides a range of tools for managing processes, including the `top` command, which displays information about running processes, and the `kill` command, which is used to terminate a running process. The `ps` command is used to list processes and their attributes, while the `nice` and `renice` commands are used to adjust process priorities.

In Linux, processes are organized into a hierarchical structure, with each process having a parent process and the root process (`init`) as the ultimate parent. This structure helps to ensure that processes are properly managed and terminated when they are no longer needed.

Another key feature of Linux process management is the ability to create and manage threads. Threads are lightweight processes that share memory and other resources with their parent process. Linux provides a range of threading models, including POSIX threads, Native Posix threads library (NPTL), and LinuxThreads.

Overall, Linux's process management capabilities are a major strength of the operating system. They enable efficient multitasking and concurrency, ensuring that the system can handle multiple tasks simultaneously without becoming bogged down or crashing. Linux's process management features have also been influential in the development of other operating systems, making them an important area of study for anyone interested in operating system design and implementation.

7.2 Comparison with process management in other operating systems

Linux process management has several unique features and capabilities that set it apart from process management in other operating systems. In this chapter, we will compare Linux process management with process management in other popular operating systems, including Windows and macOS.

7.2.1 Windows Process Management:

In Windows, the process management system is similar to that of Linux, where each process has its own virtual address space. However, there are some notable differences between the two. For example, in Windows, processes are assigned a priority value, which determines the order in which they are executed. This priority value can be changed by the operating system or the user, depending on the needs of the system. Additionally, Windows uses a system of "job objects" to group processes together and apply policies such as CPU time limits, memory limits, and more.

7.2.2 macOS Process Management:

Like Linux and Windows, macOS also uses a similar process management system. However, there are some notable differences between the three. For example, macOS uses a concept called "launchd" to manage processes. Launchd is a daemon that manages system services, user applications, and other processes. It provides a single point of control for starting, stopping, and monitoring processes on the system. Additionally, macOS uses a system of "sandboxing" to limit the resources available to individual processes, providing an additional layer of security.

7.2.3 Linux Process Management:

Linux process management is highly flexible and customizable. Each process has its own virtual address space and is managed by the kernel. The Linux kernel provides a variety of scheduling algorithms, including the Completely Fair Scheduler (CFS) and the Round Robin Scheduler. Additionally, Linux supports a wide range of IPC mechanisms, including shared memory, message passing, and semaphores.

Linux also has a unique process management feature called "cgroups" (control groups). Cgroups allow processes to be organized into

hierarchical groups, with each group having its own set of resource limits (CPU, memory, etc.). This feature is particularly useful for managing large-scale deployments such as web servers, where it is essential to limit resource usage.

In conclusion, Linux process management provides a high degree of flexibility and customization, allowing it to be adapted to a wide range of use cases. While other operating systems have similar process management systems, Linux stands out for its support for cgroups, its variety of scheduling algorithms, and its range of IPC mechanisms.

7.3 Impact on Linux Operating System's performance, reliability, and functionality

Linux is known for its excellent process management system, which allows for efficient multi-tasking and concurrency. The process management system in Linux is responsible for creating, managing, and terminating processes, as well as allocating resources to these processes.

The impact of Linux's process management system on its performance is significant. The kernel's scheduler is designed to be efficient and can quickly switch between processes, allowing for smooth multi-tasking. This means that users can run multiple programs simultaneously without any noticeable lag or slowdown.

Moreover, Linux's process management system is designed with reliability in mind. It includes several mechanisms for ensuring that processes run smoothly and without interruption. For example, Linux uses signals to communicate with processes and notify them of events such as errors or resource availability.

In addition to performance and reliability, Linux's process management system also has an impact on the system's functionality. The system is designed to be flexible and can adapt to different requirements. For

example, it allows for the creation of real-time processes that require immediate attention and prioritization.

Furthermore, Linux's process management system supports various synchronization and interprocess communication mechanisms, including shared memory, message passing, semaphores, and monitors. These mechanisms enable processes to communicate and synchronize their activities efficiently, which is essential for multi-tasking and concurrency.

Overall, Linux's process management system is a critical component of the operating system, and its impact on performance, reliability, and functionality cannot be overstated. It is a testament to the power and flexibility of open-source software development and community-driven innovation.

8 Conclusion

In conclusion, processes are the fundamental building blocks of operating systems. They allow users to run multiple tasks concurrently and efficiently use the resources of a computer system. Process management includes various activities such as process creation, scheduling, synchronization, and communication. Operating systems use a range of scheduling algorithms to manage processes effectively, depending on the needs of the system and the tasks being performed. Interprocess communication and synchronization play a crucial role in maintaining the integrity of processes and preventing errors or conflicts in multi-tasking and concurrent environments.

As we have seen, Linux operating system provides a robust and efficient process management system, allowing users to run multiple processes simultaneously and effectively utilize system resources. The Linux process management system is superior to other operating systems in terms of scalability, flexibility, and reliability.

Overall, understanding processes and their management is essential for anyone interested in operating systems and computer science. By understanding how processes work, how they communicate with each other, and how they interact with system resources, we can build more efficient and reliable operating systems that meet the demands of modern computing.