FirstRanker.com

www.FirstRanker.com



Definition of an operating system:-An Operating System (OS) is an interface between a computer user and computer hardware. An operating system is a software which performs all the basic tasks like file management, memory management, process management, handling input and output, and controlling peripheral devices such as disk drives and printers. Some popular Operating Systems include UNIX, MS-DOS, MS-Windows - 98/XP/Vista, Windows-NT/2000, OS/2 and Mac OS..

It is software that works as an interface between a user and the computer hardware. The primary objective of an *operating system* is to make computer system convenient to use and to utilize computer hardware in an efficient manner. The operating system performs the basic tasks such as receiving input from the keyboard, processing instructions and sending output to the screen. It controls the execution of all kinds of programs.

A computer system can be roughly divided into four components.

- 1. Hardware
- 2. Operating system
- 3. Application programs
- 4. Users

The abstract view of the components of a computer system is shown in following figure.





Functions of an operating system:-Following are some of important functions of an operating System.

- 1. Memory Management
- 2. Processor Management
- 3. Device Management
- 4. File Management
- 5. Security
- 6. Control over system performance
- 7. Job accounting
- 8. Error detecting aids
- 9. Coordination between other software and users

1. Memory Management:-Memory management refers to management of Primary Memory or Main Memory. Main memory is a large array of words or bytes where each word or byte has its own address. Main memory provides a fast storage that can be accessed directly by the CPU. For a program to be executed, it must in the main memory. An Operating System does the following activities for memory management .

- Keeps tracks of primary memory, i.e., what part of it are in use by whom, what part is not in use.
- In multiprogramming, the OS decides which process will get memory when and how much.
- Allocates the memory when a process requests it to do so.
- De-allocates the memory when a process no longer needs it or has been terminated.

2. Processor Management:-In multiprogramming environment, the OS decides which process gets the processor when and for how much time. This function is called **process scheduling**. An Operating System does the following activities for processor management.

- Keeps tracks of processor and status of process. The program responsible for this task is known as **traffic controller**.
- Allocates the processor (CPU) to a process.
- De-allocates processor when a process is no longer required.

3. Device Management:-An Operating System manages device communication via their respective drivers. It does the following activities for device management.

- Keeps tracks of all devices. Program responsible for this task is known as the I/O controller.
- Decides which process gets the device when and for how much time.
- Allocates the device in the efficient way.
- De-allocates devices.

4. File Management:-A file system is normally organized into directories for easy navigation and usage. These directories may contain files and other directions. An Operating System does the following activities for file management.

- Keeps track of information, location, uses, status etc. The collective facilities are often known as **file system**.
- Decides who gets the resources.
- Allocates the resources.
- De-allocates the resources.

5. Security:-By means of password and similar other techniques, it prevents unauthorized access to programs and data.

6. Control over system performance:- Recording delays between request for a service and response from the system.

7. Job accounting:-Keeping track of time and resources used by various jobs and users.

8. Error detecting aids:-Production of dumps, traces, error messages, and other debugging and error detecting aids.

9. Coordination between other software's and users:-Coordination and assignment of compilers, interpreters, assemblers and other software to the various users of the computer systems.

TYPES OF OPERATING SYSTEMS; Operating systems are there from the very first computer generation and they keep evolving with time. In this chapter, we will discuss some of the important types of operating systems which are most commonly used.

1. **Batch operating system:**-The users of a batch operating system do not interact with the computer directly. Each user prepares his job on an off-line device like punch cards and submits it to the computer operator. To speed up processing, jobs with similar needs are batched together and run as a group. The programmers leave their programs with the operator and the operator then sorts the programs with similar requirements into batches. The problems with Batch Systems are as follows.

- Lack of interaction between the user and the job.
- CPU is often idle, because the speed of the mechanical I/O devices is slower than the CPU.
- Difficult to provide the desired priority.

FirstRanker.com

2. **Time-sharing systems**:-operating Time-sharing is a technique which enables many people, located at various terminals, to use a particular computer system at the same time. Time-sharing or multitasking is a logical extension of multiprogramming. Processor's time which is shared among multiple users simultaneously is termed as time-sharing.

The main difference between Multiprogrammed Batch Systems and Time-Sharing Systems is that in case of Multiprogrammed batch systems, the objective is to maximize processor use, whereas in Time-Sharing Systems, the objective is to minimize response time.

Multiple jobs are executed by the CPU by switching between them, but the switches occur so frequently. Thus, the user can receive an immediate response. For example, in a transaction processing, the processor executes each user program in a short burst or quantum of computation. That is, if \mathbf{n} users are present, then each user can get a time quantum. When the user submits the command, the response time is in few seconds at most.

The operating system uses CPU scheduling and multiprogramming to provide each user with a small portion of a time. Computer systems that were designed primarily as batch systems have been modified to time-sharing systems.

Advantages of Timesharing operating systems are as follows.

- Provides the advantage of quick response.
- Avoids duplication of software.
- Reduces CPU idle time.



Disadvantages of Time-sharing operating systems are as follows.

- Problem of reliability.
- Question of security and integrity of user programs and data.
- Problem of data communication.

3. Distributed operating System:-Distributed systems use multiple central processors to serve multiple real-time applications and multiple users. Data processing jobs are distributed among the processors accordingly.

The processors communicate with one another through various communication lines (such as high-speed buses or telephone lines). These are referred as **loosely coupled systems** or distributed systems. Processors in a distributed system may vary in size and function. These processors are referred as sites, nodes, computers, and so on.

The advantages of distributed systems are as follows.

- With resource sharing facility, a user at one site may be able to use the resources available at another.
- Speedup the exchange of data with one another via electronic mail.
- If one site fails in a distributed system, the remaining sites can potentially continue operating.
- Better service to the customers.
- Reduction of the load on the host computer.
- Reduction of delays in data processing.

4. Network operating System:-A Network Operating System runs on a server and provides the server the capability to manage data, users, groups, security, applications, and other networking functions. The primary purpose of the network operating system is to allow shared file and printer access among multiple computers in a network, typically a local area network (LAN), a private network or to other networks.

Examples of network operating systems include Microsoft Windows Server 2003, Microsoft Windows Server 2008, UNIX, Linux, Mac OS X, Novell NetWare, and BSD.

The advantages of network operating systems are as follows -

\ •

- Centralized servers are highly stable.
- Security is server managed.
- Upgrades to new technologies and hardware can be easily integrated into the system.
- Remote access to servers is possible from different locations and types of systems.

The disadvantages of network operating systems are as follows.

- High cost of buying and running a server.
- Dependency on a central location for most operations.
- Regular maintenance and updates are required.

5. Real Time operating System:-A real-time system is defined as a data processing system in which the time interval required to process and respond to inputs is so small that it controls the environment. The



time taken by the system to respond to an input and display of required updated information is termed as the response time. So in this method, the response time is very less as compared to online processing.

Real-time systems are used when there are rigid time requirements on the operation of a processor or the flow of data and real-time systems can be used as a control device in a dedicated application. A real-time operating system must have well-defined, fixed time constraints, otherwise the system will fail. For example, scientific experiments, medical image systems, industrial control systems, weapon systems, robots, air traffic control systems, etc.

There are two types of real-time operating systems.

1. Hard real-time systems:-Hard real-time systems guarantee that critical tasks complete on time. In hard real-time systems, secondary storage is limited or missing and the data is stored in ROM. In these systems, virtual memory is almost never found.

2. Soft real-time systems:-Soft real-time systems are less restrictive. A critical real-time task gets priority over other tasks and retains the priority until it completes. Soft real-time systems have limited utility than hard real-time systems. For example, multimedia virtual reality, Advanced Scientific Projects like undersea exploration and planetary rovers, etc.

OPERATING SYSTEM SERVICES:-An Operating System provides services to both the users and to the programs.

- It provides programs an environment to execute.
- It provides users the services to execute the programs in a convenient manner.

Following are a few common services provided by an operating system.

- 1. Program execution
- 2. I/O operations
- 3. File System manipulation
- 4. Communication
- 5. Error Detection
- 6. Resource Allocation
- 7. Protection

1. Program execution:-Operating systems handle many kinds of activities from user programs to system programs like printer spooler, name servers, file server, etc. Each of these activities is encapsulated as a process.

A process includes the complete execution context (code to execute, data to manipulate, registers, OS resources in use). Following are the major activities of an operating system with respect to program management –

- Loads a program into memory.
- Executes the program.
- Handles program's execution.
- Provides a mechanism for process synchronization.
- Provides a mechanism for process communication.

• Provides a mechanism for deadlock handling.

FirstRanker.com

2. I/O Operation:-An I/O subsystem comprises of I/O devices and their corresponding driver software. Drivers hide the peculiarities of specific hardware devices from the users.

An Operating System manages the communication between user and device drivers.

- I/O operation means read or write operation with any file or any specific I/O device.
- Operating system provides the access to the required I/O device when required.

3. File system manipulation:-A file represents a collection of related information. Computers can store files on the disk (secondary storage), for long-term storage purpose. Examples of storage media include magnetic tape, magnetic disk and optical disk drives like CD, DVD. Each of these media has its own properties like speed, capacity, data transfer rate and data access methods.

A file system is normally organized into directories for easy navigation and usage. These directories may contain files and other directions. Following are the major activities of an operating system with respect to file management -

- Program needs to read a file or write a file.
- The operating system gives the permission to the program for operation on file.
- Permission varies from read-only, read-write, denied and so on.
- Operating System provides an interface to the user to create/delete files.
- Operating System provides an interface to the user to create/delete directories.
- Operating System provides an interface to create the backup of file system.

4. Communication:-In case of distributed systems which are a collection of processors that do not share memory, peripheral devices, or a clock, the operating system manages communications between all the processes. Multiple processes communicate with one another through communication lines in the network.

The OS handles routing and connection strategies, and the problems of contention and security. Following are the major activities of an operating system with respect to communication –

- Two processes often require data to be transferred between them
- Both the processes can be on one computer or on different computers, but are connected through a computer network.
- Communication may be implemented by two methods, either by Shared Memory or by Message Passing.

5. Error handling:-Errors can occur anytime and anywhere. An error may occur in CPU, in I/O devices or in the memory hardware. Following are the major activities of an operating system with respect to error handling –

- The OS constantly checks for possible errors.
- The OS takes an appropriate action to ensure correct and consistent computing.

FirstRanker.com

www.FirstRanker.com

6. Resource Management:-In case of multi-user or multi-tasking environment, resources such as main memory, CPU cycles and files storage are to be allocated to each user or job. Following are the major activities of an operating system with respect to resource management –

- The OS manages all kinds of resources using schedulers.
- CPU scheduling algorithms are used for better utilization of CPU.

7. Protection:-Considering a computer system having multiple users and concurrent execution of multiple processes, the various processes must be protected from each other's activities.

Protection refers to a mechanism or a way to control the access of programs, processes, or users to the resources defined by a computer system. Following are the major activities of an operating system with respect to protection.

- The OS ensures that all access to system resources is controlled.
- The OS ensures that external I/O devices are protected from invalid access attempts.
- The OS provides authentication features for each user by means of passwords.

SYSTEM CALLS:-A **system call** is the programmatic way in which a computer program requests a service from the kernel of the operating system it is executed on. A system call is a way for programs to **interact with the operating system**. A computer program makes a system call when it makes a request to the operating system's kernel. System call **provides** the services of the operating system to the user programs via Application Program Interface (API). It provides an interface between a process and operating system to allow user-level processes to request services of the operating system. System calls are the only entry points into the kernel system. All programs needing resources must use system calls.

Services Provided by System Calls:

- 1. Process creation and management
- 2. Main memory management
- 3. File Access, Directory and File system management
- 4. Device handling(I/O)
- 5. Protection
- 6. Networking, etc.

Types of System Calls:- There are 5 different categories of system calls.

- 1. Process control: end, abort, create, terminate, allocate and free memory.
- 2. File management: create, open, close, delete, read file etc.
- 3. Device management
- 4. Information maintenance
- 5. Communication



www.FirstRanker.com

Examples of Windows and UNIX System Calls are

Types	Windows	Unix	
Process Control	CreateProcess() ExitProcess() WaitForSingleObject()	fork() exit	:() wait()
File Manipulation	CreateFile() ReadFile() WriteFile() CloseHandle()	open() write() clo	read() ose()
Device Manipulation	SetConsoleMode() ReadConsole() WriteConsole()	ioctl() write()	read()
Information Maintenance	GetCurrentProcessID() SetTimer() Sleeo()	getpid() sleep()	alarm()
Communication	CreatePipe() CreateFileMapping() MapViewOfFile()	pipe() mmap()	shmget()
Protection	SetFileSecurity()InitlializeSecurityDescriptor()SetSecurityDescriptorGroup()	chmod() chown()	umask()
	www.FirstRanker.com		



UNIT-II

Process concept: -A process is basically a program in execution. The execution of a process must progress in a sequential fashion. A process is defined as an entity which represents the basic unit of work to be implemented in the system.

To put it in simple terms, we write our computer programs in a text file and when we execute this program, it becomes a process which performs all the tasks mentioned in the program.

When a program is loaded into the memory and it becomes a process, it can be divided into four sections stack, heap, text and data. The following image shows a simplified layout of a process inside main memory.

Stack		
	ſ	
	Ļ	
Data		
Text		



Components and description of a process:-

1. Stack:-The process Stack contains the temporary data such as method/function parameters, return address and local variables.

2. Heap:-This is dynamically allocated memory to a process during its run time.

3. Text:-This includes the current activity represented by the value of Program Counter and the contents of the processor's registers.

4. Data:-This section contains the global and static variables.

Program:-A program is a piece of code which may be a single line or millions of lines. A computer program is usually written by a computer programmer in a programming language. For example, here is a simple program written in C programming language.

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

```
int main() {
    printf("Hello, World! \n");
    return 0;
}
```



A computer program is a collection of instructions that performs a specific task when executed by a computer. When we compare a program with a process, we can conclude that a process is a dynamic instance of a computer program.

A part of a computer program that performs a well-defined task is known as an **algorithm**. A collection of computer programs, libraries and related data are referred to as software.

Process Life Cycle (or) process state diagram:-When a process executes, it passes through different states. These stages may differ in different operating systems, and the names of these states are also not standardized. In general, a process can have one of the following five states at a time.



State & Description

1. Start:-This is the initial state when a process is first started/created.

2. Ready:-The process is waiting to be assigned to a processor. Ready processes are waiting to have the processor allocated to them by the operating system so that they can run. Process may come into this state after **Start** state or while running it by but interrupted by the scheduler to assign CPU to some other process.

3. Running:-Once the process has been assigned to a processor by the OS scheduler, the process state is set to running and the processor executes its instructions.

4. Waiting:-Process moves into the waiting state if it needs to wait for a resource, such as waiting for user input, or waiting for a file to become available.

5. Terminated or Exit:-Once the process finishes its execution, or it is terminated by the operating system, it is moved to the terminated state where it waits to be removed from main memory.

Process Control Block (PCB):-A Process Control Block is a data structure maintained by the Operating System for every process. The PCB is identified by an integer process ID (PID). A PCB keeps all the information needed to keep track of a process as listed below in the table. The architecture of a PCB is completely dependent on Operating System and may contain different information in different operating systems. Here is a simplified diagram of a PCB.





Information & Description:-

1. Process State:-The current state of the process i.e., whether it is ready, running, waiting, or whatever.

- 2. Process privileges:-This is required to allow/disallow access to system resources.
- 3. Process ID:-Unique identification for each of the process in the operating system.
- **4. Pointer:-**A pointer to parent process.

5. Program Counter:-Program Counter is a pointer to the address of the next instruction to be executed for this process.

6. CPU registers:-Various CPU registers where process need to be stored for execution for running state.

7. CPU Scheduling Information:-Process priority and other scheduling information which is required to schedule the process.

8. Memory management information:-This includes the information of page table, memory limits, Segment table depending on memory used by the operating system.

9.Accounting information:-This includes the amount of CPU used for process execution, time limits, execution ID etc.

10. IO status information:-This includes a list of I/O devices allocated to the process.

The architecture of a PCB is completely dependent on Operating System and may contain different information in different operating systems. Here is a simplified diagram of a PCB. The PCB is maintained for a process throughout its lifetime, and is deleted once the process terminates.



Process Scheduling

Definition:-The process scheduling is the activity of the process manager that handles the removal of the running process from the CPU and the selection of another process on the basis of a particular strategy.

Process scheduling is an essential part of a Multiprogramming operating systems. Such operating systems allow more than one process to be loaded into the executable memory at a time and the loaded process shares the CPU using time multiplexing.

Process Scheduling Queues:-The OS maintains all PCBs in Process Scheduling Queues. The OS maintains a separate queue for each of the process states and PCBs of all processes in the same execution state are placed in the same queue. When the state of a process is changed, its PCB is unlinked from its current queue and moved to its new state queue.

The Operating System maintains the following important process scheduling queues.

- **Job queue**: This queue keeps all the processes in the system.
- **Ready queue**:-This queue keeps a set of all processes residing in main memory, ready and waiting to execute. A new process is always put in this queue.
- Device queues: The processes which are blocked due to unavailability of an I/O device constitute this queue.



The OS can use different policies to manage each queue (FIFO, Round Robin, Priority, etc.). The OS scheduler determines how to move processes between the ready and run queues which can only have one entry per processor core on the system; in the above diagram, it has been merged with the CPU.



Scheduling Queues:-

- 1. All processes, upon entering into the system, are stored in the **Job Queue**.
- 2. Processes in the Ready state are placed in the Ready Queue.
- 3. Processes waiting for a device to become available are placed in **Device Queues**. There are unique device queues available for each I/O device.

A new process is initially put in the **Ready queue**. It waits in the ready queue until it is selected for execution (or dispatched). Once the process is assigned to the CPU and is executing, one of the following several events can occur:

- 1. The process could issue an I/O request, and then be placed in the I/O queue.
- 2. The process could create a new sub process and wait for its termination.
- 3. The process could be removed forcibly from the CPU, as a result of an interrupt, and be put back in the ready queue.



Process Schedulers:-Schedulers are special system software which handle process scheduling in various ways. Their main task is to select the jobs to be submitted into the system and to decide which process to run. Schedulers are of three types.

- 1. Long-Term Scheduler
- 2. Short-Term Scheduler
- 3. Medium-Term Scheduler

1. Long Term Scheduler:-It is also called a **job scheduler**. A long-term scheduler determines which programs are admitted to the system for processing. It selects processes from the queue and loads them into memory for execution. Process loads into the memory for CPU scheduling.



The primary objective of the job scheduler is to provide a balanced mix of jobs, such as I/O bound and processor bound. It also controls the degree of multiprogramming. If the degree of multiprogramming is stable, then the average rate of process creation must be equal to the average departure rate of processes leaving the system. On some systems, the long-term scheduler may not be available or minimal. Time-sharing operating systems have no long term scheduler. When a process changes the state from new to ready, then there is use of long-term scheduler.

2. Short Term Scheduler:-It is also called as CPU scheduler. Its main objective is to increase system performance in accordance with the chosen set of criteria. It is the change of ready state to running state of the process. CPU scheduler selects a process among the processes that are ready to execute and allocates CPU to one of them.Short-term schedulers, also known as dispatchers, make the decision of which process to execute next. Short-term schedulers are faster than long-term schedulers.

3. Medium Term Scheduler:-Medium-term scheduling is a part of **swapping**. It removes the processes from the memory. It reduces the degree of multiprogramming. The medium-term scheduler is in-charge of handling the swapped out-processes. A running process may become suspended if it makes an I/O request. Suspended processes cannot make any progress towards completion. In this condition, to remove the process from memory and make space for other processes, the suspended process is moved to the secondary storage. This process is called **swapping**, and the process is said to be swapped out or rolled out. Swapping may be necessary to improve the process mix.

Comparison among Scheduler

S.N. Long-Term Scheduler

1

2

Short-Term Scheduler

Medium-Term Scheduler

It is a job schedulerIt is a CPU schedulerIt is a process swapping scheduler.Speed is lesser than short term Speed is fastest among otherSpeed is in between both shortschedulertwoand long term scheduler.

ercorr

- 3 It controls the degree of It provides lesser control over It reduces the degree of multiprogramming degree of multiprogramming multiprogramming.
- 4 It is almost absent or minimal in It is also minimal in time It is a part of Time sharing system systems.
- It selects processes from pool and loads them into memory for execution
 It selects those processes which are ready to execute which are ready to execute be continued.



Context Switch: - A context switch is the mechanism to store and restore the state or context of a CPU in Process Control block so that a process execution can be resumed from the same point at a later time. Using this technique, a context switcher enables multiple processes to share a single CPU. Context switching is an essential part of a multitasking operating system features.

When the scheduler switches the CPU from executing one process to execute another, the state from the current running process is stored into the process control block. After this, the state for the process to run next is loaded from its own PCB and used to set the PC, registers, etc. At that point, the second process can start executing.



Context switches are computationally intensive since register and memory state must be saved and restored. To avoid the amount of context switching time, some hardware systems employ two or more sets of processor registers. When the process is switched, the following information is stored for later use.

- Program Counter
- Scheduling information
- Base and limit register value
- Currently used register
- Changed State
- I/O State information
- Accounting information



Scheduling algorithms

A Process Scheduler schedules different processes to be assigned to the CPU based on particular scheduling algorithms. There are six popular process scheduling algorithms.

- 1. First-Come, First-Served (FCFS) Scheduling
- 2. Shortest-Job-Next (SJN) Scheduling
- 3. Priority Scheduling
- 4. Shortest Remaining Time
- 5. Round Robin(RR) Scheduling
- 6. Multiple-Level Queues Scheduling

These algorithms are either **non-preemptive or preemptive**. Non-preemptive algorithms are designed so that once a process enters the running state; it cannot be preempted until it completes its allotted time, whereas the preemptive scheduling is based on priority where a scheduler may preempt a low priority running process anytime when a high priority process enters into a ready state.

1. First Come First Serve (FCFS):-

- Jobs are executed on first come, first serve basis.
- It is a non-preemptive, pre-emptive scheduling algorithm.
- Easy to understand and implement.
- Its implementation is based on FIFO queue.
- Poor in performance as average wait time is high.

Process Arrival Time Execute Time Service Time PO 0 5 0 P1 1 5 3 P2 2 8 8 3 P3 6 16

P	0	P1	P2	P3	
0	5	8		16	22



Wait time of each process is as follows.

Process	Wait Time : Service Time - Arrival Time
PO	0 - 0 = 0
P1	5 - 1 = 4
P2	8 - 2 = 6
Р3	16 - 3 = 13

Average Wait Time: (0+4+6+13) / 4 = 5.75

2. Shortest Job Next (SJN):-

- This is also known as **shortest job first**, or SJF
- This is a non-preemptive, pre-emptive scheduling algorithm.
- Best approach to minimize waiting time.
- Easy to implement in Batch systems where required CPU time is known in advance.
- Impossible to implement in interactive systems where required CPU time is not known.
- The processer should know in advance how much time process will take.

Process	Arrival Time	Execute Time	Service Time
PO	0	5	3
P1	1	3	0
P2	2	8	16
P3	3	6	8

P	1	PO	P3	P2	
0	3	8		16	22

Wait time of each process is as follows.

Wait Time : Service Time - Arrival Time
3 - 0 = 3
0 - 0 = 0
16 - 2 = 14
8 - 3 = 5

Average Wait Time: (3+0+14+5) / 4 = 5.50



3. Priority Based Scheduling:-

- Priority scheduling is a non-preemptive algorithm and one of the most common scheduling algorithms in batch systems.
- Each process is assigned a priority. Process with highest priority is to be executed first and so on.
- Processes with same priority are executed on first come first served basis. •
- Priority can be decided based on memory requirements, time requirements or any other • resource requirement.

Process	Arrival Time	Execute Time	Priority	Service Time
PO	0	5	1	9
P1	1	3	2	6
P2	2	8	1	14
P3	3	6	3	0

P	3	P1	PO	P2	
0	6		9	14	22

Wait time	of each process is as follows.
Process	Wait Time : Service Time - Arrival Time
P0	9 - 0 = 9
P1	6 - 1 = 5
P2	14 - 2 = 12
Р3	0 - 0 = 0

Average Wait Time: (9+5+12+0) / 4 = 6.5

4. Shortest Remaining Time:-

- Shortest remaining time (SRT) is the preemptive version of the SJN algorithm.
- The processor is allocated to the job closest to completion but it can be preempted by a newer ready job with shorter time to completion.
- Impossible to implement in interactive systems where required CPU time is not known.
- It is often used in batch environments where short jobs need to give preference. •



5. Round Robin Scheduling:-

- Round Robin is the preemptive process scheduling algorithm.
- Each process is provided a fix time to execute, it is called a **quantum**.
- Once a process is executed for a given time period, it is preempted and other process executes for a given time period.
- Context switching is used to save states of preempted processes.

Quantum = 3

F	20	P1	P2	P3	PO	P2	P3	P2
 0	3	6	9	12	14	17	20	22

Wait time of each process is as follows.

Process	Wait Time : Service Time - Arrival Time
P0	(0 - 0) + (12 - 3) = 9
P1	(3 - 1) = 2
P2	(6 - 2) + (14 - 9) + (20 - 17) = 12
Р3	(9 - 3) + (17 - 12) = 11

Average Wait Time: (9+2+12+11) / 4 = 8.5

6. Multiple-Level Queues Scheduling:-

Multiple-level queues are not an independent scheduling algorithm. They make use of other existing algorithms to group and schedule jobs with common characteristics.

- Multiple queues are maintained for processes with common characteristics.
- Each queue can have its own scheduling algorithms.
- Priorities are assigned to each queue.

For example, CPU-bound jobs can be scheduled in one queue and all I/O-bound jobs in another queue. The Process Scheduler then alternately selects jobs from each queue and assigns them to the CPU based on the algorithm assigned to the queue.



Multi-Threading

Thread:-A thread is a flow of execution through the process code, with its own program counter that keeps track of which instruction to execute next, system registers which hold its current working variables, and a stack which contains the execution history.

A thread shares with its peer threads few information like code segment, data segment and open files. When one thread alters a code segment memory item, all other threads see that.

A thread is also called a **lightweight process**. Threads provide a way to improve application performance through parallelism. Threads represent a software approach to improving performance of operating system by reducing the overhead thread is equivalent to a classical process.

Each thread belongs to exactly one process and no thread can exist outside a process. Each thread represents a separate flow of control. Threads have been successfully used in implementing network servers and web server. They also provide a suitable foundation for parallel execution of applications on shared memory multiprocessors. The following figure shows the working of a single-threaded and a multithreaded process.



Single Process P with single thread

Single Process P with three threads

Types of Thread:-Threads are implemented in following two ways.

- 1. User Level Threads: User managed threads.
- 2. **Kernel Level Threads**:-Operating System managed threads acting on kernel, an operating system core.

User Level Threads:-In this case, the thread management kernel is not aware of the existence of threads. The thread library contains code for creating and destroying threads, for passing message and



data between threads, for scheduling thread execution and for saving and restoring thread contexts. The application starts with a single thread.

ſ	Process P1	Process P2	
User Space	Thread Library	Thread Library	
Karaal Caraa	ļ		
Kernel Space			

Advantages

- Thread switching does not require Kernel mode privileges.
- User level thread can run on any operating system.
- Scheduling can be application specific in the user level thread.
- User level threads are fast to create and manage.

Disadvantages

- In a typical operating system, most system calls are blocking.
- Multithreaded application cannot take advantage of multiprocessing.

Kernel Level Threads:-In this case, thread management is done by the Kernel. There is no thread management code in the application area. Kernel threads are supported directly by the operating system. Any application can be programmed to be multithreaded. All of the threads within an application are supported within a single process.

The Kernel maintains context information for the process as a whole and for individuals threads within the process. Scheduling by the Kernel is done on a thread basis. The Kernel performs thread creation, scheduling and management in Kernel space. Kernel threads are generally slower to create and manage than the user threads.



Advantages

- Kernel can simultaneously schedule multiple threads from the same process on multiple processes.
- If one thread in a process is blocked, the Kernel can schedule another thread of the same process.
- Kernel routines themselves can be multithreaded.

Disadvantages

- Kernel threads are generally slower to create and manage than the user threads.
- Transfer of control from one thread to another within the same process requires a mode switch to the Kernel.

Multithreading Models:-Some operating system provides a combined user level thread and Kernel level thread facility. Solaris is a good example of this combined approach. In a combined system, multiple threads within the same application can run in parallel on multiple processors and a blocking system call need not block the entire process. Multithreading models are three types

- 1. Many to many relationship.
- 2. Many to one relationship.
- 3. One to one relationship.

1. Many to Many Models:-The many-to-many model multiplexes any number of user threads onto an equal or smaller number of kernel threads.

The following diagram shows the many-to-many threading model where 6 user level threads are multiplexing with 6 kernel level threads. In this model, developers can create as many user threads as necessary and the corresponding Kernel threads can run in parallel on a multiprocessor machine. This model provides the best accuracy on concurrency and when a thread performs a blocking system call, the kernel can schedule another thread for execution.



FirstRanker.com

www.FirstRanker.com

2. Many to One Model:-Many-to-one model maps many user level threads to one Kernel-level thread. Thread management is done in user space by the thread library. When thread makes a blocking system call, the entire process will be blocked. Only one thread can access the Kernel at a time, so multiple threads are unable to run in parallel on multiprocessors.

If the user-level thread libraries are implemented in the operating system in such a way that the system does not support them, then the Kernel threads use the many-to-one relationship modes.



4.One to One Model:-There is one-to-one relationship of user-level thread to the kernel-level thread. This model provides more concurrency than the many-to-one model. It also allows another thread to run when a thread makes a blocking system call. It supports multiple threads to execute in parallel on microprocessors. Disadvantage of this model is that creating user thread requires the corresponding Kernel thread. OS/2, windows NT and windows 2000 use one to one relationship model.



Difference between User-Level & Kernel-Level Thread



s.no	User-Level Threads	Kernel-Level Thread
1	User-level threads are faster to create and manage.	Kernel-level threads are slower to create and manage.
2	Implementation is by a thread library at the user level.	Operating system supports creation of Kernel threads.
3	User-level thread is generic and can run on any operating system.	Kernel-level thread is specific to the operating system.
4	Multi-threaded applications cannot take advantage of multiprocessing.	Kernel routines themselves can be multithreaded.

Interprocess Communication

- Processes within a system may be independent or cooperating
- Cooperating process can affect or be affected by other processes, including sharing data •
- Reasons for cooperating processes: •
- Information sharing
- **Computation speedup**
- Modularity •
- Convenience
- Cooperating processes need inter process communication (IPC) •
- Two models of IPC
- Shared memory •
- Message passing .

Communications Models





Cooperating Processes

Independent process cannot affect or be affected by the execution of another process •

• Cooperating process can affect or be affected by the execution of another process Advantages of process cooperation

- Information sharing •
- Computation speed-up •
- Modularity
- Convenience •

Producer-Consumer Problem

- Paradigm for cooperating processes, producer process produces information that is consumed • by a consumer process
- unbounded-buffer places no practical limit on the size of the buffer •
- bounded-buffer assumes that there is a fixed buffer size

Bounded-Buffer – Shared-Memory Solution

Shared data

#define BUFFER_SIZE 10

typedef struct {

. . .

} item;

item buffer[BUFFER_SIZE];

int in = 0;

int out = 0;

FirstRanker.com Solution is correct, but can only use BUFFER_SIZE-1 elements

Bounded-Buffer – Producer

while (true) {

/* Produce an item */

while (((in = (in + 1) % BUFFER SIZE count) == out)

; /* do nothing -- no free buffers */

buffer[in] = item;



www.FirstRanker.com

in = (in + 1) % BUFFER SIZE;

}

Bounded Buffer – Consumer

while (true) {

while (in == out)

; // do nothing -- nothing to consume

// remove an item from the buffer

item = buffer[out];

out = (out + 1) % BUFFER SIZE;

return item;

}

Interprocess Communication – Message Passing

- Mechanism for processes to communicate and to synchronize their actions
- Message system processes communicate with each other without resorting to shared variables
- IPC facility provides two operations:
- send(message) message size fixed or variable
- **receive**(*message*)
- If *P* and *Q* wish to communicate, they need to:
- establish a communication link between them
- exchange messages via send/receive
- Implementation of communication link
- physical (e.g., shared memory, hardware bus)
- logical (e.g., logical properties)

Direct Communication

- Processes must name each other explicitly:
- send (P, message) send a message to process P
- receive(Q, message) receive a message from process Q
- Properties of communication link
- Links are established automatically
- A link is associated with exactly one pair of communicating processes
- Between each pair there exists exactly one link
- The link may be unidirectional, but is usually bi-directional



Indirect Communication

- Messages are directed and received from mailboxes (also referred to as ports)
- Each mailbox has a unique id
- Processes can communicate only if they share a mailbox
- Properties of communication link
- Link established only if processes share a common mailbox
- A link may be associated with many processes
- Each pair of processes may share several communication links
- Link may be unidirectional or bi-directional
- Operations
- create a new mailbox
- send and receive messages through mailbox
- destroy a mailbox
- Primitives are defined as:
- **send**(*A*, *message*) send a message to mailbox A
- receive(A, message) receive a message from mailbox A
- Mailbox sharing
- P1, P2, and P3 share mailbox A
- P1, sends; P2 and P3 receive
- Who gets the message?
- Solutions
- Allow a link to be associated with at most two processes
- Allow only one process at a time to execute a receive operation
- Allow the system to select arbitrarily the receiver. Sender is notified who the receiver was.

Synchronization

- Message passing may be either blocking or non-blocking
- Blocking is considered synchronous
- Blocking send has the sender block until the message is received
- Blocking receive has the receiver block until a message is available
- Non-blocking is considered asynchronous
- Non-blocking send has the sender send the message and continue
- Non-blocking receive has the receiver receive a valid message or null

Buffering

Queue of messages attached to the link; implemented in one of three ways

1. Zero capacity – 0 messages

Sender must wait for receiver (rendezvous)

2. Bounded capacity – finite length of *n* messages

Sender must wait if link full



3. Unbounded capacity – infinite length Sender never waits

Examples of IPC systems

- 1. Posix : uses shared memory method.
- 2. Mach : uses message passing
- 3. Windows XP : uses message passing using local procedural calls

<u>Scheduling</u> Criteria:-There are many different criteria's to check when considering the "best" scheduling algorithm, they are:

1. *CPU Utilization:* - To make out the best use of CPU and not to waste any CPU cycle, CPU would be working most of the time (Ideally 100% of the time). Considering a real system, CPU usage should range from 40% (lightly loaded) to 90% (heavily loaded.)

2. Throughput:-It is the total number of processes completed per unit time or rather say total amount of work done in a unit of time. This may range from 10/second to 1/hour depending on the specific processes.

3. Turnaround Time:-It is the amount of time taken to execute a particular process, i.e. The interval from time of submission of the process to the time of completion of the process (Wall clock time).

4. Waiting Time:-The sum of the periods spent waiting in the ready queue amount of time a process has been waiting in the ready queue to acquire get control on the CPU.

5. Load Average:-It is the average number of processes residing in the ready queue waiting for their turn to get into the CPU.

6. Response Time:- Amount of time it takes from when a request was submitted until the first response is produced. Remember, it is the time till the first response and not the completion of process execution (final response).

In general CPU utilization and Throughput are maximized and other factors are reduced for proper optimization.

Operations on Process:-Below we have discussed the two major operations Process

Creation and Process Termination. 1.Process Creation:-

- **Parent** process create **children** processes, which, in turn create other processes, forming a tree of processes
- Generally, process identified and managed via a process identifier (pid)
- Resource sharing
- Parent and children share all resources
- Children share subset of parent's resources
- Parent and child share no resources
- Execution
- Parent and children execute concurrently
- Parent waits until children terminate
- Address space



- Child duplicate of parent
- Child has a program loaded into it
- UNIX examples
- fork system call creates new process
- exec system call used after a fork to replace the process' memory space with a new program

```
Process Creation:-
```



C Program Forking Separate Process

```
int main()
```

```
{
```

```
pid_t pid;
    /* fork another process */
    pid = fork();
    if (pid < 0) { /* error occurred */
        fprintf(stderr, "Fork Failed");
        exit(-1);
    }
    else if (pid == 0) { /* child process */
        execlp("/bin/ls", "ls", NULL);
    }
    else { /* parent process */
        /* parent will wait for the child to complete */
        wait (NULL);
        printf ("Child Complete");
```





A tree of processes on a typical Solaris



2. Process Termination:-

- Process executes last statement and asks the operating system to delete it (exit)
- Output data from child to parent (via wait)
- Process' resources are deallocated by operating system
- Parent may terminate execution of children processes (abort)
- Child has exceeded allocated resources
- Task assigned to child is no longer required
- If parent is exiting
 - Some operating system do not allow child to continue if its parent terminates

All children terminated - cascading termination



Multithreading Issues: The following are few issues related to multithreading.

1. *Thread Cancellation:-Thread cancellation means terminating a thread before it has finished working. There can be two approaches for this; one is Asynchronous cancellation, which terminates the target thread immediately. The other is Deferred cancellation allows the target thread to periodically check if it should be cancelled.*

2. Signal Handling:-Signals are used in UNIX systems to notify a process that a particular event has occurred. Now in when a multithreaded process receives a signal, to which thread it must be delivered? It can be delivered to all, or a single thread.

3.fork() System Call:-fork() is a system call executed in the kernel through which a process creates a copy of itself. Now the problem in Multithreaded process is, if one thread forks, will the entire process be copied or not?

4. Security Issues:-Yes, there can be security issues because of extensive sharing of resources between multiple threads.

There are many other issues that you might face in a multithreaded process, but there are appropriate solutions available for them. Pointing out some issues here was just to study both sides of the coin.

www.FirstRanker.com



Unit-III

Memory Management

- To provide a detailed description of various ways of organizing memory hardware
- To discuss various memory-management techniques, including paging and segmentation
- To provide a detailed description of the Intel Pentium, which supports both pure segmentation and segmentation with paging
- Program must be brought (from disk) into memory and placed within a process for it to be run
- Main memory and registers are only storage CPU can access directly
- Register access in one CPU clock (or less)
- Main memory can take many cycles
- Cache sits between main memory and CPU registers
- Protection of memory required to ensure correct operation

Base and Limit Registers:-



Binding of Instructions and Data to Memory

- Address binding of instructions and data to memory addresses can happen at three different stages
- **Compile time**: If memory location known a priori, **absolute code** can be generated; must recompile code if starting location changes
- Load time: Must generate relocatable code if memory location is not known at compile time
- **Execution time**: Binding delayed until run time if the process can be moved during its execution from one memory segment to another. Need hardware support for address maps (e.g., base and limit registers)



Multistep Processing of a User Program



Logical vs. Physical Address Space

- The concept of a logical address space that is bound to a separate **physical address space** is central to proper memory management
- Logical address generated by the CPU; also referred to as virtual address
- Physical address address seen by the memory unit
- Logical and physical addresses are the same in compile-time and load-time address-binding schemes; logical (virtual) and physical addresses differ in execution-time address-binding scheme

Memory-Management Unit (MMU)

- Hardware device that maps virtual to physical address
- In MMU scheme, the value in the relocation register is added to every address generated by a user process at the time it is sent to memory
- The user program deals with *logical* addresses; it never sees the *real* physical addresses **Dynamic relocation using a relocation register**





Dynamic Loading

- Routine is not loaded until it is called
- Better memory-space utilization; unused routine is never loaded
- Useful when large amounts of code are needed to handle infrequently occurring cases
- No special support from the operating system is required implemented through program design

Dynamic Linking

- Linking postponed until execution time
- Small piece of code, *stub*, used to locate the appropriate memory-resident library routine
- Stub replaces itself with the address of the routine, and executes the routine
- Operating system needed to check if routine is in processes' memory address
- Dynamic linking is particularly useful for libraries
- System also known as shared libraries

Swapping:-

- A process can be swapped temporarily out of memory to a backing store, and then brought back into memory for continued execution .
- **Backing store** fast disk large enough to accommodate copies of all memory images for all users; must provide direct access to these memory images
- Roll out, roll in swapping variant used for priority-based scheduling algorithms; lower-priority process is swapped out so higher-priority process can be loaded and executed
- Major part of swap time is transfer time; total transfer time is directly proportional to the amount of memory swapped
- Modified versions of swapping are found on many systems (i.e., UNIX, Linux, and Windows)

• System maintains a **ready queue** of ready-to-run processes which have memory images on disk **Schematic View of Swapping**





Contiguous memory Allocation:-

- Main memory usually into two partitions:
- Resident operating system, usually held in low memory with interrupt vector
- User processes then held in high memory
- Relocation registers used to protect user processes from each other, and from changing operating-system code and data
- Base register contains value of smallest physical address
- Limit register contains range of logical addresses each logical address must be less than the limit register
- MMU maps logical address dynamically

Hardware Support for Relocation and Limit Registers



- Multiple-partition allocation
- Hole block of available memory; holes of various size are scattered throughout memory
- When a process arrives, it is allocated memory from a hole large enough to accommodate it
- Operating system maintains information about:
- a) allocated partitions b) free partitions (hole)



www.FirstRanker.com



- First-fit: Allocate the first hole that is big enough
- **Best-fit**: Allocate the *smallest* hole that is big enough; must search entire list, unless ordered by size Produces the smallest leftover hole
- Worst-fit: Allocate the *largest* hole; must also search entire list
- Produces the largest leftover hole
- First-fit and best-fit better than worst-fit in terms of speed and storage utilization

Fragmentation:-

- External Fragmentation total memory space exists to satisfy a request, but it is not contiguous
- Internal Fragmentation allocated memory may be slightly larger than requested memory; this size difference is memory internal to a partition, but not being used
- Reduce external fragmentation by compaction
- Shuffle memory contents to place all free memory together in one large block
- Compaction is possible *only* if relocation is dynamic, and is done at execution time.
- I/O problem

ILatch job in memory while it is involved in I/O

Do I/O only into OS buffers

PAGING

- Logical address space of a process can be noncontiguous; process is allocated physical memory whenever the latter is available
- Divide physical memory into fixed-sized blocks called **frames** (size is power of 2, between 512 bytes and 8,192 bytes)
- Divide logical memory into blocks of same size called pages N Keep track of all free frames
- To run a program of size *n* pages, need to find *n* free frames and load program
- Set up a page table to translate logical to physical addresses
- Internal fragmentation

Address Translation Scheme:-

- Address generated by CPU is divided into
- **Page number (***p***)** used as an index into a *page table* which contains base address of each page in physical memory
- Page offset (d) combined with base address to define the physical memory address that is sent to the memory unit
- For given logical address space 2^m and page size 2ⁿ

Paging Hardware








32-byte memory and 4-byte pages

Free Frames



Implementation of Page Table

- Page table is kept in main memory
- Page-table base register (PTBR) points to the page table
- Page-table length register (PRLR) indicates size of the page table
- In this scheme every data/instruction access requires two memory accesses. One for the page table and one for the data/instruction.
- The two memory access problem can be solved by the use of a special fast-lookup hardware cache called **associative memory** or **translation look-aside buffers (TLBs)**
- Some TLBs store **address-space identifiers (ASIDs)** in each TLB entry uniquely identifies each process to provide address-space protection for that process

Associative Memory

- Associative memory parallel search
- Address translation (p, d)
- If p is in associative register, get frame # out
- Otherwise get frame # from page table in memory

Page #	Frame #



Paging Hardware With TLB



- Memory protection implemented by associating protection bit with each frame
- Valid-invalid bit attached to each entry in the page table:
- "valid" indicates that the associated page is in the process' logical address space, and is thus a legal page

= 2 + e – a

- "invalid" indicates that the page is not in the process' logical address space
- Valid (v) or Invalid (i) Bit In A Page Table



www.FirstRanker.com



Shared code

- One copy of read-only (reentrant) code shared among processes (i.e., text editors, compilers, window systems).
- Shared code must appear in same location in the logical address space of all processes •

Private code and data

- Each process keeps a separate copy of the code and data •
- The pages for the private code and data can appear anywhere in the logical address space •

Shared Pages Example



- Break up the logical address space into multiple page tables





Two-Level Paging Example

- A logical address (on 32-bit machine with 1K page size) is divided into:
- a page number consisting of 22 bits
- a page offset consisting of 10 bits
- Since the page table is paged, the page number is further divided into:
- a 12-bit page number
- a 10-bit page offset
- Thus, a logical address is as follows:

where p_i is an index into the outer page table, and p_2 is the displacement within the page of the outer page table





2.Hashed Page Tables:-

- Common in address spaces > 32 bits
- The virtual page number is hashed into a page table
- This page table contains a chain of elements hashing to the same location
- Virtual page numbers are compared in this chain searching for a match
- If a match is found, the corresponding physical frame is extracted

Hashed Page Table



3. Inverted Page 1

- One entry for each real page of memory
- Entry consists of the virtual address of the page stored in that real memory location, with information about the process that owns that page
- Decreases memory needed to store each page table, but increases time needed to search the table when a page reference occurs
- Use hash table to limit the search to one or at most a few page-table entries

Inverted Page Table Architecture





Segmentation

Memory-management scheme that supports user view of memory

- A program is a collection of segments
- A segment is a logical unit such as:
- main program
- procedure function
- method
- object
- local variables, global variables
- common block
- stack
- symbol table
- arrays



User's View of a Program





Segmentation Architecture

- Logical address consists of a two tuple:
 - <segment-number, offset>,
- Segment table maps two-dimensional physical addresses; each table entry has:
- **base** contains the starting physical address where the segments reside in memory
- **limit** specifies the length of the segment
- Segment-table base register (STBR) points to the segment table's location in memory
- Segment-table length register (STLR) indicates number of segments used by a program; segment number *s* is legal if *s* < STLR
 - Protection
 - With each entry in segment table associate:
 - Ivalidation bit = 0 P illegal segment
 - Pread/write/execute privileges
 - Protection bits associated with segments; code sharing occurs at segment level
 - Since segments vary in length, memory allocation is a dynamic storage-allocation problem
 - A segmentation example is shown in the following diagram



www.FirstRanker.com



Example: The Intel Pentium

- Supports both segmentation and segmentation with paging
- CPU generates logical address
- Given to segmentation unit 2Which produces linear addresses
- Linear address given to paging unit
 Which generates physical address in main memory

Paging units form equivalent of MMU

Logical to Physical Address Translation in Pentium





Linear Address in Linux

global	middle	page	offset
directory	directory	table	

Three-level Paging in Linux



<u>Virtual memory:</u>-A computer can address more memory than the amount physically installed on the system. This extra memory is actually called **virtual memory** and it is a section of a hard disk that's set up to emulate the computer's RAM.

The main visible advantage of this scheme is that programs can be larger than physical memory. Virtual memory serves two purposes. First, it allows us to extend the use of physical memory by using disk. Second, it allows us to have memory protection, because each virtual address is translated to a physical address.

Following are the situations, when entire program is not required to be loaded fully in main memory.

- User written error handling routines are used only when an error occurred in the data or computation.
- Certain options and features of a program may be used rarely.
- Many tables are assigned a fixed amount of address space even though only a small amount of the table is actually used.
- The ability to execute a program that is only partially in memory would counter many benefits.
- Less number of I/O would be needed to load or swap each user program into memory.
- A program would no longer be constrained by the amount of physical memory that is available.
- Each user program could take less physical memory, more programs could be run the same time, with a corresponding increase in CPU utilization and throughput.

Modern microprocessors intended for general-purpose use, a memory management unit, or MMU, is built into the hardware. The MMU's job is to translate virtual addresses into physical addresses. A basic example is given below.



Virtual memory is commonly implemented by demand paging. It can also be implemented in a segmentation system. Demand segmentation can also be used to provide virtual memory.

Demand Paging:-A demand paging system is quite similar to a paging system with swapping where processes reside in secondary memory and pages are loaded only on demand, not in advance. When a context switch occurs, the operating system does not copy any of the old program's pages out to the disk or any of the new program's pages into the main memory Instead, it just begins executing the new program after loading the first page and fetches that program's pages as they are referenced.



While executing a program, if the program references a page which is not available in the main memory because it was swapped out a little ago, the processor treats this invalid memory reference as a **page fault** and transfers control from the program to the operating system to demand the page back into the memory.

Advantages:-Following are the advantages of Demand Paging.

- Large virtual memory.
- More efficient use of memory.
- There is no limit on degree of multiprogramming.

Disadvantages:-

• Number of tables and the amount of processor overhead for handling page interrupts are greater than in the case of the simple paged management techniques.

Page Replacement Algorithm

Page replacement algorithms are the techniques using which an Operating System decides which memory pages to swap out, write to disk when a page of memory needs to be allocated. Paging happens whenever a page fault occurs and a free page cannot be used for allocation purpose accounting to reason that pages are not available or the number of free pages is lower than required pages.

When the page that was selected for replacement and was paged out, is referenced again, it has to read in from disk, and this requires for I/O completion. This process determines the quality of the page replacement algorithm: the lesser the time waiting for page-ins, the better is the algorithm.

A page replacement algorithm looks at the limited information about accessing the pages provided by hardware, and tries to select which pages should be replaced to minimize the total number of page misses, while balancing it with the costs of primary storage and processor time of the algorithm itself. There are many different page replacement algorithms. We evaluate an algorithm by running it on a particular string of memory reference and computing the number of page faults,

Reference String:-

The string of memory references is called reference string. Reference strings are generated artificially or by tracing a given system and recording the address of each memory reference. The latter choice produces a large number of data, where we note two things.

- For a given page size, we need to consider only the page number, not the entire address.
- If we have a reference to a page **p**, then any immediately following references to page **p** will never cause a page fault. Page p will be in memory after the first reference; the immediately following references will not fault.
- For example, consider the following sequence of addresses 123,215,600,1234,76,96
- If page size is 100, then the reference string is 1,2,6,12,0,0

1. First In First Out (FIFO) algorithm:



- Oldest page in main memory is the one which will be selected for replacement.
- Easy to implement, keep a list, replace pages from the tail and add new pages at the head.

Reference String: 0, 2, 1, 6, 4, 0, 1, 0, 3, 1, 2, 1

```
Misses
```



Fault Rate = 9 / 12 = 0.75

- 2. Optimal Page algorithm:-
 - An optimal page-replacement algorithm has the lowest page-fault rate of all algorithms. An optimal page-replacement algorithm exists, and has been called OPT or MIN.
 - Replace the page that will not be used for the longest period of time. Use the time when a page is to be used.

Reference String: 0, 2, 1, 6, 4, 0, 1, 0, 3, 1, 2, 1

Misses :x x x x x x



Fault Rate = 6 / 12 = 0.50



3. Least Recently Used (LRU) algorithm:-

- Page which has not been used for the longest time in main memory is the one which will be • selected for replacement.
- Easy to implement, keep a list, replace pages by looking back into time.

Reference String: 0, 2, 1, 6, 4, 0, 1, 0, 3, 1, 2, 1 Misses :xx x x x x x x x



Fault Rate = 8 / 12 = 0.67

4. Page Buffering algorithm:-

- anker.co • To get a process start quickly, keep a pool of free frames.
- On page fault, select a page to be replaced.
- Write the new page in the frame of free pool, mark the page table and restart the process.
- Now write the dirty page out of disk and place the frame holding replaced page in free pool. •

5. Least frequently Used (LFU) algorithm:-

- The page with the smallest count is the one which will be selected for replacement.
- This algorithm suffers from the situation in which a page is used heavily during the initial phase of a process, but then is never used again.

6. Most frequently Used (MFU) algorithm:-

This algorithm is based on the argument that the page with the smallest count was probably just • brought in and has yet to be used.



Thrashing: A process that is spending more time paging than executing is said to be thrashing. In other words it means, that the process doesn't have enough frames to hold all the pages for its execution, so it is swapping pages in and out very frequently to keep executing. Sometimes, the pages which will be required in the near future have to be swapped out.

Initially when the CPU utilization is low, the process scheduling mechanism, to increase the level of multiprogramming loads multiple processes into the memory at the same time, allocating a limited amount of frames to each process. As the memory fills up, process starts to spend a lot of time for the required pages to be swapped in, again leading to low CPU utilization because most of the processes are waiting for pages. Hence the scheduler loads more processes to increase CPU utilization, as this continues at a point of time the complete system comes to a stop



To prevent thrashing we must provide processes with as many frames as they really need right now.

Scanned by CamScanner



www.FirstRanker.com

UNIT-IV

SYLLABUS

Concurrency: Process Synchronization, The Critical- Section Problem, Synchronization Hardware, Semaphores, Classic Problems of Synchronization, Monitors, Synchronization examples **Principles of deadlock** – System Model, Deadlock Characterization, Deadlock Prevention, Detection and Avoidance, Recovery form Deadlock

Process Synchronization

Process Synchronization means sharing system resources by processes in a such a way that, Concurrent access to shared data is handled thereby minimizing the chance of inconsistent data. Maintaining data consistency demands mechanisms to ensure synchronized execution of cooperating processes.

Critical Section Problem

A Critical Section is a code segment that accesses shared variables and has to be executed as an atomic action. It means that in a group of cooperating processes, at a given point of time, only one process must be executing its critical section. If any other process also wants to execute its critical section, it must wait until the first one finishes.



Solution to Critical Section Problem

A solution to the critical section problem must satisfy the following three conditions:

1. Mutual Exclusion

Out of a group of cooperating processes, only one process can be in its critical section at a given point of time.

2. Progress

If no process is in its critical section, and if one or more threads want to execute their critical section then any one of these threads must be allowed to get into its critical section.



3. Bounded Waiting

After a process makes a request for getting into its critical section, there is a limit for how many other processes can get into their critical section, before this process's request is granted. So after the limit is reached, system must grant the process permission to get into its critical section.

Synchronization Hardware

Many systems provide hardware support for critical section code. The critical section problem could be solved easily in a single-processor environment if we could disallow interrupts to occur while a shared variable or resource is being modified.

In this manner, we could be sure that the current sequence of instructions would be allowed to execute in order without pre-emption. Unfortunately, this solution is not feasible in a multiprocessor environment.

Disabling interrupt on a multiprocessor environment can be time consuming as the message is passed to all the processors.

This message transmission lag, delays entry of threads into critical section and the system efficiency decreases.

Introduction to Semaphores

In 1965, Dijkstra proposed a new and very significant technique for managing concurrent processes by using the value of a simple integer variable to synchronize the progress of interacting processes. This integer variable is called **semaphore**. So it is basically a synchronizing tool and is accessed only through two low standard atomic operations, **wait** and **signal** designated by P(S) and V(S) respectively.

In very simple words, **semaphore** is a variable which can hold only a non-negative Integer value, shared between all the threads, with operations **wait** and **signal**, which work as follow:



The classical definitions of wait and signal are:

Wait: Decrements the value of its argument s, as soon as it would become non-

negative(greater than or equal to 1).

• **Signal**: Increments the value of its argument **s**, as there is no more process blocked on the queue.



Properties of Semaphores

- 1. It's simple and always have a non-negative Integer value.
- 2. Works with many processes.
- 3. Can have many different critical sections with different semaphores.
- 4. Each critical section has unique access semaphores.
- 5. Can permit multiple processes into the critical section at once, if desirable.

Types of Semaphores

Semaphores are mainly of two types:

1. Binary Semaphore:

It is a special form of semaphore used for implementing mutual exclusion, hence it is often called a **Mutex**. A binary semaphore is initialized to **1** and only takes the values **0** and **1** during execution of a program.

2. <u>Counting Semaphores:</u>

These are used to implement bounded concurrency.

Example of Use

Here is a simple step wise implementation involving declaration and usage of semaphore.



Limitations of Semaphores

- 1. **Priority Inversion** is a big limitation of semaphores.
- 2. Their use is not enforced, but is by convention only.
- 3. With improper use, a process may block indefinitely.

Classical Problems of Synchronization:



Classical problems used to test newly-proposed synchronization schemes \circ

Bounded-Buffer Problem

- $\circ~$ Readers and Writers Problem
- o Dining-Philosophers Problem

Bounded-Buffer Problem

N buffers, each can hold one item

Semaphore mutex initialized to the value 1

Semaphore full initialized to the value 0

Semaphore empty initialized to the value N

The structure of the producer process Ranker.com

do {

produce an item in nextp

wait (empty);

wait (mutex);

add the item to the buffer

signal (mutex);

signal (full);

} while (TRUE);

The structure of the consumer process

do {

www.FirstRanker.com



www.FirstRanker.com

wait (full);

wait (mutex);

// remove an item from buffer to nextc

signal (mutex);

signal (empty);

consume the item in nextc

} while (TRUE);

Readers-Writers Problem:

A data set is shared among a number of concurrent processes

- Readers only read the data set; they do **not** perform any updates
- Writers can both read and write

Problem – allow multiple readers to read at the same time

Only one single writer can access the shared data at the same time

Several variations of how readers and writers are treated - all involve priorities

Shared Data

Data set

Semaphore mutex initialized to 1

Semaphore wrt initialized to 1



www.FirstRanker.com

```
The structure of a writer process
```

do {

wait (wrt);

writing is performed

signal (wrt);

} while (TRUE);

The structure of a reader process			
do {			
wait (mutex);			
readcount ++ ;			
if (readcount == 1)			
wait (wrt);			
signal (mutex)			
reading is performed			
wait (mutex); readcount			
;			
if (readcount $== 0$)			



www.FirstRanker.com

signal (wrt);

signal (mutex);

} while (TRUE);

Dining-Philosophers Problem

Philosophers spend their lives thinking and eating

Don't interact with their neighbors, occasionally try to pick up 2 chopsticks (one at a time) to eat from bowl

Need both to eat, then release both when done

In the case of 5 philosophers

Shared data

Bowl of rice (data set) net com [5] initialized to 1

Semaphore chopstick [5] initialized to 1

The structure of Philosopher i

do {

it (chopstick[i]);

wait (chopStick[(i + 1) % 5]);

// eat

signal (chopstick[i]);

signal (chopstick[(i + 1) % 5]);



think

} while (TRUE);

Monitors

A high-level abstraction that provides a convenient and effective mechanism for process synchronization

Abstract data type, internal variables only accessible by code within the procedure

Only one process may be active within the monitor at a time

But not powerful enough to model some synchronization schemes

monitor monitor-name

{

arations arations

}

Classical Problems of Synchronization

Semaphore can be used in other synchronization problems besides Mutual Exclusion.

Below are some of the classical problem depicting flaws of process synchronaization in systems where cooperating processes are present.

We will discuss the following three problems:



- 2. The Readers Writers Problem
- 3. Dining Philosophers Problem

Bounded Buffer Problem

• This problem is generalised in terms of the **Producer Consumer problem**, where a **finite** buffer pool is used to exchange messages between producer and consumer processes.

Because the buffer pool has a maximum size, this problem is often called the **Bounded buffer problem**.

• Solution to this problem is, creating two counting semaphores "full" and "empty" to keep track of the current number of full and empty buffers respectively.

The Readers Writers Problem

• In this problem there are some processes(called **readers**) that only read the shared data, and never change it, and there are other processes(called **writers**) who may change the data in addition to reading, or instead of reading it.

0

• There are various type of readers-writers problem, most centred on relative priorities of readers and writers.

Dining Philosophers Problem

- The dining philosopher's problem involves the allocation of limited resources to a group of processes in a deadlock-free and starvation-free manner.
- There are five philosophers sitting around a table, in which there are five chopsticks/forks kept beside them and a bowl of rice in the centre, When a philosopher wants to eat, he uses two chopsticks one from their left and one from their right. When a philosopher wants to think, he keeps down both chopsticks at their original place.

PRINCIPLES OF DEAD LOCKS:-



What is a Deadlock?

Deadlocks are a set of blocked processes each holding a resource and waiting to acquire a resource held by another process.



How to avoid Deadlocks

Deadlocks can be avoided by avoiding at least one of the four conditions, because all this four conditions are required simultaneously to cause deadlock.

Q

1. Mutual Exclusion

Resources shared such as read-only files do not lead to deadlocks but resources, such as printers and tape drives, requires exclusive access by a single process.

2. Hold and Wait

In this condition processes must be prevented from holding one or more resources while simultaneously waiting for one or more others.

3. No Preemption

Preemption of process resource allocations can avoid the condition of deadlocks, where ever possible.

4. Circular Wait

Circular wait can be avoided if we number all resources, and require that processes request

resources only in strictly increasing(or decreasing) order.

Λ.



Handling Deadlock

The above points focus on preventing deadlocks. But what to do once a deadlock has occured. Following three strategies can be used to remove deadlock after its occurrence.

1. **Preemption**

We can take a resource from one process and give it to other. This will resolve the deadlock situation, but sometimes it does causes problems.

2. Rollback

In situations where deadlock is a real possibility, the system can periodically make a record of the state of each process and when deadlock occurs, roll everything back to the last checkpoint, and restart, but allocating resources differently so that deadlock does not occur.

3. Kill one or more processes

This is the simplest way, but it works.

Deadlock avoidance

As you saw already, most prevention algorithms have poor resource utilization, and hence result in reduced throughputs. Instead, we can try to avoid deadlocks by making use prior knowledge about the usage of resources by processes including resources available, resources allocated, future requests and future releases by processes. Most deadlock avoidance algorithms need every process to tell in advance the maximum number of resources of each type that it may need. Based on all these info we may decide if a process should wait for a resource or not, and thus avoid chances for circular wait.

If a system is already in a safe state, we can try to stay away from an unsafe state and avoid deadlock. Deadlocks cannot be avoided in an unsafe state. A system can be considered to be in safe state if it is not in a state of deadlock and can allocate resources upto the maximum available. A safe sequence of processes and allocation of resources ensures a safe state. Deadlock avoidance algorithms try not to allocate resources to a process if it will make the system in an unsafe state. Since resource allocation is not done **right away in some cases**, deadlock avoidance algorithms also suffer from low resource utilization problem.

A resource allocation graph is generally used to avoid deadlocks. If there are no cycles in the resource allocation graph, then there are no deadlocks. If there are cycles, there may be a deadlock. If there is only one instance of every resource, then a cycle implies a deadlock. Vertices of the resource allocation graph are resources and processes. The resource allocation graph has request edges and assignment edges. An edge from a process to resource is a request edge and an edge from a resource to process is an allocation edge. A calm edge denotes that a request may be made in future and is represented as a dashed line. Based on calm edges we can see if there is a chance for a cycle and then grant requests if the system will again be in a safe state.

www.FirstRanker.com



Consider the image with calm edges as below:



If R2 is allocated to p2 and if P1 request for R2, there will be a deadlock.



The resource allocation graph is not much useful if there are multiple instances for a resource. In such a case, we can use Banker's algorithm. In this algorithm, every process must tell upfront the maximum resource of each type it need, subject to the maximum available instances for each type. Allocation of resources is made only, if the allocation ensures a safe state; else the processes need to wait. The Banker's algorithm can be divided into two parts: Safety algorithm if a system is in a safe state or not. The resource request algorithm make an assumption of allocation and see if the system will be in a safe state. If the new state is unsafe, the resources are not allocated and the data structures are restored to their previous state; in this case the processes must wait for the resource.

Deadlock Detection

If deadlock prevention and avoidance are not done properly, as deadlock may occur and only things left to do is to detect the recover from the deadlock.

If all resource types has only single instance, then we can use a graph called wait-for-graph, which is a variant of resource allocation graph. Here, vertices represent processes and a directed edge from P1 to P2 indicate that P1 is waiting for a resource held by P2. Like in the case of resource allocation graph, a cycle in a wait-for-graph indicate a deadlock. So the system can maintain a wait-for-graph and check for cycles periodically to detect any deadlocks.



www.FirstRanker.com



The wait-for-graph is not much useful if there are multiple instances for a resource, as a cycle may not imply a deadlock. In such a case, we can use an algorithm similar to Banker's algorithm to detect deadlock. We can see if further allocations can be made on not based on current allocations. *You can refer to any operating system text books for details of these algorithms*.

Deadlock Recovery

Once a deadlock is detected, you will have to break the deadlock. It can be done through different ways, including, aborting one or more processes to break the circular wait condition causing the deadlock and preempting resources from one or more processes which are deadlocked.

Deadlock Characterization

Deadlock can arise if four conditions hold simultaneously.

Mutual exclusion: only one process at a time can use a resource

- Hold and wait: a process holding at least one resource is waiting to acquire additional resources held by other processes
- **No preemption:** a resource can be released only voluntarily by the process holding it, after that process has completed its task
- **Circular wait:** there exists a set $\{P_0, P_1, ..., P_n\}$ of waiting processes such that P_0 is waiting for a resource that is held by P_1, P_1 is waiting for a resource that is held by

 P_2, \ldots, P_{n-1} is waiting for a resource that is held by P_n , and P_n is waiting for a resource that is held by P_0 .

Deadlock Prevention

Restrain the ways request can be made



- **utual Exclusion** not required for sharable resources; must hold for nonsharable resources
- **Hold and Wait** must guarantee that whenever a process requests a resource, it does not hold any other resources

Require process to request and be allocated all its resources before it begins execution, or allow process to request resources only when the process has none

Low resource utilization; starvation possible

No Preemption -

If a process that is holding some resources requests another resource that cannot be immediately allocated to it, then all resources currently being held are released

Preempted resources are added to the list of resources for which the process is waiting

Process will be restarted only when it can regain its old resources, as well as the new ones that it is requesting

Circular Wait – impose a total ordering of all resource types, and require that each process requests resources in an increasing order of enumeration

Deadlock Avoidance

Requires that the system has some additional a priori information available

Simplest and most useful model requires that each process declare the *maximum number* of resources of each type that it may need

The deadlock-avoidance algorithm dynamically examines the resource-allocation state to ensure that there can never be a circular-wait condition

Resource-allocation *state* is defined by the number of available and allocated resources, and the maximum demands of the processes

Safe State

When a process requests an available resource, system must decide if immediate allocation leaves the system in a safe state



System is in safe state if there exists a sequence $\langle P_1, P_2, ..., P_n \rangle$ of ALL the processes in the systems such that for each P_i, the resources that P_i can still request can be satisfied by currently available resources + resources held by all the P_j , with j < I

That is:

If P_i resource needs are not immediately available, then P_i can wait until all P_ihave finished

When P_i is finished, P_i can obtain needed resources, execute, return allocated resources, and terminate

then P_i terminates, P_{i+1} can obtain its needed resources, and so on

If a system is in safe state no deadlocks

If a system is in unsafe state possibility of deadlock

Avoidance ensure that a system will never enter an unsafe state

Avoidance algorithms

Single instance of a resource type

Her.con Use a resource-allocation grap

Multiple instances of a resource type

Use the banker's algorithm

Resource-Allocation Graph Scheme

Claim edge $P_i \not\models R_j$ indicated that process P_j may request resource R_j ; represented by a dashed line

Claim edge converts to request edge when a process requests a resource

Request edge converted to an assignment edge when the resource is allocated to the process

When a resource is released by a process, assignment edge reconverts to a claim edge





Banker's Algorithm

Multiple instances

Each process must a priori claim maximum use

When a process requests a resource it may have to wait

When a process gets all its resources it must return them in a finite amount of time Let n = number of processes, and m = number of resources types.

- Available: Vector of length *m*. If available [j] = k, there are *k* instances of resource type R_j available
- **Max**: *n* x *m* matrix. If Max[i,j] = k, then process P_i may request at most k instances of resource type R_j
- Allocation: $n \ge m$ matrix. If Allocation[i,j] = k then P_i is currently allocated k instances of R_j
- **Need**: $n \ge m$ matrix. If Need[i,j] = k, then P_i may need k more instances of R_j to complete its task

Need
$$[i,j] = Max[i,j] - Allocation [i,j]$$

safety Algorithm



www.FirstRanker.com

Let Work and Finish be vectors of length m and n, respectively.

Initialize: Work = Available

Finish [i] = false for $i = 0, 1, \dots, n-1$

Find an isuch that both:

Finish [i] = false

Need_i£Work

If no such iexists, go to step 4

Work = Work + Allocation_i Finish[i] = true go to step 2

4.IfFinish [i] == true for all i, then the system is in a safe state **Resource-Request Algorithm for Process** P_i

Request = request vector for process P_i . If *Request*_i[j] = k then process P_i wants k instances of resource type R_j

If $Request_i$ £Need_igo to step 2. Otherwise, raise error condition, since process has exceeded its maximum claim

If $Request_i$ $\pounds Available$, go to step 3. Otherwise P_i must wait, since resources are not available

Pretend to allocate requested resources to P_i by modifying the state as follows:

Available = Available – Request; Allocation_i= Allocation_i + Request_i; Need_i=Need_i – Request_i;

• If safe the resources are allocated to Pi



www.FirstRanker.com

Example of Banker's Algorithm

```
5 processes P_0 through P_4;
```

3 resource types:

A (10 instances), B (5instances), and C (7 instances)

Snapshot at time T_0 :

	ABC	ABC	ABC
P_0	010	753	332
<i>P</i> ₁	200	322	
P_2	302	902	
P 3	211	222	-
P_4	002	433	col

The content of the matrix *Need* is defined to be *Max – Allocation*

0



The system is in a safe state since the sequence $\langle P_1, P_3, P_4, P_2, P_0 \rangle$ satisfies safety criteria

*P*₁ Request (1,0,2)



Allocation	Need	Available		
		A B C	A B C A B C	
	P_0	010	743	230
	<i>P</i> ₁ 3 0) 2	020	
	<i>P</i> ₂	302	600	
	<i>P</i> ₃	211	011	
	P_4	002	431	

Executing safety algorithm shows that sequence $\langle P_1, P_3, P_4, P_0, P_2 \rangle$ satisfies safety :Ranker.com

Recovery from Deadlock:

Process Termination

Abort all deadlocked processes

Abort one process at a time until the deadlock cycle is eliminated

In which order should we choose to abort?

Priority of the process

How long process has computed, and how much longer to completion

Resources the process has used

Resources process needs to complete

How many processes will need to be terminated

Is process interactive or batch?

FirstRanker.com

www.FirstRanker.com

www.FirstRanker.com

Selecting a victim – minimize cost

Rollback - return to some safe state, restart process for that state

Starvation – same process may always be picked as victim, include number of rollback in cost factor

UNIT-V

SYLLABUS

File system Interface- the concept of a file, Access Methods, Directory structure, File system mounting, file sharing, protection.

File System implementation- File system structure, allocation methods, free-space management

Mass-storage structure overview of Mass-storage structure, Disk scheduling, Device drivers.

The concept of a file

File

A file is a named collection of related information that is recorded on secondary storage such as magnetic disks, magnetic tapes and optical disks. In general, a file is a sequence of bits, bytes, lines or records whose meaning is defined by the files creator and user.

File Structure

A File Structure should be according to a required format that the operating system can understand.

- A file has a certain defined structure according to its type.
- A text file is a sequence of characters organized into lines.
- A source file is a sequence of procedures and functions.
- An object file is a sequence of bytes organized into blocks that are understandable by the machine.
- When operating system defines different file structures, it also contains the code to support these file structure. Unix, MS-DOS support minimum number of file structure.

File Type

File type refers to the ability of the operating system to distinguish different types of file such as text files source files and binary files etc. Many operating systems support many types of files. Operating system like MS-DOS and UNIX have the following types of files –

Ordinary files

- These are the files that conmmuterFirstRanker.com
- These may have text, databases or executable program.


The user can apply various operations on such files like add, modify, delete or even remove the entire file.

Directory files

These files contain list of file names and other information related to these files.

Special files

- These files are also known as device files.
- These files represent physical device like disks, terminals, printers, networks, tape drive etc.

These files are of two types –

- Character special files data is handled character by character as in case of terminals or printers.
- **Block special files** data is handled in blocks as in the case of disks and tapes.

File Access Methods

File access mechanism refers to the manner in which the records of a file may be accessed. There are several ways to access files – anker.com

- Sequential access
- Direct/Random access
- Indexed sequential access

Sequential access

A sequential access is that in which the records are accessed in some sequence, i.e., the information in the file is processed in order, one record after the other. This access method is the most primitive one. Example: Compilers usually access files in this fashion.

Direct/Random access

- Random access file organization provides, accessing the records directly.
- Each record has its own address on the file with by the help of which it can be directly accessed for reading or writing.
- The records need not be in any sequence within the file and they need not be in adjacent locations on the storage medium.

Indexed sequential access

- This mechanism is built up on base of sequential access.
- An index is created for each file which contains pointers to various blocks.
- Index is searched sequentially and its pointer is used to access the file directly.

Space Allocation

Files are allocated disk space wwwo first Ranken comperating systems deploy following three main ways to allocate disk space to files.



- Contiguous Allocation
- Linked Allocation
- Indexed Allocation

Contiguous Allocation

- Each file occupies a contiguous address space on disk.
- Assigned disk address is in linear order.
- Easy to implement.
- External fragmentation is a major issue with this type of allocation technique.

Linked Allocation

- Each file carries a list of links to disk blocks.
- Directory contains link / pointer to first block of a file.
- No external fragmentation
- Effectively used in sequential access file.
- Inefficient in case of direct access file.

Indexed Allocation

- Provides solutions to problems of contiguous and linked allocation.
- A index block is created having all pointers to files.
- Each file has its own index block which stores the addresses of disk space occupied by the file.
- Directory contains the addresses of index blocks of files.

Directory structure

Directory is a symbol table of files that stores all the related information about the file it hold with the contents. Directory is a list of files. Each entry of a directory define a file information like a file name, type, its version number, size ,owner of file, access rights, date of creation and date of last backup.



Unit-VI

Components of Linux System

Linux Operating System has primarily three components

- **Kernel** Kernel is the core part of Linux. It is responsible for all major activities of this operating system. It consists of various modules and it interacts directly with the underlying hardware. Kernel provides the required abstraction to hide low level hardware details to system or application programs.
- **System Library** System libraries are special functions or programs using which application programs or system utilities accesses Kernel's features. These libraries implement most of the functionalities of the operating system and do not requires kernel module's code access rights.
- **System Utility** System Utility programs are responsible to do specialized, individual level tasks.

System Softwares	User Process	User Utility	Compilers
	System	Libraries	
	Kei	rnel	
	Kernel	Modules	
			=
	4		
ardware	CPU	RAM	1/0

Kernel Mode vs User Mode

Kernel component code executes in a special privileged mode called **kernel mode** with full access to all resources of the computer. This code represents a single process, executes in single address space and do not require any context switch and hence is very efficient and fast. Kernel runs each processes and provides system services to processes, provides protected access to hardware to processes.

Support code which is not required to run in kernel mode is in System Library. User programs and other system programs works in **User Mode** which has no access to system hardware and kernel code. User programs/ utilities use System libraries to access Kernel functions to get system's low level tasks.

www.FirstRanker.com



Basic Features

Following are some of the important features of Linux Operating System.

- **Portable** Portability means software can works on different types of hardware in same way. Linux kernel and application programs supports their installation on any kind of hardware platform.
- **Open Source** Linux source code is freely available and it is community based development project. Multiple teams work in collaboration to enhance the capability of Linux operating system and it is continuously evolving.
- **Multi-User** Linux is a multiuser system means multiple users can access system resources like memory/ ram/ application programs at same time.
- **Multiprogramming** Linux is a multiprogramming system means multiple applications can run at same time.
- **Hierarchical File System** Linux provides a standard file structure in which system files/ user files are arranged.
- **Shell** Linux provides a special interpreter program which can be used to execute commands of the operating system. It can be used to do various types of operations, call application programs. etc.
- Security Linux provides user security using authentication features like password protection/ controlled access to specific files/ encryption of data.

Architecture

The following illustration shows the architecture of a Linux system -





The architecture of a Linux System consists of the following layers -

- **Hardware layer** Hardware consists of all peripheral devices (RAM/ HDD/ CPU etc).
- **Kernel** It is the core component of Operating System, interacts directly with hardware, provides low level services to upper layer components.
- Shell An interface to kernel, hiding complexity of kernel's functions from users. The shell takes commands from the user and executes kernel's functions.
- Utilities Utility programs that provide the user most of the functionalities of an operating systems.

Interprocess Communication

Processes communicate with each other and with the kernel to coordinate their activities. Linux supports a number of Inter-Process Communication (IPC) mechanisms. Signals and pipes are two of them but Linux also supports the System V IPC mechanisms named after the Unix release in which they first appeared.

1 Signals:-

Signals are one of the oldest inter-process communication methods used by Unix TM systems. They are used to signal asynchronous events to one or more processes. A signal could be generated by a keyboard interrupt or an error condition such as the process attempting to access a non-existent location in its virtual memory. Signals are also used by the shells to signal job control commands to their child processes.

There are a set of defined signals that the kernel can generate or that can be generated by other processes in the system, provided that they have the correct privileges. You can list a system's set of signals using the kill command (kill -l), on my Intel Linux box this gives:

```
SIGHUP 2) SIGINT 3) SIGQUIT 4) SIGILL
SIGTRAP 6) SIGIOT 7) SIGBUS 8) SIGFPE
SIGKILL10) SIGUSR111) SIGSEGV12) SIGUSR2
SIGPIPE14) SIGALRM15) SIGTERM17) SIGCHLD
SIGCONT19) SIGSTOP20) SIGTSTP21) SIGTTIN
SIGTTOU23) SIGURG24) SIGXCPU25) SIGXFSZ
SIGVTALRM27) SIGPROF28) SIGWINCH29) SIGIO
SIGPWR
```

The numbers are different for an Alpha AXP Linux box. Processes can choose to ignore most of the signals that are generated, with two notable exceptions: neither the SIGSTOP signal which causes a process to halt its execution nor the SIGKILL signal which causes a process to exit can be ignored. Otherwise though, a process can choose just how it wants to handle the various signals. Processes can block the signals and, if they do not block them, they can either choose to handle them themselves or allow the kernel to handle them. If the kernel handles the signals, it will do the default actions required for this signal. For example, the default action when a process receives the SIGFPE (floating point exception) signal is to core dump and then exit. Signals have no inherent relative priorities. If two signals are generated for a process at the same time then they may be presented to the process or handled in any order. Also there is no mechanism for handling multiple signals of the same kind. There is no way that a process can tell if it received 1 or 42 SIGCONT signals.

Linux implements signals using information stored in the task_struct for the process. The number of supported signals is limited to the Word size of the processor. Processes with a word size of 32 bits can have 32 signals whereas 64 bit processors like the Alpha AXP may



have up to 64 signals. The currently pending signals are kept in the signal field with a mask of blocked signals held in blocked. With the exception of SIGSTOP and SIGKILL, all signals can be blocked. If a blocked signal is generated, it remains pending until it is unblocked. Linux also holds information about how each process handles every possible signal and this is held in an array of sigaction data structures pointed at by the task_struct for each process. Amongst other things it contains either the address of a routine that will handle the signal or a flag which tells Linux that the process either wishes to ignore this signal or let the kernel handle the signal for it. The process modifies the default signal handling by making system calls and these calls alter the signation for the appropriate signal as well as the blocked mask.

Not every process in the system can send signals to every other process, the kernel can and super users can. Normal processes can only send signals to processes with the same *uid* and gid or to processes in the same process group. Signals are generated by setting the appropriate bit in the task struct's signal field. If the process has not blocked the signal and is waiting but interruptible (in state Interruptible) then it is woken up by changing its state to Running and making sure that it is in the run queue. That way the scheduler will consider it a candidate for running when the system next schedules. If the default handling is needed, then Linux can optimize the handling of the signal. For example if the signal SIGWINCH (the X window changed focus) and the default handler is being used then there is nothing to be done.

Signals are not presented to the process immediately they are generated, they must wait until the process is running again. Every time a process exits from a system call its signal and blocked fields are checked and, if there are any unblocked signals, they can now be delivered. This might seem a very unreliable method but every process in the system is making system calls, for example to write a character to the terminal, all of the time. Processes can elect to wait for signals if they wish, they are suspended in state Interruptible until a signal is presented. The Linux signal processing code looks at the signation structure for each of the current unblocked signals.

If a signal's handler is set to the default action then the kernel will handle it. The SIGSTOP signal's default handler will change the current process's state to Stopped and then run the scheduler to select a new process to run. The default action for the SIGFPE signal will core dump the process and then cause it to exit. Alternatively, the process may have specified its own signal handler. This is a routine which will be called whenever the signal is generated and the sigaction structure holds the address of this routine. The kernel must call the process's signal handling routine and how this happens is processor specific but all CPUs must cope with the fact that the current process is running in kernel mode and is just about to return to the process that called the kernel or system routine in user mode. The problem is solved by manipulating the stack and registers of the process. The process's program counter is set to the address of its signal handling routine and the parameters to the routine are added to the call frame or passed in registers. When the process resumes operation it appears as if the signal handling routine were called normally.

Linux is POSIX compatible and so the process can specify which signals are blocked when a particular signal handling routine is called. This means changing the blocked mask during the call to the processes signal handler. The blocked mask must be returned to its original value when the signal handling routine has finished. Therefore Linux adds a call to a tidy up routine which will restore the original blocked mask onto the call stack of the signalled process. Linux also optimizes the case where several signal handling routines need to be called by stacking them so that each time one handling routine exits, the next one is called until the tidy up routine is called. www.FirstRanker.com



www.FirstRanker.com

2.Pipes:-

The common Linux shells all allow redirection. For example

\$ ls | pr | lpr

pipes the output from the ls command listing the directory's files into the standard input of the pr command which paginates them. Finally the standard output from the pr command is piped into the standard input of the lpr command which prints the results on the default printer. Pipes then are unidirectional byte streams which connect the standard output from one process into the standard input of another process. Neither process is aware of this redirection and behaves just as it would normally. It is the shell which sets up these temporary pipes between the processes.



Figure 5.1: Pipes

In Linux, a pipe is implemented using two file data structures which both point at the same temporary VFS inode which itself points at a physical page within memory. Figure <u>5.1</u> shows that each file data structure contains pointers to different file operation routine vectors; one for writing to the pipe, the other for reading from the pipe.

This hides the underlying differences from the generic system calls which read and write to ordinary files. As the writing process writes to the pipe, bytes are copied into the shared data page and when the reading process reads from the pipe, bytes are copied from the shared data page. Linux must synchronize access to the pipe. It must make sure that the reader and the writer of the pipe are in step and to do this it preselects on the shared by the state of the pipe.



When the writer wants to write to the pipe it uses the standard write library functions. These all pass file descriptors that are indices into the process's set of file data structures, each one representing an open file or, as in this case, an open pipe. The Linux system call uses the write routine pointed at by the file data structure describing this pipe. That write routine uses information held in the VFS inode representing the pipe to manage the write request.

If there is enough room to write all of the bytes into the pipe and, so long as the pipe is not locked by its reader, Linux locks it for the writer and copies the bytes to be written from the process's address space into the shared data page. If the pipe is locked by the reader or if there is not enough room for the data then the current process is made to sleep on the pipe inode's wait queue and the scheduler is called so that another process can run. It is interruptible, so it can receive signals and it will be woken by the reader when there is enough room for the write data or when the pipe is unlocked. When the data has been written, the pipe's VFS inode is unlocked and any waiting readers sleeping on the inode's wait queue will themselves be woken up.

Reading data from the pipe is a very similar process to writing to it.

Processes are allowed to do non-blocking reads (it depends on the mode in which they opened the file or pipe) and, in this case, if there is no data to be read or if the pipe is locked, an error will be returned. This means that the process can continue to run. The alternative is to wait on the pipe inode's wait queue until the write process has finished. When both processes have finished with the pipe, the pipe inode is discarded along with the shared data page.

Linux also supports named pipes, also known as FIFOs because pipes operate on a First In, First Out principle. The first data written into the pipe is the first data read from the pipe. Unlike pipes, FIFOs are not temporary objects, they are entities in the file system and can be created using the mkfifo command. Processes are free to use a FIFO so long as they have appropriate access rights to it. The way that FIFOs are opened is a little different from pipes. A pipe (its two file data structures, its VFS inode and the shared data page) is created in one go whereas a FIFO already exists and is opened and closed by its users. Linux must handle readers opening the FIFO before writers open it as well as readers reading before any writers have written to it. That aside, FIFOs are handled almost exactly the same way as pipes and they use the same data structures and operations.

Synchronization

In a shared memory application, developers must ensure that shared resources are protected from concurrent access. The kernel is no exception. Shared resources require protection from concurrent access because if multiple threads of execution access and manipulate the data at the same time, the threads may overwrite each other's changes or access data while it is in an inconsistent state. Concurrent access of shared data often results in instability is hard to track down and debug.

The term *threads of execution* implies any instance of executing code. For example, this includes any of the following:

- A task in the kernel
- An interrupt handler
- A bottom half
- A kernel thread

www.FirstRanker.com This chapter may shorten *threads of execution* to simply *threads*. Keep in mind that this term describes any executing code.

Symmetrical multiprocessing support was introduced in the 2.0 kernel. Multiprocessing support implies that kernel code can simultaneously run on two or more processors. Consequently, without protection, code in the kernel, running on two different processors, can simultaneously access shared data at exactly the same time. With the introduction of the 2.6 kernel, the Linux kernel is preemptive. This implies that (in the absence of protection) the scheduler can preempt kernel code at virtually any point and reschedule another task. Today, a number of scenarios enable for concurrency inside the kernel, and they all require protection.

This chapter discusses the issues of concurrency and synchronization in the abstract, as they exist in any operating system kernel.

Critical Regions and Race Conditions:-

FirstRanker.com

- Code paths that access and manipulate shared data are called **critical regions** (also called **critical sections**). It is usually unsafe for multiple threads of execution to access the same resource simultaneously.
- To prevent concurrent access during critical regions, the programmer must ensure that code executes *atomically*, which means that operations complete without interruption as if the entire critical region were one indivisible instruction.
- It is a bug if it is possible for two threads of execution to be simultaneously executing within the same critical region. When this occurs, it is called a **race condition**, so-named because the threads raced to get there first. Debugging race conditions is often difficult because they are not easily reproducible.
- Ensuring that unsafe concurrency is prevented and that race conditions do not occur is called **synchronization**.

Android - Architecture

Android operating system is a stack of software components which is roughly divided into five sections and four main layers as shown below in the architecture diagram.

		Applications		
Home Dia	ler SMS/MMS	IM Brows er	Camera Ala	erm Calculator
Contacts Voice	e Dial Email	Calendar Media Player	Albums Cit	adk
		Application Framewo	rk	
Activity Manager	Window Manager	Content Providers	View System	Notification Manager
Package Manager	Telephony Manager	Resource Manager	Lo cation Manager	XMPP Service
	Libraries		Andro	oid Runtime
Surface Manager	Media Framework	SQLite	Core Libraries	
OpenGLIES	FreeType	LibWebCore	Dalvk Virtual Machine	
SGL	SSL	Liba		
		Linux Kernel		
Display Driver	Camera Driver	Bluetooth Driver	Flash Memory Driver	Binder (IPC) Driver



Linux kernel

At the bottom of the layers is Linux - Linux 3.6 with approximately 115 patches. This provides a level of abstraction between the device hardware and it contains all the essential hardware drivers like camera, keypad, display etc. Also, the kernel handles all the things that Linux is really good at such as networking and a vast array of device drivers, which take the pain out of interfacing to peripheral hardware.

Libraries

On top of Linux kernel there is a set of libraries including open-source Web browser engine WebKit, well known library libc, SQLite database which is a useful repository for storage and sharing of application data, libraries to play and record audio and video, SSL libraries responsible for Internet security etc.

Android Libraries

This category encompasses those Java-based libraries that are specific to Android development. Examples of libraries in this category include the application framework libraries in addition to those that facilitate user interface building, graphics drawing and database access. A summary of some key core Android libraries available to the Android developer is as follows –

- **android.app** Provides access to the application model and is the cornerstone of all Android applications.
- **android.content** Facilitates content access, publishing and messaging between applications and application components.
- **android.database** Used to access data published by content providers and includes SQLite database management classes.
- android.opengl A Java interface to the OpenGL ES 3D graphics rendering API.
- **android.os** Provides applications with access to standard operating system services including messages, system services and inter-process communication.
- android.text Used to render and manipulate text on a device display.
- android.view The fundamental building blocks of application user interfaces.
- **android.widget** A rich collection of pre-built user interface components such as buttons, labels, list views, layout managers, radio buttons etc.
- **android.webkit** A set of classes intended to allow web-browsing capabilities to be built into applications.

Having covered the Java-based core libraries in the Android runtime, it is now time to turn our attention to the C/C++ based libraries contained in this layer of the Android software stack.

Android Runtime

This is the third section of the architecture and available on the second layer from the bottom. This section provides a key component called **Dalvik Virtual Machine** which is a kind of Java Virtual Machine specially designed and optimized for Android.

The Dalvik VM makes use of Linux core features like memory management and multithreading, which is intrinsic in the Java language. The Dalvik VM enables every Android application to run in its own process, with its own instance of the Dalvik virtual machine. **www.FirstRanker.com**

The Android runtime also provides a set of core libraries which enable Android application developers to write Android applications using standard Java programming language.

Application Framework

FirstRanker.com

The Application Framework layer provides many higher-level services to applications in the form of Java classes. Application developers are allowed to make use of these services in their applications.

The Android framework includes the following key services -

- Activity Manager Controls all aspects of the application lifecycle and activity stack.
- **Content Providers** Allows applications to publish and share data with other applications.
- **Resource Manager** Provides access to non-code embedded resources such as strings, color settings and user interface layouts.
- Notifications Manager Allows applications to display alerts and notifications to the user.
- View System An extensible set of views used to create application user interfaces.

Android Services

Android service is a component that is *used to perform operations on the background* such as playing music, handle network transactions, interacting content providers etc. It doesn't has any UI (user interface).

, Her.con

The service runs in the background indefinitely even if application is destroyed.

Moreover, service can be bounded by a component to perform interactivity and inter process communication (IPC).

The android.app.Service is subclass of ContextWrapper class.

Note: Android service is not a thread or separate process.

Life Cycle of Android Service

There can be two forms of a service. The lifecycle of service can follow two different paths: started or bound.

- 1. Started
- 2. Bound

1) Started Service

A service is started when component (like activity) calls **startService**() method, now it runs in the background indefinitely. It is stopped by **stopService**() method. The service can stop itself by calling the **stopSelf**() method.



2) Bound Service

A service is bound when another component (e.g. client) calls **bindService**() method. The client can unbind the service by calling the **unbindService**() method.

The service cannot be stopped until all clients unbind the service.

