

UNIT I: Introduction

- What is an Operating System?
- Mainframe Systems
- Desktop Systems
- Multiprocessor Systems
- Distributed Systems
- Clustered System
- Real -Time Systems
- Handheld Systems
- Computing Environments

What is an Operating System?

A program that acts as an intermediary between a user of a computer and the computer hardware.

- Operating system goals:
 - ☞ Execute user programs and make solving user problems easier.
 - ☞ Make the computer system convenient to use.
- Use the computer hardware in an efficient manner.

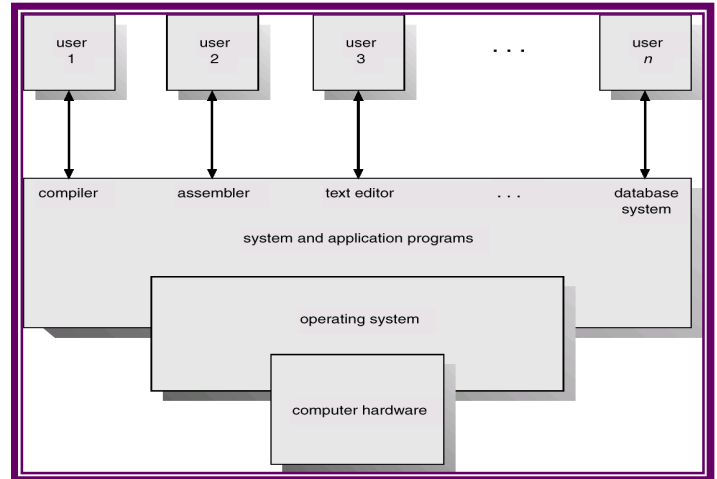
Computer System Components

1. Hardware – provides basic computing resources (CPU, memory, I/O devices).
2. Operating system – controls and coordinates the use of the hardware among the various application programs for the various users.
3. Applications programs – define the ways in which the system resources are used to solve the computing problems of the users

(compilers, database systems, video games, business programs).

4. Users (people, machines, other computers).

Abstract View of System Components




Operating System Definitions

- Resource allocator – manages and allocates resources.
- Control program – controls the execution of user programs and operations of I/O devices .
- Kernel – the one program running at all times (all else being application programs).


Mainframe Systems

- Reduce setup time by batching similar jobs
 - Automatic job sequencing – automatically transfers control from one job to another.
- First rudimentary operating system.
- Resident monitor
 - ☞ initial control in monitor
 - ☞ control transfers to job
 - ☞ when job completes control transfers back to monitor

Memory Layout for a Simple Batch System



Multiprogrammed Batch Systems



OS Features Needed for

Multiprogramming

- I/O routine supplied by the system.
- Memory management – the system must allocate the memory to several jobs.
- CPU scheduling – the system must choose

among several jobs ready to run.

- Allocation of devices.

Time-Sharing Systems–Interactive

Computing

- The CPU is multiplexed among several jobs that are kept in memory and on disk (the CPU is allocated to a job only if the job is in memory).
- A job swapped in and out of memory to the disk.
- On-line communication between the user and the system is provided; when the operating system finishes the execution of one command, it seeks the next “control statement” from the user’s keyboard.
- On-line system must be available for users to access data and code.


Desktop Systems

- Personal computers – computer system dedicated to a single user.
- I/O devices – keyboards, mice, display screens, small printers.
- User convenience and responsiveness.
- Can adopt technology developed for larger operating system’ often individuals have sole use of computer and do not need advanced CPU utilization or protection features.
- May run several different types of operating systems (Windows, MacOS, UNIX, Linux)

Parallel Systems

- Multiprocessor systems with more than one CPU in close communication.
- Tightly coupled system – processors share memory and a clock; communication usually takes place through the shared memory.
- Advantages of parallel system:
 - ☞ Increased throughput
 - ☞ Economical
 - ☞ Increased reliability
 - ▣ graceful degradation
 - ▣ fail-soft systems
- Symmetric multiprocessing (SMP)
 - ☞ Each processor runs an identical copy of the operating system.
 - ☞ Many processes can run at once without performance deterioration.
 - ☞ Most modern operating systems support SMP
- Asymmetric multiprocessing
 - ☞ Each processor is assigned a specific task; master processor schedules and allocates work to slave processors.
 - ☞ More common in extremely large systems


Symmetric Multiprocessing Architecture



Distributed Systems

- Distribute the computation among several physical processors.
- Loosely coupled system – each processor has its own local memory; processors communicate with one another through various communication lines, such as high-speed buses or telephone lines.
- Advantages of distributed systems.
 - ☞ Resources Sharing
 - ☞ Computation speed up – load sharing
 - ☞ Reliability
 - ☞ Communications
- Requires networking infrastructure.
- Local area networks (LAN) or Wide area networks (WAN)
- May be either client-server or peer-to-peer systems.

General Structure of Client-Server



Clustered Systems

- Clustering allows two or more systems to share storage.
- Provides high reliability.
- Asymmetric clustering: one server runs the application while other servers stand by.

- Symmetric clustering: all N hosts are running the application.

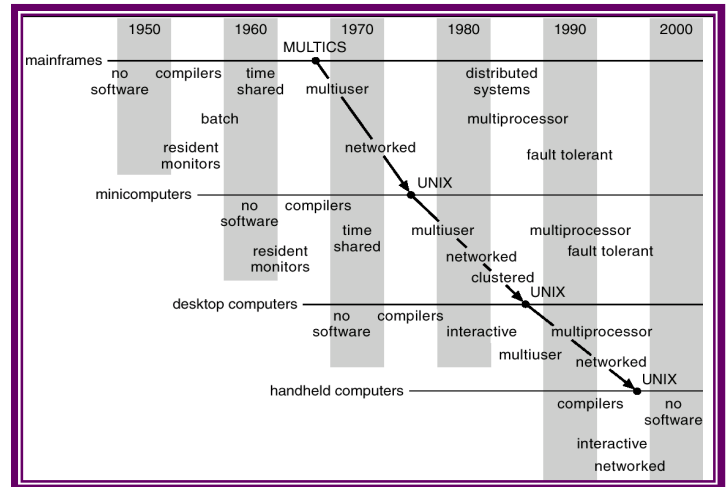
Real-Time Systems

- Often used as a control device in a dedicated application such as controlling scientific experiments, medical imaging systems, industrial control systems, and some display systems.
- Well-defined fixed-time constraints.
- Real-Time systems may be either hard or soft real-time.
- Hard real-time:
 - ☞ Secondary storage limited or absent, data stored in short term memory, or read-only memory (ROM)
 - ☞ Conflicts with time-sharing systems, not supported by general-purpose operating systems.
- Soft real-time
 - ☞ Limited utility in industrial control of robotics
 - ☞ Useful in applications (multimedia, virtual reality) requiring advanced operating-system features.

Handheld Systems

- Personal Digital Assistants (PDAs)
- Cellular telephones
- Issues:
 - ☞ Limited memory
 - ☞ Slow processors & Small display screens.

Migration of Operating-System Concepts and Features



Computing Environments

- Traditional computing
- Web-Based Computing

Operating-System Structures

- System Components
- Operating System Services
- System Calls
- System Programs
- System Structure
- Virtual Machines
- System Design and Implementation
- System Generation

Common System Components

- Process Management
- Main Memory Management
- File Management
- I/O System Management
- Secondary Management
- Networking
- Protection System
- Command-Interpreter System

Process Management

- A process is a program in execution. A process needs certain resources, including CPU time, memory, files, and I/O devices, to accomplish its task.
- The operating system is responsible for the following activities in connection with process management.
 - ☞ Process creation and deletion.
 - ☞ process suspension and resumption.
 - ☞ Provision of mechanisms for:
 - ▣ process synchronization
 - ▣ process communication

Main-Memory Management

- Memory is a large array of words or bytes, each with its own address. It is a repository of quickly accessible data shared by the CPU and I/O devices.
- Main memory is a volatile storage device. It loses its contents in the case of system failure.
- The operating system is responsible for the following activities in connections with memory management:
 - ☞ Keep track of which parts of memory are currently being used and by whom.
 - ☞ Decide which processes to load when memory space becomes available.
 - ☞ Allocate and deallocate memory space as needed.

File Management

- A file is a collection of related information defined by its creator. Commonly, files represent programs (both source and object forms) and data.
- The operating system is responsible for the following activities in connections with file management:
 - ☞ File creation and deletion.
 - ☞ Directory creation and deletion.
 - ☞ Support of primitives for manipulating

files and directories.

- ☞ Mapping files onto secondary storage.
- ☞ File backup on stable (nonvolatile) storage media.

I/O System Management

- The I/O system consists of:
 - ☞ A buffer-caching system
 - ☞ A general device-driver interface
 - ☞ Drivers for specific hardware devices

Secondary-Storage Management

- Since main memory (primary storage) is volatile and too small to accommodate all data and programs permanently, the computer system must provide secondary storage to back up main memory.
- Most modern computer systems use disks as the principle on-line storage medium, for both programs and data.
- The operating system is responsible for the following activities in connection with disk management:
 - ☞ Free space management
 - ☞ Storage allocation
 - ☞ Disk scheduling

Networking (Distributed Systems)

- A distributed system is a collection processors that do not share memory or a clock. Each processor has its own local memory.
- The processors in the system are connected through a communication network.
- Communication takes place using a protocol.
- A distributed system provides user access to various system resources.
- Access to a shared resource allows:
 - ☞ Computation speed-up
 - ☞ Increased data availability
 - ☞ Enhanced reliability

Protection System

- Protection refers to a mechanism for controlling access by programs, processes, or users to both system and user resources.
- The protection mechanism must:
 - ☞ distinguish between authorized and unauthorized usage.
 - ☞ specify the controls to be imposed.
 - ☞ provide a means of enforcement.

Command-Interpreter System

- Many commands are given to the operating system by control statements which deal with:
 - ☞ process creation and management
 - ☞ I/O handling
 - ☞ secondary-storage management
 - ☞ main-memory management
 - ☞ file-system access
 - ☞ protection
 - ☞ networking
- The program that reads and interprets control statements is called variously:
 - ☞ command-line interpreter
 - ☞ shell (in UNIX)Its function is to get and execute the next command statement.

Operating System Services

- Program execution – system capability to load a program into memory and to run it.
- I/O operations – since user programs cannot execute I/O operations directly, the operating system must provide some means to perform I/O.
- File-system manipulation – program capability to read, write, create, and delete files.
- Communications – exchange of information between processes executing either on the same computer or on different systems tied together by a network. Implemented via shared memory or message passing.

- Error detection – ensure correct computing by detecting errors in the CPU and memory hardware, in I/O devices, or in user programs.

Additional Operating System Functions

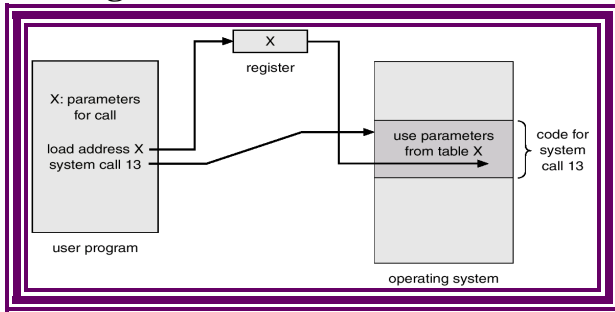
Additional functions exist not for helping the user, but rather for ensuring efficient system operations.

- Resource allocation – allocating resources to multiple users or multiple jobs running at the same time.
- Accounting – keep track of and record which users use how much and what kinds of computer resources for account billing or for accumulating usage statistics.
- Protection – ensuring that all access to system resources is controlled.

System Calls

- System calls provide the interface between a running program and the operating system.
 - ☞ Generally available as assembly-language instructions.
 - ☞ Languages defined to replace assembly language for systems programming allow system calls to be made directly (e.g., C, C++)
- Three general methods are used to pass parameters between a running program and the operating system.
 - ☞ Pass parameters in registers.
 - ☞ Store the parameters in a table in memory, and the table address is passed as a parameter in a register.
 - ☞ Push (store) the parameters onto the stack by the program, and pop off the stack by operating system.


Passing of Parameters As A Table




Types of System Calls

- Process control
- File management
- Device management
- Information maintenance
- Communications

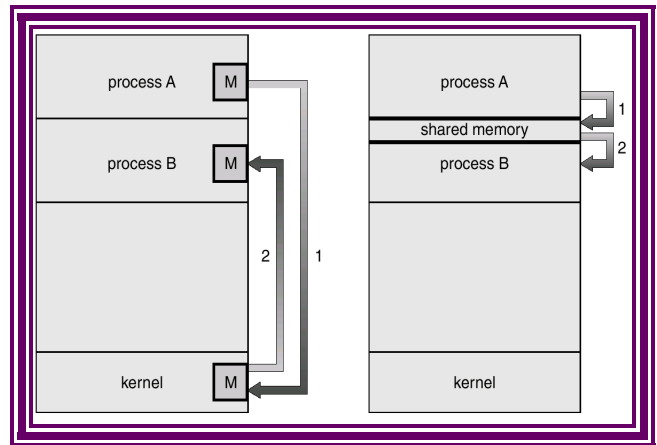
MS-DOS Execution



UNIX Running Multiple Programs



Communication Models Communication may take place using either message passing or shared memory.




System Programs

- System programs provide a convenient environment for program development and execution. They can be divided into:
 - ☞ File manipulation
 - ☞ Status information
 - ☞ File modification
 - ☞ Programming language support
 - ☞ Program loading and execution
 - ☞ Communications
 - ☞ Application programs
- Most users' view of the operation system is defined by system programs, not the actual system calls.

MS-DOS System Structure

- MS-DOS – written to provide the most functionality in the least space
 - ☞ not divided into modules
 - ☞ Although MS-DOS has some structure, its interfaces and levels of functionality are not well separated

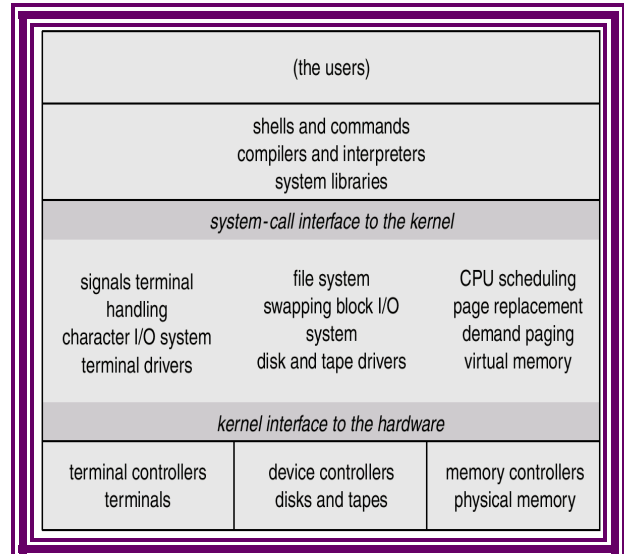
MS-DOS Layer Structure



UNIX System Structure

- UNIX – limited by hardware functionality, the original UNIX operating system had limited structuring. The UNIX OS consists of two separable parts.
 - ☞ Systems programs
 - ▣ Consists of everything below the system-call interface and above the physical hardware
 - ▣ Provides the file system, CPU scheduling, memory management, and other operating-system functions; a large number of functions for one level.
 - ☞ The kernel

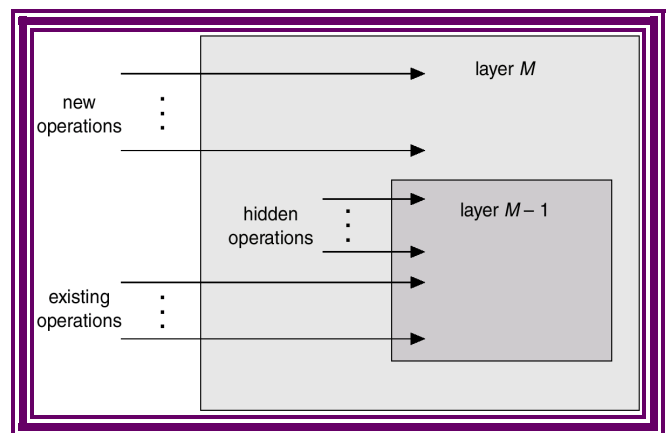
UNIX System Structure



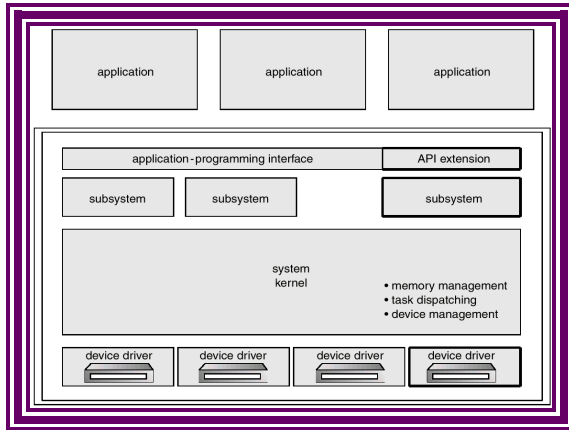
Layered Approach

- The operating system is divided into a number of layers (levels), each built on top of lower layers. The bottom layer (layer 0), is the hardware; the highest (layer N) is the user interface.
- With modularity, layers are selected such that each uses functions (operations) and services of only lower-level layers.

An Operating System Layer




OS/2 Layer Structure



Microkernel System Structure

- Moves as much from the kernel into “user” space.
- Communication takes place between user modules using message passing.
- Benefits:
 - easier to extend a microkernel
 - easier to port the operating system to new architectures
 - more reliable (less code is running in kernel mode)
 - more secure

Windows NT Client-Server Structure



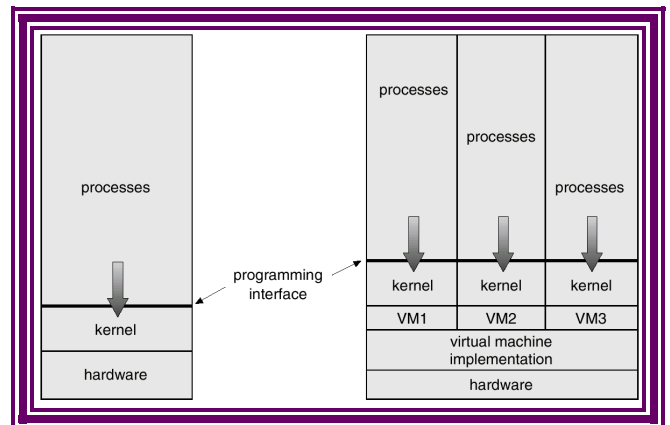
Virtual Machines

- A virtual machine takes the layered approach to its logical conclusion. It treats hardware and the operating system kernel as though

they were all hardware.

- A virtual machine provides an interface identical to the underlying bare hardware.
- The operating system creates the illusion of multiple processes, each executing on its own processor with its own (virtual) memory.
- The resources of the physical computer are shared to create the virtual machines.
 - ☞ CPU scheduling can create the appearance that users have their own processor.
 - ☞ Spooling and a file system can provide virtual card readers and virtual line printers.
 - ☞ A normal user time-sharing terminal serves as the virtual machine operator’s console.

System Models



Advantages/Disadvantages of Virtual Machines

- The virtual-machine concept provides complete protection of system resources since each virtual machine is isolated from all other virtual machines. This isolation, however, permits no direct sharing of resources.
- A virtual-machine system is a perfect vehicle for operating-systems research and development. System development is done on the virtual machine, instead of on a physical machine and so does not disrupt


normal system operation.

- The virtual machine concept is difficult to implement due to the effort required to provide an exact duplicate to the underlying machine.

Java Virtual Machine

- Compiled Java programs are platform-neutral bytecodes executed by a Java Virtual Machine (JVM).
- JVM consists of
 - class loader ;
 - class verifier ;
 - runtime interpreter
- Just-In-Time (JIT) compilers increase performance

Java Virtual Machine



System Design Goals

- User goals – operating system should be convenient to use, easy to learn, reliable, safe, and fast.
- System goals – operating system should be easy to design, implement, and maintain, as well as flexible, reliable, error-free, and efficient.

Mechanisms and Policies

- Mechanisms determine how to do something, policies decide what will be done.
- The separation of policy from mechanism is a

very important principle, it allows maximum flexibility if policy decisions are to be changed later.

System Implementation

- Traditionally written in assembly language, operating systems can now be written in higher-level languages.
- Code written in a high-level language:
 - ☞ can be written faster.
 - ☞ is more compact.
 - ☞ is easier to understand and debug.
- An operating system is far easier to port (move to some other hardware) if it is written in a high-level language.

System Generation (SYSGEN)

- Operating systems are designed to run on any of a class of machines; the system must be configured for each specific computer site.
- SYSGEN program obtains information concerning the specific configuration of the hardware system.
- Booting – starting a computer by loading the kernel.
- Bootstrap program – code stored in ROM that is able to locate the kernel, load it into memory, and start its execution.

Processes

- Process Concept
- Process Scheduling
- Operations on Processes
- Cooperating Processes
- Interprocess Communication
- Communication in Client-Server Systems

Process Concept

- An operating system executes a variety of programs:

- ☞ Batch system – jobs
- ☞ Time-shared systems – user programs or tasks
- Textbook uses the terms job and process almost interchangeably.
- Process – a program in execution; process execution must progress in sequential fashion.
- A process includes:
 - ☞ program counter
 - ☞ stack
 - ☞ data section

Process State

- As a process executes, it changes state
 - ☞ **new**: The process is being created.
 - ☞ **running**: Instructions are being executed.
 - ☞ **waiting**: The process is waiting for some event to occur.
 - ☞ **ready**: The process is waiting to be assigned to a process.
 - ☞ **terminated**: The process has finished execution.

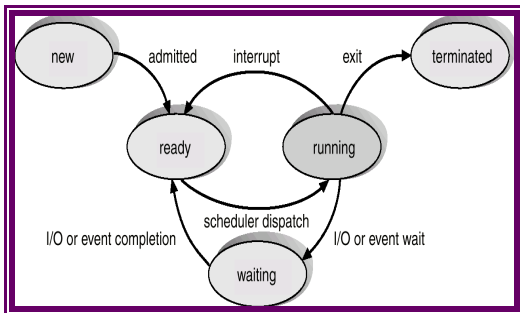


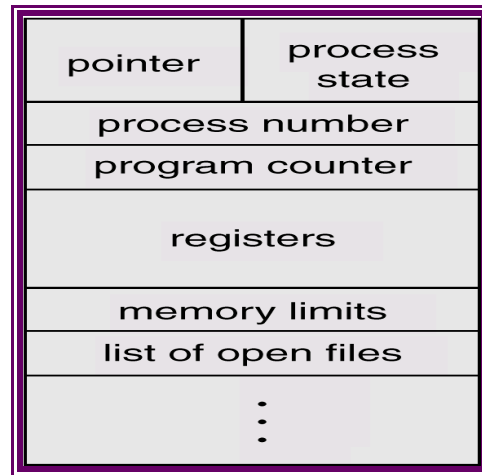
Diagram of Process State

Process Control Block (PCB)

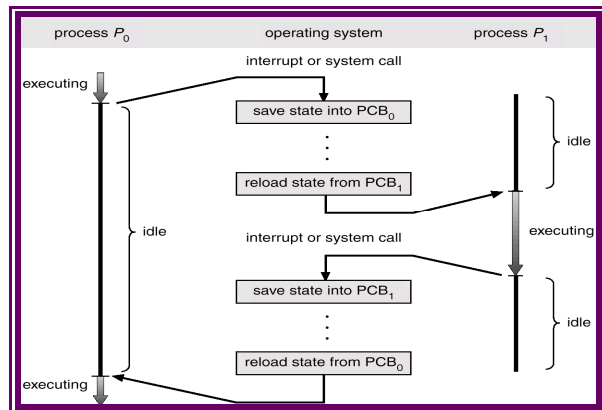
Information associated with each process.

- Process state
- Program counter
- CPU registers
- CPU scheduling information
- Memory-management information
- Accounting information
- I/O status information

Process Control Block (PCB)



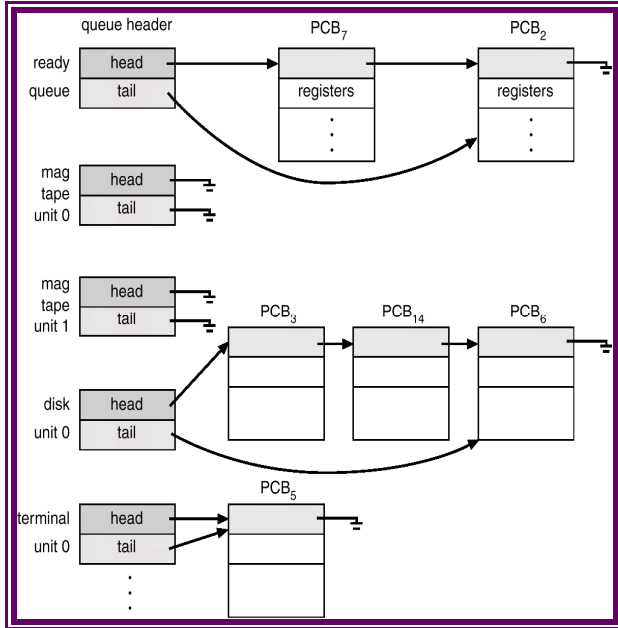
CPU Switch From Process to Process



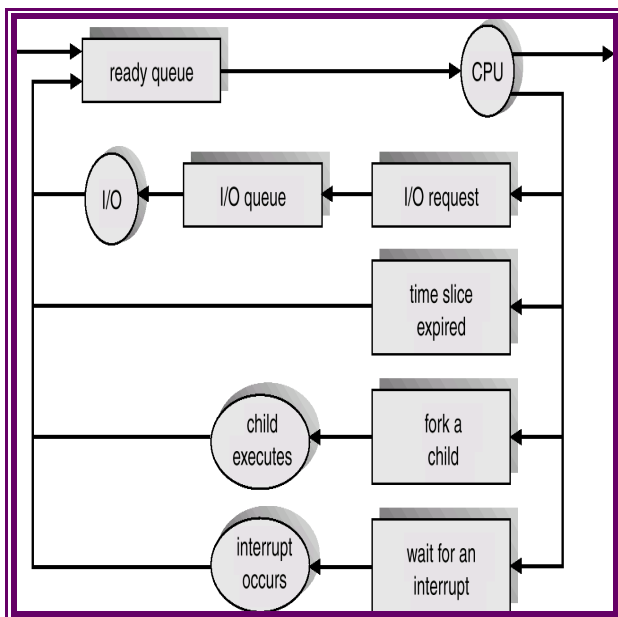
Process Scheduling Queues

- Job queue – set of all processes in the system.
- Ready queue – set of all processes residing in main memory, ready and waiting to execute.
- Device queues – set of processes waiting for an I/O device.
- Process migration between the various queues.

Ready Queue And Various I/O Device Queues



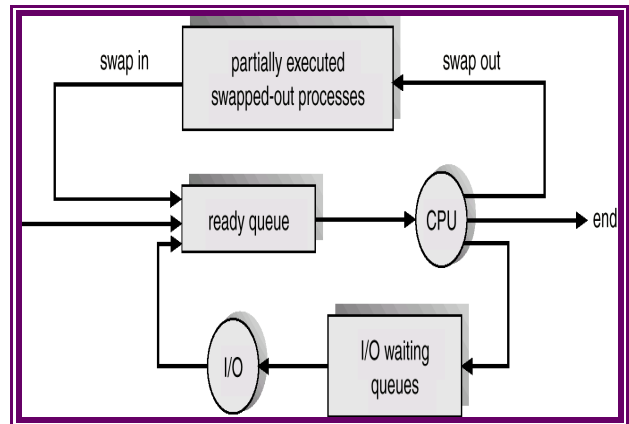
Representation of Process Scheduling



Schedulers

- Long-term scheduler (or job scheduler) – selects which processes should be brought into the ready queue.
- Short-term scheduler (or CPU scheduler) – selects which process should be executed next and allocates CPU.

Addition of Medium Term Scheduling



- Short-term scheduler is invoked very frequently (milliseconds) ⇒ (must be fast).
- Long-term scheduler is invoked very infrequently (seconds, minutes) ⇒ (may be slow).
- The long-term scheduler controls the degree of multiprogramming.
- Processes can be described as either:
 - ☞ I/O-bound process – spends more time doing I/O than computations, many short CPU bursts.
 - ☞ CPU-bound process – spends more time doing computations; few very long CPU bursts.

Context Switch


- When CPU switches to another process, the system must save the state of the old process and load the saved state for the new process.
- Context-switch time is overhead; the system does no useful work while switching.
- Time dependent on hardware support.

Process Creation

- Parent process create children processes, which, in turn create other processes, forming a tree of processes.

- 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.

Processes Tree on a UNIX System



Process Termination

- Process executes last statement and asks the operating system to decide 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.
 - ☞ Parent is exiting.
 - ☐ Operating system does not allow child to continue if its parent terminates.
 - ☐ Cascading termination.

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;
```
- Solution is correct, but can only use

BUFFER_SIZE-1 elements

Bounded-Buffer – Producer Process

```
item nextProduced;
while (1) {
    while (((in + 1) % BUFFER_SIZE) ==
out)
        ; /* do nothing */
    buffer[in] = nextProduced;
    in = (in + 1) % BUFFER_SIZE;
}
```

Bounded-Buffer – Consumer Process

```
item nextConsumed;

while (1) {
    while (in == out)
        ; /* do nothing */
    nextConsumed = buffer[out];
    out = (out + 1) % BUFFER_SIZE;
}
```

Interprocess Communication (IPC)

- 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)

Implementation Questions

- How are links established?
- Can a link be associated with more than two processes?
- How many links can there be between every pair of communicating processes?
- What is the capacity of a link?
- Is the size of a message that the link can accommodate fixed or variable?
- Is a link unidirectional or bi-directional?

Direct Communication

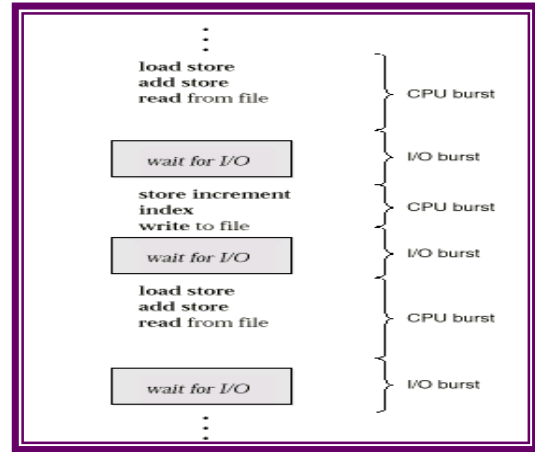
- Processes must name each other explicitly:
 - ☞ **send** (P, message) – send a message to process P
 - ☞ **receive**(Q, message) – receive a message from process Q
- Properties of communication link
 - ☞ Links are established automatically.
 - ☞ A link is associated with exactly one pair of communicating 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.

Alternating Sequence of CPU And I/O Bursts




CPU Scheduling

- Basic Concepts
- Scheduling Criteria
- Scheduling Algorithms
- Multiple-Processor Scheduling
- Real-Time Scheduling
- Algorithm Evaluation

Basic Concepts

- Maximum CPU utilization obtained with multiprogramming
- CPU-I/O Burst Cycle – Process execution consists of a *cycle* of CPU execution and I/O wait.
- CPU burst distribution

Histogram of CPU-burst Times



CPU Scheduler

- Selects from among the processes in memory that are ready to execute, and allocates the CPU to one of them.

- CPU scheduling decisions may take place when a process:
 1. Switches from running to waiting state.
 2. Switches from running to ready state.
 3. Switches from waiting to ready.
 4. Terminates.
- Scheduling under 1 and 4 is *nonpreemptive*.
- All other scheduling is *preemptive*.

Dispatcher

- Dispatcher module gives control of the CPU to the process selected by the short-term scheduler; this involves:
 - ☞ switching context
 - ☞ switching to user mode
 - ☞ jumping to the proper location in the user program to restart that program
- **Dispatch latency** – time it takes for the dispatcher to stop one process and start another running.

Scheduling Criteria

- CPU utilization – keep the CPU as busy as possible
- Throughput – # of processes that complete their execution per time unit
- Turnaround time – amount of time to execute a particular process
- Waiting time – amount of time a process has been waiting in the ready queue
- Response time – amount of time it takes from when a request was submitted until the first response is produced, **not** output (for time-sharing environment)


Optimization Criteria

- Max CPU utilization
- Max throughput
- Min turnaround time
- Min waiting time
- Min response time

First-Come, First-Served (FCFS) Scheduling

<u>Process</u>	<u>Burst Time</u>
P_1	24
P_2	3
P_3	3

-
- Suppose that the processes arrive in the order: P_1, P_2, P_3



- The Gantt Chart for the schedule is:
 - Waiting time for $P_1 = 0; P_2 = 24; P_3 = 27$
 - Average waiting time: $(0 + 24 + 27)/3 = 17$
- Suppose that the processes arrive in the order

$$P_2, P_3, P_1.$$

- The Gantt chart for the schedule is:
Waiting time for $P_1 = 6; P_2 = 0; P_3 = 3$
- Average waiting time: $(6 + 0 + 3)/3 = 3$
- Much better than previous case.
- *Convoy effect* short process behind long process

Shortest-Job-First (SJF) Scheduling

- Associate with each process the length of its next CPU burst. Use these lengths to schedule the process with the shortest time.
- Two schemes:
 - ☞ nonpreemptive – once CPU given to the process it cannot be preempted until completes its CPU burst.
 - ☞ preemptive – if a new process arrives with CPU burst length less than remaining time of current executing process, preempt. This scheme is know as the Shortest-Remaining-Time-First (SRTF).
- SJF is optimal – gives minimum average waiting time for a given set of processes.

Example of Non-Preemptive SJF

Process	Arrival Time	Burst Time
7	P_1	0.0
4	P_2	2.0
1	P_3	4.0
4	P_4	5.0

SJF (non-preemptive)

- Average waiting time = $(0 + 6 + 3 + 7)/4 - 4$

Example of Preemptive SJF

Process	Arrival Time	Burst Time
P_1	0.0	7
P_2	2.0	4
P_3	4.0	1
P_4	5.0	4

- SJF (preemptive)




Average waiting time = $(9 + 1 + 0 + 2)/4 - 3$

Determining Length of Next CPU Burst

- Can only estimate the length.
- Can be done by using the length of previous CPU bursts, using exponential averaging.

Prediction of the Length of the Next CPU Burst



increase the priority of the process.

Examples of Exponential Averaging

- $\alpha = 0$
 - ☞ $\tau_{n+1} = \tau_n$
 - ☞ Recent history does not count.
- $\alpha = 1$
 - ☞ $\tau_{n+1} = t_n$
 - ☞ Only the actual last CPU burst counts.
- If we expand the formula, we get:

$$\tau_{n+1} = \alpha t_n + (1 - \alpha) \alpha t_{n-1} + \dots$$

$$+ (1 - \alpha)^j \alpha t_{n-j} + \dots$$

$$+ (1 - \alpha)^{n-1} t_n \tau_0$$
- Since both α and $(1 - \alpha)$ are less than or equal to 1, each successive term has less weight than its predecessor.


Priority Scheduling

- A priority number (integer) is associated with each process
- The CPU is allocated to the process with the highest priority (smallest integer \equiv highest priority).
 - ☞ Preemptive
 - ☞ nonpreemptive
- SJF is a priority scheduling where priority is the predicted next CPU burst time.
- Problem \equiv Starvation – low priority processes may never execute.
- Solution \equiv Aging – as time progresses

Round Robin (RR)

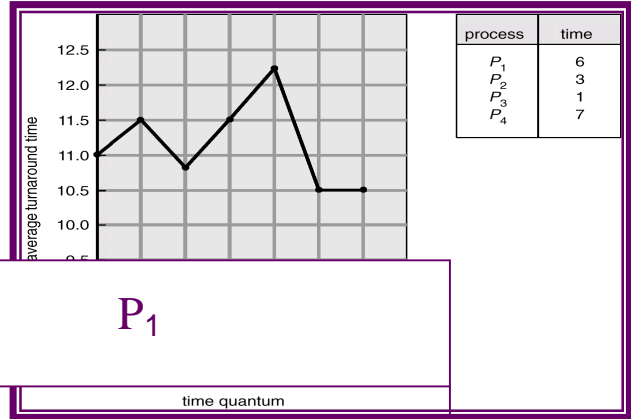
- Each process gets a small unit of CPU time (*time quantum*), usually 10-100 milliseconds. After this time has elapsed, the process is preempted and added to the end of the ready queue.
- If there are n processes in the ready queue and the time quantum is q , then each process gets $1/n$ of the CPU time in chunks of at most q time units at once. No process waits more than $(n-1)q$ time units.
- Performance
 - ☞ q large \Rightarrow FIFO
 - ☞ q small $\Rightarrow q$ must be large with respect to context switch, otherwise overhead is too high.

Example of RR with Time Quantum = 20




<u>Process</u>	<u>Burst Time</u>
P_1	53
P_2	17
P_3	68
P_4	24

- The Gantt chart is:



- Typically, higher average turnaround than SJF, but better *response*.

Time Quantum and Context Switch Time



Turnaround Time Varies With The Time Quantum

Multilevel Queue

- Ready queue is partitioned into separate queues:
 - foreground (interactive)
 - background (batch)
- Each queue has its own scheduling algorithm,
 - foreground – RR
 - background – FCFS
- Scheduling must be done between the queues.
 - ☞ Fixed priority scheduling; (i.e., serve all from foreground then from background). Possibility of starvation.
 - ☞ Time slice – each queue gets a certain amount of CPU time which it can

schedule amongst its processes; i.e., 80%
to foreground in RR

☞ 20% to background in FCFS

process will enter when that process needs
service

Example of Multilevel Feedback

Queue

■ Three queues:

☞ Q_0 – time quantum 8 milliseconds

☞ Q_1 – time quantum 16 milliseconds

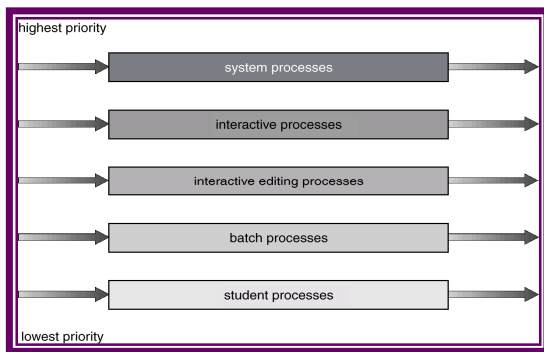
☞ Q_2 – FCFS

■ Scheduling

☞ A new job enters queue Q_0 which is served FCFS. When it gains CPU, job receives 8 milliseconds. If it does not finish in 8 milliseconds, job is moved to queue Q_1 .

☞ At Q_1 job is again served FCFS and receives 16 additional milliseconds. If it still does not complete, it is preempted and moved to queue Q_2 .


Multilevel Queue Scheduling



Multilevel Feedback Queue

- A process can move between the various queues; aging can be implemented this way.
- Multilevel-feedback-queue scheduler defined by the following parameters:
 - ☞ number of queues
 - ☞ scheduling algorithms for each queue
 - ☞ method used to determine when to upgrade a process
 - ☞ method used to determine when to demote a process
 - ☞ method used to determine which queue a

Multilevel Feedback Queues



Multiple-Processor Scheduling

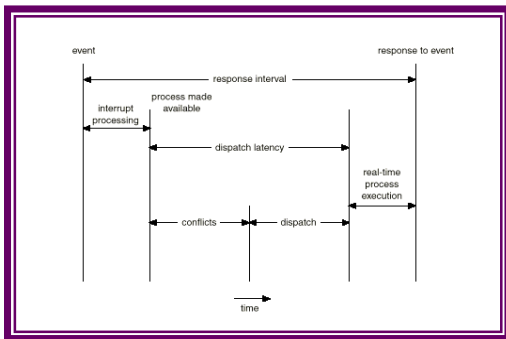
- CPU scheduling more complex when multiple CPUs are available.
- *Homogeneous processors* within a multiprocessor.

- *Load sharing*
- *Asymmetric multiprocessing* – only one processor accesses the system data structures, alleviating the need for data sharing.

Real-Time Scheduling

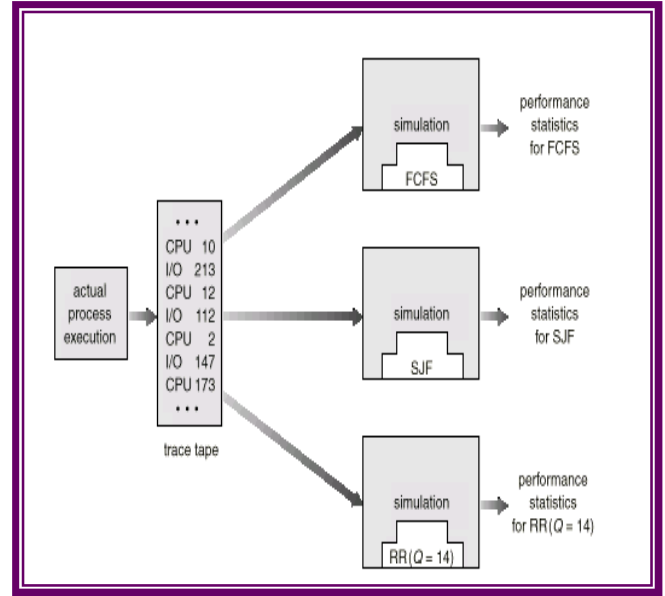
- *Hard real-time* systems – required to complete a critical task within a guaranteed amount of time.
- *Soft real-time* computing – requires that critical processes receive priority over less fortunate ones.

Dispatch Latency




Algorithm Evaluation

- *Deterministic modeling* – takes a particular predetermined workload and defines the performance of each algorithm for that workload.
- *Queueing models*
- *Implementation Evaluation of CPU Schedulers by Simulation*



Solaris 2 Scheduling



Windows 2000 Priorities

	real-time	high	above normal	normal	below normal	idle priority
time-critical	31	15	15	15	15	15
highest	26	15	12	10	8	6
above normal	25	14	11	9	7	5
normal	24	13	10	8	6	4
below normal	23	12	9	7	5	3
lowest	22	11	8	6	4	2
idle	16	1	1	1	1	1

UNIT I

2 marks with Answers

1. What are the advantages of Multiprocessor System?

- i) Increased throughput
- ii) Economy of scale
- iii) Increased reliability

2. Define Hand Held System.

Hand Held System includes PDAs, Palm Pilots or cellular phones with connectivity to a network such as Internet.

3. List the principle advantages of the multiprogramming.

Multiprogramming system increases the CPU utilization by organizing jobs so that the CPU always has a one to execute.

4. What is a control program?

A control program manages the execution of the user programs to prevent errors and improper use of the computer. It is concerned with the operation and control of I/O devices.

5. Define the term Fault tolerant.

Functions can be distributed among several processors, then the failure of one process will not halt the system, only slow it down. The ability to continue providing service proportional to the level of surviving hardware is called graceful degradation. System designed for graceful degradation are called fault tolerant.

6. List the OS Services.

- Program execution
- I/o operations
- File-system manipulations
- Communications
- Error detection
- Resource allocation

Accounting
Protection

7. Define Operating System.

Os is a program that manages the computer hardware. It also provides a basis for application programs and act as an intermediary between a user of a computer and the computer hardware.

8. Mention the Categories of a System program.

- File management
- Status information
- File modification
- Programming language support
- Program loading and execution
- Communications

9. Define Interprocess communication.

IPC provides a mechanism to allow processes to communicate and to synchronize their actions without sharing the same address space. IPC is particularly useful in a distributed environment where the communicating processes may reside on different computers connected with a network.

10. What is scheduler?

The os must select, for scheduling purpose, processes from the queue in some fashion. The selection process is carried out by the appropriate scheduler.

Or The processes are kept in the ready queue. The scheduler is used to decide which process is to assign the cpu from the queue.

11. What is a process control block?

Each process is represented in the os by a process control block(PCB) also called task control block. A PCB contains many pieces of information associated with a specific process.

12. Define system call.

System call provides a interface between a process and the operating system. System calls are generally available as a assembly language instructions.

PART-B

1. Write short notes on Batch system and multiprogrammed system. (8)
2. List and explain the operating system services. (8)
3. Describe multiprocessor system. (8)
4. Discuss about Distributed System. (8)
5. Describe about the system components. (16)
6. Explain the various Process scheduling algorithms. (16)