

1 (a)	<p>Explain process management and various process schedule algorithm in Linux operating system</p> <p><b>Process Management</b></p> <p>A process is the basic context within which all user-requested activity is serviced within the OS. To be compatible with other UNIX systems, Linux must use a process model similar to those of other versions of UNIX.</p> <p><b>8.11.1 The Fork/Exec Process Model</b></p> <p>UNIX process management separates the creation of processes and the running of a new program into two distinct operations. The <i>fork</i> system call creates a new process. A new program is run after a call to <i>exec</i>. Under UNIX, a process encompasses all the information that the operating system must maintain to track the context of a single execution of a single program.</p> <p>Under Linux, process properties fall into three groups:</p> <ul style="list-style-type: none"><li>• <b>Process Identity:</b> The process identity consists of mainly following items:</li><li>• <b>Process ID (PID):</b> The unique identifier for the process; used to specify processes to the OS when an application makes a system call to signal, modify, or wait for another process.</li><li>• <b>Credentials:</b> Each process must have an associated user ID and one or more group IDs that determine the process's rights to access system resources and files.</li><li>• <b>Personality:</b> Not traditionally found on UNIX systems, but under Linux each process has an associated personality identifier that can slightly modify the semantics of certain system calls. Used primarily by emulation libraries to request that system calls be compatible with certain specific flavors of UNIX.</li><li>• <b>Process environment:</b> The process's environment is inherited from its parent, and is composed of two null-terminated vectors:<ul style="list-style-type: none"><li>o The <b>argument vector</b> lists the command-line arguments used to invoke the running program; conventionally starts with the name of the program itself</li><li>o The <b>environment vector</b> is a list of "NAME=VALUE" pairs that associates named environment variables with arbitrary textual values.</li></ul>Passing environment variables among processes and inheriting variables by a process's children are flexible means of passing information to components of the user-mode system software. The environment-variable mechanism provides a customization of the operating system that can be set on a per-process basis, rather than being configured for the system as a whole.</li></ul> <p><b>Context:</b> It is the (constantly changing) state of a running program at any point in time. There are various parts:</p> <p><b>Scheduling context</b> is the most important part of the process context; it is the information that the scheduler needs to suspend and restart the process.</p> <p><b>Accounting:</b> The kernel maintains accounting information about the resources currently being consumed by each process, and the total resources consumed by the process in its lifetime so far.</p> <ul style="list-style-type: none"><li>o <b>File table</b> is an array of pointers to kernel file structures. When making file I/O system calls, processes refer to files by their index into this table.</li></ul>	10
-------	--	----

- o **File System Context:** Whereas the file table lists the existing open files, the file-system context applies to requests to open new files. The current root and default directories to be used for new file searches are stored here.
- o **Signal-handler table** defines the routine in the process's address space to be called when specific signals arrive.
- o **Virtual-memory context** of a process describes the full contents of its private address space.

### Processes and Threads

Linux uses the same internal representation for processes and threads; a thread is simply a new process that happens to share the same address space as its parent. A distinction is only made when a new thread is created by the **clone** system call:

- fork** creates a new process with its own entirely new process context
- clone** creates a new process with its own identity, but that is allowed to share the data structures of its parent

Using **clone** gives an application fine-grained control over exactly what is shared between two threads.

### SCHEDULING

The job of allocating CPU time to different tasks within an OS. While scheduling is normally thought of as the running and interrupting of processes, in Linux, scheduling also includes the running of the various kernel tasks. Running kernel tasks encompasses both tasks that are requested by a running process and tasks that execute internally on behalf of a device driver.

Various aspects of scheduling in Linux are discussed here.

- Kernel Synchronization:** A request for kernel-mode execution can occur in two ways:
  - o A running program may request an operating system service, either explicitly via a system call, or implicitly, for example, when a page fault occurs.
  - o A device driver may deliver a hardware interrupt that causes the CPU to start executing a kernel-defined handler for that interrupt.

Kernel synchronization requires a framework that will allow the kernel's critical sections to run without interruption by another critical section.

Linux uses two techniques to protect critical sections:

- o Normal kernel code is non-preemptable : when a time interrupt is received while a process is executing a kernel system service routine, the kernel's **need\_resched** flag is set so that the scheduler will run once the system all has completed and control is about to be returned to user mode.
- o The second technique applies to critical sections that occur in an interrupt service routines. By using the processor's interrupt control hardware to disable interrupts during a critical section, the kernel guarantees that it can proceed without the risk of concurrent access of shared data structures.

- Process Scheduling:** Linux uses two process-scheduling algorithms:

- o A time-sharing algorithm for fair preemptive scheduling between multiple processes
- o A real-time algorithm for tasks where absolute priorities are more important than fairness

A process's scheduling class defines which algorithm to apply. For time-sharing processes, Linux uses a prioritized, credit based algorithm.

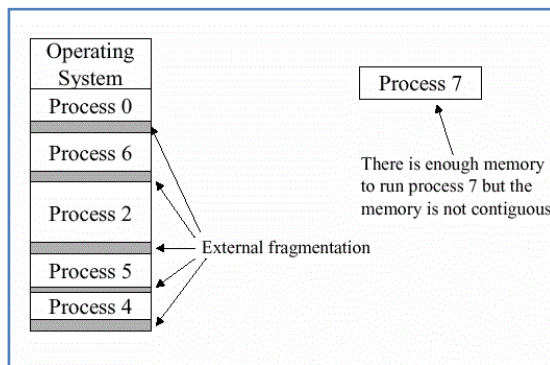
	<p>□ <b>Symmetric Multiprocessing:</b> Linux 2.0 was the first Linux kernel to support SMP hardware; separate processes or threads can execute in parallel on separate processors. To preserve the kernel's non-preemptible synchronization requirements, SMP imposes the restriction, via a single kernel spinlock, that only one processor at a time may execute kernel-mode code.</p>	
2 (a)	<p><u>Explain hardware solution for critical section problem</u></p> <p>There are many software algorithms for enforcing mutual exclusion. But they have high processing overhead and the risk of logical errors. Hence, we will here consider few hardware approaches.</p> <p><b>4.3.1 Interrupt Disabling</b></p> <p>In a single-processor system, concurrent processes cannot be overlapped, but can be interleaved. Moreover, a process will continue to execute until it is interrupted. Hence, to guarantee mutual exclusion, a process has to be prevented from being interrupted. This capability can be provided in the form of primitives defined by OS kernel for disabling and enabling interrupts. So, a process should not be interrupted till it completes its critical section. But, by doing so, the efficiency of OS will slow-down, as OS cannot switch between the processes. Moreover, this approach does not guarantee mutual exclusion in case of multi-processor systems.</p> <p>Normally, access to a memory location excludes any other access to that same location. Hence, processor designers have proposed several machine instructions that carry out two actions atomically (that is, an instruction is a single step and not interrupted), such as reading and writing or reading and testing, of a single memory location with one instruction fetch cycle. During execution of the instruction, access to the memory location is blocked for any other instruction referencing that location. Two most commonly implemented instructions are explained below:</p> <p>□ <b>Compare and Swap Instruction:</b> It compares the contents of a memory location to a given value. If they are same, it modifies the contents of that memory location to a given new value. This is done as a single atomic operation. The atomicity guarantees that the new value is calculated based on up-to-date information; if the value had been updated by another process in the meantime, the write would fail. The result of the operation must indicate whether it performed the substitution; this can be done either with a simple boolean response, or by returning the value read from the memory location.</p> <p>□ <b>Exchange Instruction:</b> The instruction exchanges the contents of a register with that of a memory location. The use of a special machine instruction to enforce mutual exclusion has a number of advantages:</p> <ul style="list-style-type: none"> <li>□ It is applicable to any number of processes on either a single processor or multiple processors sharing main memory.</li> <li>□ It is simple and therefore easy to verify.</li> <li>□ It can be used to support multiple critical sections; each critical section can be defined by its own variable. There are some serious disadvantages as well:</li> <li>□ <b>Busy waiting is employed:</b> Thus, while a process is waiting for access to a critical section, it continues to consume processor time.</li> <li>□ <b>Starvation is possible:</b> When a process leaves a critical section and more than one process is waiting, the selection of a waiting process is arbitrary. Thus, some process could indefinitely be denied access.</li> <li>□ <b>Deadlock is possible:</b> Consider the following scenario on a single-processor system. Process P1 executes the special instruction and enters its critical section. P1 is then interrupted to give the processor to P2, which has higher priority. If P2 now attempts to use the same resource as P1, it will be denied access because of the mutual exclusion mechanism. Thus, it will go into a busy waiting loop. However, P1 will never be dispatched because it is of lower priority than another ready process, P2.</li> </ul>	5

2.b) Explain fragmentation and its types with neat diagram

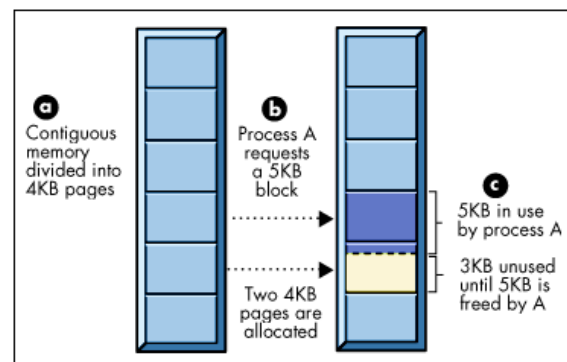
**External Fragmentation** – when total memory space exists to satisfy a request, but the available spaces are not contiguous.

- Storage is fragmented into a large number of small holes caused by first-fit and best-fit strategies for memory allocation plus variable-partition scheme

**Internal Fragmentation** – when allocated memory may be slightly larger than requested memory; this size difference is memory internal to a partition, but not being used



External Fragmentation Example



Internal Fragmentation Example

(b) What are the different types of scheduler?

**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 –

- Long-Term Scheduler
- Short-Term Scheduler
- Medium-Term Scheduler

**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.

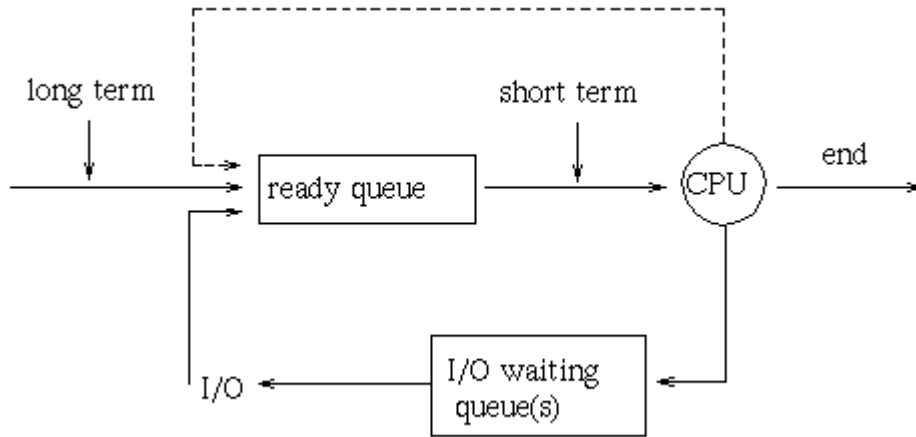
**Short Term Scheduler :**

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.

5

Short-term schedulers are *faster* than long-term schedulers.



**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. A 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.

3 (a) What is Critical Section and Mutual Exclusion

5

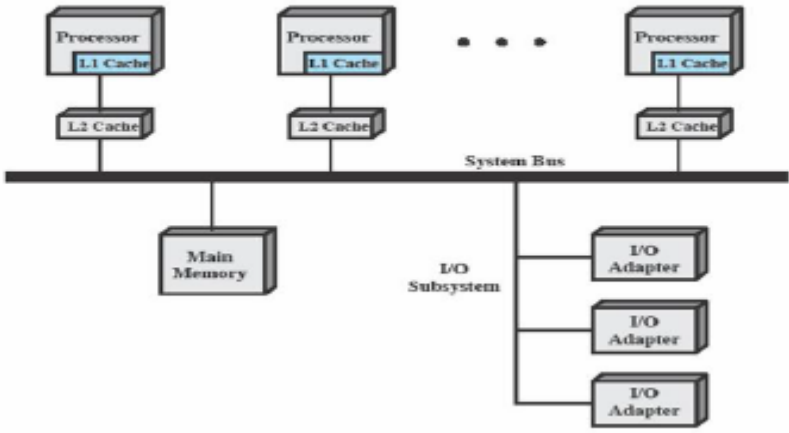
**Critical Section**

Suppose two processes require an access to single non-sharable resource like printer. During the course of execution, each process will be sending commands to the I/O device, receiving status information, sending data, and/or receiving data. We will refer to such a resource as a **critical resource**, and the portion of the program that uses it a **critical section** of the program.

It is important that only one program at a time be allowed in its critical section. We cannot simply rely on the OS to understand and enforce this restriction. For example, in case of the printer, we want only one process to have control of the printer till it finishes printing the entire document. Otherwise, lines from competing processes will be interleaved.

**Mutual Exclusion**

while one process executes the shared variable, all other processes desiring to do so at the same time moment should be kept waiting; when that process has finished executing the shared variable, one of the processes waiting; while that process has finished executing the shared variable, one of the processes waiting to do so should be allowed to proceed. In this fashion, each process executing the shared data (variables) excludes all others from doing so simultaneously. This is called Mutual Exclusion.

<p>(b)</p>	<p><b>What is symmetric multiprocessor system? What are its benefits?</b></p> <p>Traditionally, the computer has been viewed as a sequential machine. That is, a processor executes instructions one at a time in a sequence and each instruction is a sequence of operations. But, as computer technology has evolved, parallel processing got importance. One of the popular approaches for providing parallelism is symmetric multiprocessors (SMPs), where processors are replicated.</p> <p><b>3.8.1 SMP Architecture</b></p> <p><b>Single instruction single data (SISD) stream:</b> A single processor executes a single instruction stream to operate on data stored in a single memory.</p> <ul style="list-style-type: none"> <li>□ <b>Single instruction multiple data (SIMD) stream:</b> Each instruction is executed on a different set of data by the different processors.</li> <li>□ <b>Multiple instruction single data (MISD) stream:</b> A sequence of data is transmitted to a set of processors, each of execute a different instruction sequence.</li> <li>□ <b>Multiple instruction multiple data (MIMD) stream:</b> A set of processors simultaneously execute different instruction sequences on different data sets. Based on the communication among processors, MIMD can be further divided. If every processor has a dedicated memory, then each processing element is a self-contained computer. Communication among the computers is either via fixed path or via some network. Such a system is known as a <i>cluster</i>. If the processors share a common memory, it is known as <i>shared-memory multiprocessor</i>. This again can be further divided into <i>master/slave</i> architecture and <i>SMP</i>.</li> </ul> <p>The master/slave architecture has disadvantages:</p> <ul style="list-style-type: none"> <li>□ A failure of the master brings down the whole system</li> <li>□ As master has to do all scheduling and process management, the performance may slow down.</li> </ul> <p>But, in SMP, the kernel can execute on any processor and it allows portions of kernel to execute in parallel. Here, each processor does self-scheduling from the pool of available process or threads.</p> <p><b>3.8.2 SMP Organization</b></p> <p>In SMP, there are multiple processors, each of which contains its own control unit, arithmetic-logic unit and registers as shown in Figure 3.14. Each processor has access to a shared main memory and the I/O devices through a shared bus.</p>  <p style="text-align: center;"><b>Figure 3.14 SMP Organization</b></p>	<p>5</p>
<p>4 (a)</p>	<p><b>What is system call? What are the types of system call? Give examples of each.</b></p> <p>The <b>system call</b> provides an interface to the operating system services. Application developers often do not have direct access to the system calls, but can access them through an application programming interface (API). The functions that are included in the API</p>	<p>6</p>

invoke the actual system calls. By using the API, certain benefits can be gained:

- Portability: as long a system supports an API, any program using that API can compile and run.
- Ease of Use: using the API can be significantly easier then using the actual system call.

## 2.2. Types of System Calls

There are 5 different categories of system calls:

process control, file manipulation, device manipulation, information maintenance and communication.

### 1.12.2.1. Process Control

A running program needs to be able to stop execution either normally or abnormally. When execution is stopped abnormally, often a dump of memory is taken and can be examined with a debugger.

### 1.12.2.2. File Management

Some common system calls are *create*, *delete*, *read*, *write*, *reposition*, or *close*. Also, there is a need to determine the file attributes – *get* and *set* file attribute. Many times the OS provides an API to make these system calls.

### 1.12.2.3. Device Management

Process usually require several resources to execute, if these resources are available, they will be granted and control returned to the user process. These resources are also thought of as devices. Some are physical, such as a video card, and others are abstract, such as a file.

User programs *request* the device, and when finished they *release* the device. Similar to files, we can *read*, *write*, and *reposition* the device.

### 1.12.2.4. Information Management

Some system calls exist purely for transferring information between the user program and the operating system. An example of this is *time*, or *date*.

The OS also keeps information about all its processes and provides system calls to report this information.

### 1.12.2.5. Communication

**Table 2.1 Types of System Calls**

Category	System Calls
Process Control	<ul style="list-style-type: none"> <li>• end, abort</li> <li>• load, execute</li> <li>• create process, terminate process</li> <li>• get process attributes, set process attributes</li> <li>• wait for time</li> <li>• wait event, signal event</li> <li>• allocate and free memory</li> </ul>
File Management	<ul style="list-style-type: none"> <li>• create file, delete file</li> <li>• open, close</li> <li>• read, write, reposition</li> <li>• get file attributes, set file attributes</li> </ul>
Device Management	<ul style="list-style-type: none"> <li>• request device, release device</li> <li>• read, write, reposition</li> <li>• get device attributes, set device attributes</li> <li>• logically attach or detach devices</li> </ul>
Information Maintenance	<ul style="list-style-type: none"> <li>• get time or date, set time or date</li> <li>• get system data, set system data</li> <li>• get process, file, or device attributes</li> <li>• set process, file, or device attributes</li> </ul>
Communications	<ul style="list-style-type: none"> <li>• create, delete communication connection</li> <li>• send, receive messages</li> <li>• transfer status information</li> <li>• attach or detach remote devices</li> </ul>

There are two models of interprocess communication, the message-passing model and the shared memory model.

- Message-passing uses a common mailbox to pass messages between processes.
- Shared memory use certain system calls to create and gain access to create and gain

	access to regions of memory owned by other processes. The two processes exchange information by reading and writing in the shared data.	
(b)	<p><b>What is a PCB? What are its contents</b></p> <p>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 –</p> <ul style="list-style-type: none"> <li><input type="checkbox"/> <b>Identifier:</b> A unique identifier associated with this process, to distinguish it from all other processes.</li> <li><input type="checkbox"/> <b>State:</b> If the process is currently executing, it is in the running state.</li> <li><input type="checkbox"/> <b>Priority:</b> Priority level relative to other processes.</li> <li><input type="checkbox"/> <b>Program counter:</b> The address of the next instruction in the program to be executed.</li> </ul> <p><b>Memory pointers:</b> Includes pointers to the program code and data associated with this process, plus any memory blocks shared with other processes.</p> <ul style="list-style-type: none"> <li><input type="checkbox"/> <b>Context data:</b> These are data that are present in registers in the processor while the process is executing.</li> <li><input type="checkbox"/> <b>I/O status information:</b> Includes outstanding I/O requests, I/O devices (e.g., disk drives) assigned to this process, a list of files in use by the process, and so on.</li> <li><input type="checkbox"/> <b>Accounting information:</b> May include the amount of processor time and clock time used, time limits, account numbers, and so on.</li> </ul> <p>The above information is stored in a data structure known as <i>process control block</i></p>	4
5 (a)	<p><b>Write and explain Banker's algorithm for deadlock avoidance</b></p> <p>The resource-allocation graph algorithm is not applicable when there are multiple instances for each resource. The banker's algorithm addresses this situation, but it is less efficient. The name was chosen because this algorithm could be used in a banking system to ensure that the bank never allocates its available cash such that it can no longer satisfy the needs of all its customers.</p> <p>When a new process enters the system, it must declare the maximum number of instances of each resource type that it may need. This number may not exceed the total number of resources in the system. When a user requests a set of resources, the system must determine whether the allocation of these resources will leave the system in a safe state. If it will, the resources are allocated; otherwise, the process must wait until some other process releases enough resources.</p> <p>Data structure for Banker's algorithms is as below –  Let <math>n</math> be the number of processes in the system and <math>m</math> be the number of resource types.</p> <ul style="list-style-type: none"> <li><input type="checkbox"/> <b>Available:</b> Vector of length <math>m</math> indicating number of available resources. If <math>Available[j] = k</math>, there are <math>k</math> instances of resource type <math>R_j</math> available.</li> <li><input type="checkbox"/> <b>Max:</b> An <math>n \times m</math> matrix defines the maximum demand of each process. If <math>Max[i,j] = k</math>, then process <math>P_i</math> may request at most <math>k</math> instances of resource type <math>R_j</math>.</li> <li><input type="checkbox"/> <b>Allocation:</b> An <math>n \times m</math> matrix defines the number of resources currently allocated to each process. If <math>Allocation[i, j] = k</math> then <math>P_i</math> is currently allocated <math>k</math> instances of <math>R_j</math>.</li> <li><input type="checkbox"/> <b>Need:</b> An <math>n \times m</math> matrix indicates remaining resource need of each process. If <math>Need[i,j] = k</math>, then <math>P_i</math> may need <math>k</math> more instances of <math>R_j</math> to complete its task. Note that, <math>Need[i,j] = Max[i,j] - Allocation[i,j]</math>.</li> </ul> <p>The Banker's algorithm has two parts:</p> <ol style="list-style-type: none"> <li>1. <b>Safety Algorithm:</b> It is for finding out whether a system is in safe state or not. The steps are as given below – <ol style="list-style-type: none"> <li>1. Let <i>Work</i> and <i>Finish</i> be vectors of length <math>m</math> and <math>n</math>, respectively. Initialize:</li> </ol> </li> </ol>	5



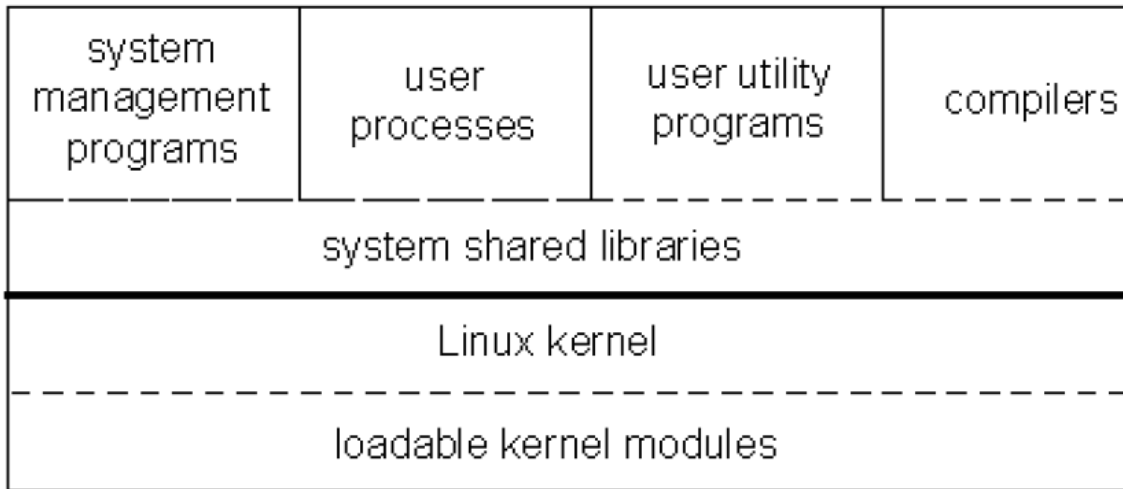
	<p><math>Work = Available</math>  <math>Finish [i] = false</math> for <math>i = 1, 2, 3, \dots, n</math>.</p> <p>2. Find an <math>i</math> such that both:  (a) <math>Finish [i] = false</math>  (b) <math>Need_i \leq Work</math></p> <p>If no such <math>i</math> exists, go to step 4.  3. <math>Work = Work + Allocation_i</math>  <math>Finish[i] = true</math>  go to step 2.</p> <p>4. If <math>Finish [i] == true</math> for all <math>i</math>, then the system is in a safe state.</p> <p><b>Resource – Request Algorithm:</b>  Let <math>Request_i</math> be the request vector for process <math>P_i</math>. If <math>Request_i [j] = k</math> then process <math>P_i</math> wants <math>k</math> instances of resource type <math>R_j</math>.</p> <ol style="list-style-type: none"> <li>If <math>Request_i \leq Need_i</math> go to step 2. Otherwise, raise error condition, since process has exceeded its maximum claim.</li> <li>If <math>Request_i \leq Available</math>, go to step 3. Otherwise <math>P_i</math> must wait, since resources are not available.</li> <li>Pretend to allocate requested resources to <math>P_i</math> by modifying the state as follows:  <math>Available = Available - Request_i</math>;  <math>Allocation_i = Allocation_i + Request_i</math>;  <math>Need_i = Need_i - Request_i</math>;</li> </ol> <p>If the resulting resource allocation is safe, then the transaction is complete and the process <math>P_i</math> is allocated its resources. If the new state is unsafe, then <math>P_i</math> must wait for <math>Request_i</math>, and the old resource-allocation state is restored</p>	
(b)	<p>Write short note about Inter process communication in Linux</p> <p>Like UNIX, Linux informs processes that an event has occurred via signals. There is a limited number of signals, and they cannot carry information: Only the fact that a signal occurred is available to a process. The Linux kernel does not use signals to communicate with processes with are running in kernel mode, rather, communication within the kernel is accomplished via scheduling states and wait-queue structures.</p> <p>Passing of Data among Processes: The pipe mechanism allows a child process to inherit a communication channel to its parent; data written to one end of the pipe can be read by the other. Shared memory offers an extremely fast way of communicating; any data written by one process to a shared memory region can be read immediately by any other process that has mapped that region into its address space. To obtain synchronization, however, shared memory must be used in conjunction with another inter process communication mechanism.</p> <p>Explain the components of Linux OS</p> <p>Ans: The Linux system is composed of three main bodies of code, in line with most traditional UNIX implementations:</p> <ol style="list-style-type: none"> <li><b>1. Kernel.</b> The kernel is responsible for maintaining all the important abstractions of the operating system, including such things as virtual memory and processes.</li> <li><b>2. System libraries.</b> The system libraries define a standard set of functions through which applications can interact with the kernel. These functions implement much of the operating-system functionality</li> </ol>	5

that does not need the full privileges of kernel code.

**3. System utilities.** The system utilities are programs that perform individual, specialized management tasks. Some system utilities may be invoked just once to initialize and configure some aspect of the system; others—

known as *daemons* in UNIX terminology—may run permanently, handling such tasks as responding to incoming network connections, accepting logon requests from terminals, and updating log files.

Figure 21.1 illustrates the various components that make up a full Linux system. The most important distinction here is between the kernel and everything else. All the kernel code executes in the processor's privileged mode with full access to all the physical resources of the computer. Linux refers to this privileged mode as **kernel mode**.



6 (a) With a neat diagram explain the steps in handling page fault

6

In Demand paging, we need to distinguish the pages which are in the memory and pages which are there on the disk. For this purpose, the valid – invalid bit is used. When this bit is set to *valid*, it indicates the page is in the memory. Whereas, the value of bit as *invalid* indicates page is on the disk.

If the process tries to access a page which is not in the memory (means, it is on the disk), **page fault** occurs. The paging hardware notices the *invalid* bit in the page table and cause a trap to the OS. This trap is the result of the failure of OS to bring the desired page into memory. This error has to be corrected. The procedure for handling this page fault is as shown in Figure 6.16. The steps are explained below:

1. We check an internal table (usually kept with the process control block) for this process, to determine whether the reference was a valid or invalid memory access.
2. If the reference was invalid, we terminate the process. If it was valid, but we have not yet brought in that page into memory, it is brought now.
3. We find a free frame.

4. We schedule a disk operation to read the desired page into the newly allocated frame.
5. When the disk read is complete, we modify the internal table kept with the process and the page table to indicate that the page is now in memory.
6. We restart the instruction that was interrupted by the illegal address trap. The process can now access the page as though it had always been in memory.

It is important to realize that, because we save the state (registers, condition code, instruction counter etc.) of the interrupted process when the page fault occurs, we can restart the process in exactly the same place and state. In an extreme situation, a process may start executing with no page in the memory. So, each time an instruction has to be executed, page fault occurs and the required page needs to be brought into the memory. This situation is called as *pure demand paging*.

That is, no page is brought into the memory until it is required.

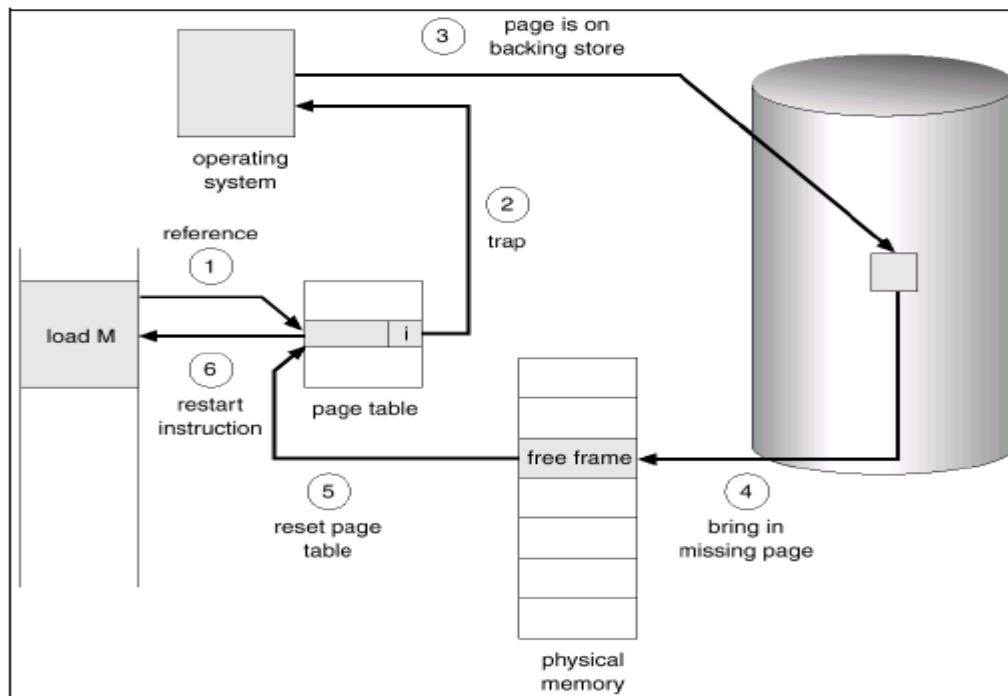


Figure 6.16 Steps in handling page fault

(b) Explain the components of Linux OS

Ans: The Linux system is composed of three main bodies of code, in line with most traditional UNIX implementations:

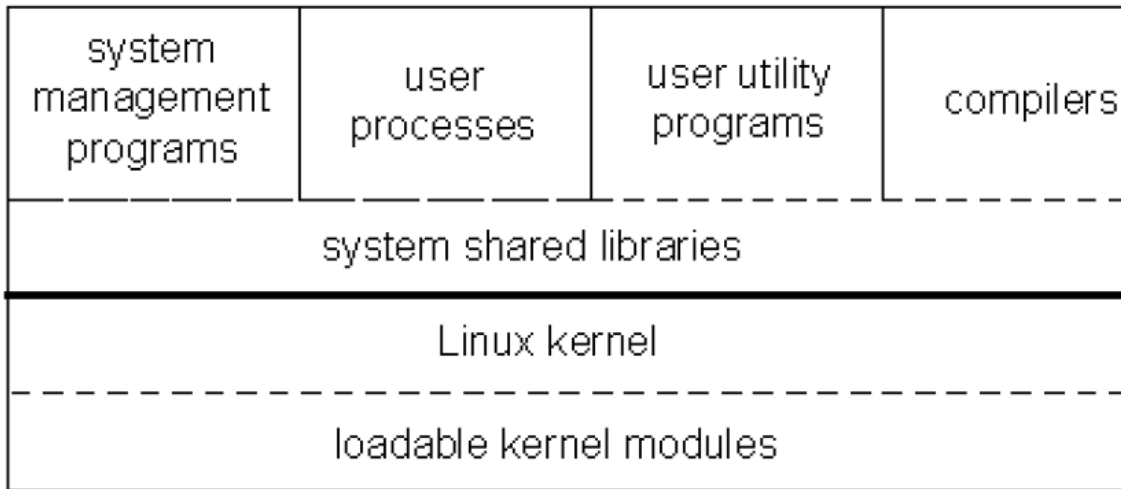
1. **Kernel.** The kernel is responsible for maintaining all the important abstractions of the operating system, including such things as virtual memory and processes.
2. **System libraries.** The system libraries define a standard set of functions through which applications can interact with the kernel. These functions implement much of the operating-system functionality

that does not need the full privileges of kernel code.

**3. System utilities.** The system utilities are programs that perform individual, specialized management tasks. Some system utilities may be invoked just once to initialize and configure some aspect of the system; others—

known as *daemons* in UNIX terminology—may run permanently, handling such tasks as responding to incoming network connections, accepting logon requests from terminals, and updating log files.

Figure 21.1 illustrates the various components that make up a full Linux system. The most important distinction here is between the kernel and everything else. All the kernel code executes in the processor's privileged mode with full access to all the physical resources of the computer. Linux refers to this privileged mode as **kernel mode**.



7

Mention the different page table structure explain in brief

There are three common techniques for structuring a page table. They are:

**Hierarchical Paging** - Break up the logical address space into multiple page tables. A simple technique is a two-level page table.

**Hashed Page**

10

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

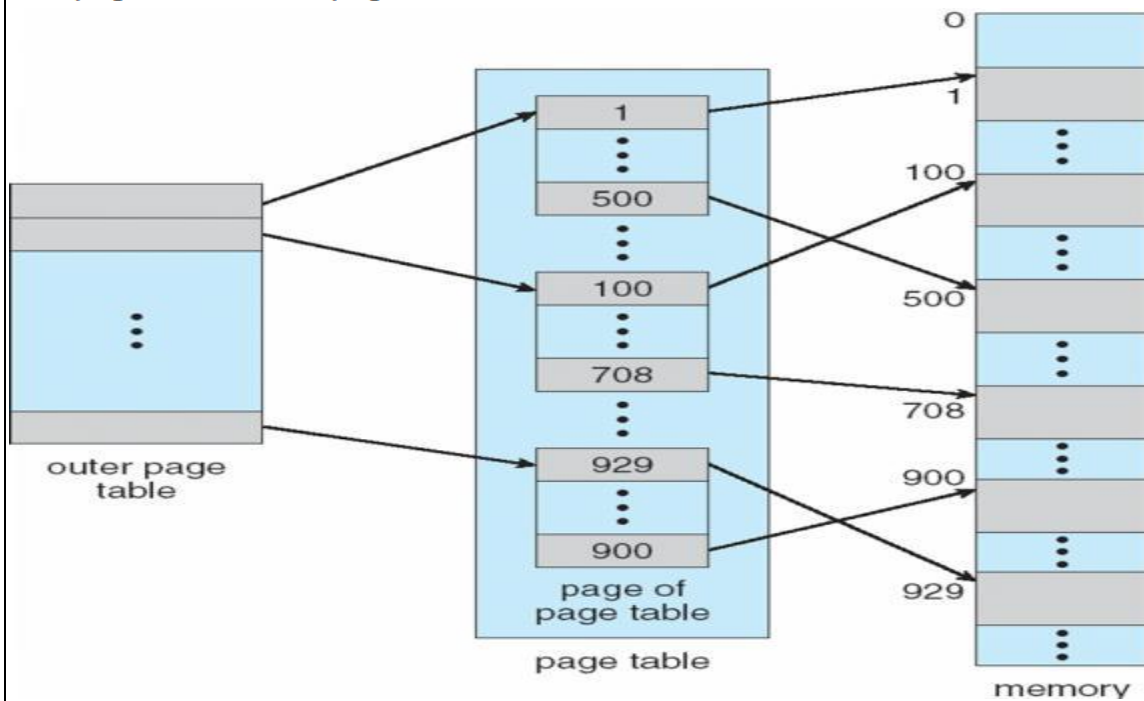
Since the page table is paged, the page number is further divided into:

- a 12-bit page number
- a 10-bit page offset

Thus, a logical address is as follows:

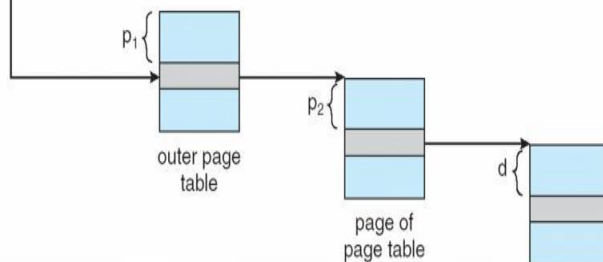
page number		page offset
$p_1$	$p_2$	$d$
12	10	10

where  $p_1$  is an index into the outer page table, and  $p_2$  is the displacement within the page of the outer page table



logical address

$p_1$	$p_2$	$d$
-------	-------	-----



outer page	inner page	offset
$p_1$	$p_2$	$d$
42	10	12

2nd outer page	outer page	inner page	offset
$p_1$	$p_2$	$p_3$	$d$

Hashed Page Table –

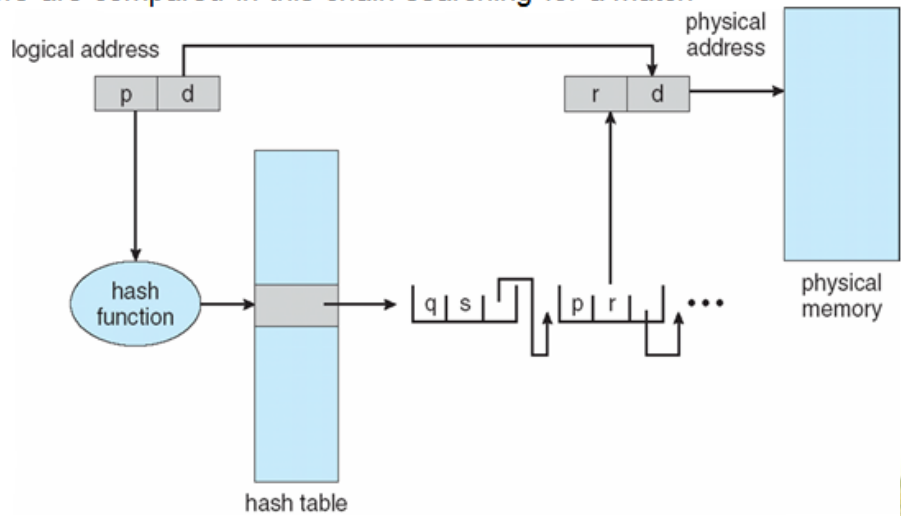
Common in address spaces > 32 bits

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



Inverted Page Table –

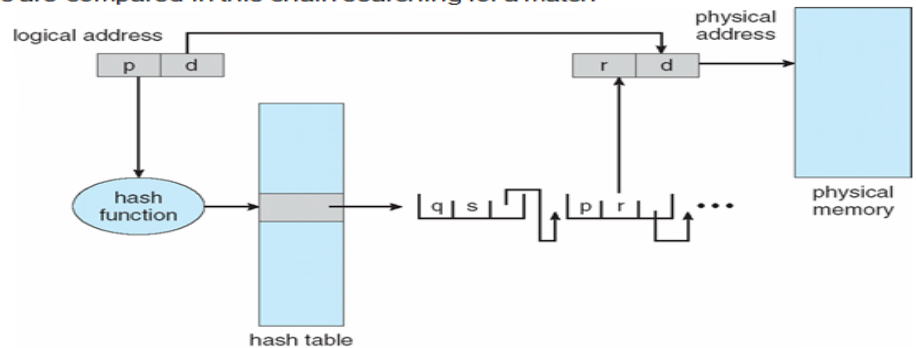
Common in address spaces > 32 bits

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



8 (a) Define Security problem and system threat

10

We have discussed earlier that each file/directory can be secured by controlling access to them. But, total security cannot be achieved by this. Security violations of the system can be categorized as intentional (malicious) or accidental. It is easier to protect against accidental misuse. Malicious access has following forms:

- Unauthorized reading of data
- Unauthorized modification of data
- Unauthorized destruction of data
- Preventing legitimate use of the system

To protect the system, we must take security measures at four levels:

1. **Physical:** The site or sites containing the computer systems must be physically

	<p>secured against armed or surreptitious entry by intruders.</p> <p>2. <b>Human:</b> Users must be screened carefully to reduce the chance of authorizing a user who then gives access to an intruder.</p> <p>3. <b>Network:</b> Data being transferred via network is prone to attack.</p> <p>4. <b>Operating System:</b> The system must protect itself security breaches.</p> <p><b>SYSTEM THREATS</b></p> <p>Most OS allows processes to spawn (creating child process) other processes. In such a situation, it is possible that OS resources and user files are misused. The methods for misusing are discussed here:</p> <ul style="list-style-type: none"> <li><input type="checkbox"/> <b>Worms:</b> A worm is a process that uses the spawn mechanism to attack system performance. It creates multiple copies of itself using system resources and locking all other processes from using the system.</li> <li><input type="checkbox"/> <b>Viruses:</b> It is a code embedded within a legitimate program. They spread into other programs and mess-up the system – modifying/destroying files, causing system crash etc.</li> <li><input type="checkbox"/> <b>Denial of Service:</b> Here, the hacker does not gain any information, but disables legitimate user from accessing system.</li> </ul>	
(b)	<p><b>Explain the role of Operating system to prevent</b></p> <p>The major security problem for OS is authentication. The protection system depends on an ability identify the programs and processes currently executing. Generally authentication is based on one or more of three items:</p> <ul style="list-style-type: none"> <li><input type="checkbox"/> User possession (a key or a card)</li> <li><input type="checkbox"/> User knowledge (identifier or password)</li> <li><input type="checkbox"/> User attribute (fingerprint, retina, signature etc)</li> </ul> <p><b>Passwords</b></p> <p>Passwords are used to protect the data in the computer, when there are no complete protection schemes. User is asked to provide the username and password during the access. If the password matches with the one stored in the system, access is allowed.</p> <p><b>8.5.2 Password Vulnerabilities</b></p> <p>Though passwords are popularly used because of its easy usage and understanding, it is vulnerable to threats. They can be guessed, accidentally exposed, illegally transferred from one to other etc. To avoid these various possibilities can be tried:</p> <ul style="list-style-type: none"> <li><input type="checkbox"/> keeping non-guessable passwords</li> <li><input type="checkbox"/> keep changing the passwords frequently</li> <li><input type="checkbox"/> keeping encrypted passwords</li> <li><input type="checkbox"/> using One – Time Passwords (OTP)</li> <li><input type="checkbox"/> using biometrics authentication</li> </ul>	
9	<p><b>Explain the following with respect to the file system: i) Contiguous allocation ii) Linked allocation iii) Indexed allocation</b></p> <p>The direct-access nature of disks allows us flexibility in the implementation of files. In almost every case, many files will be stored on the same disk. The main problem is how to allocate space to these files so that disk space is utilized effectively and files can be accessed quickly. Three major methods of allocating disk space are: contiguous, linked, and indexed.</p> <p><b>Contiguous Allocation</b></p> <p>In contiguous allocation, files are assigned to contiguous areas of secondary storage. A</p>	10

user specifies in advance the size of the area needed to hold a file to be created. If the desired amount of contiguous space is not available, the file cannot be created. A contiguous allocation of disk space is shown in Figure 7.11.

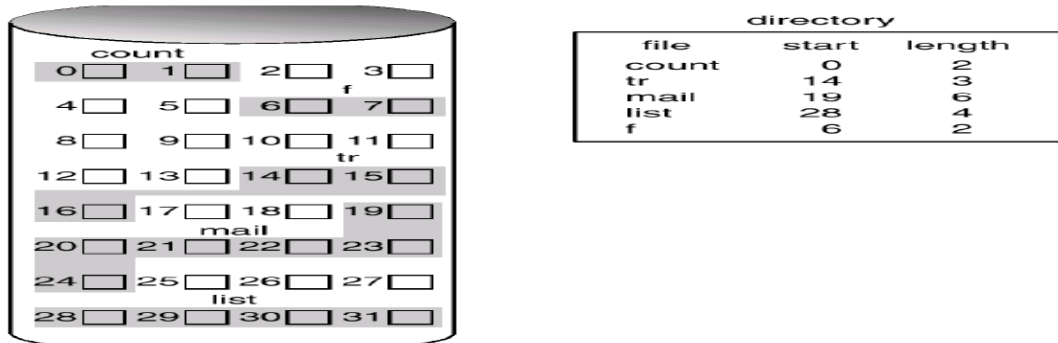


Figure 7.11 Contiguous allocation of disk space

One advantage of contiguous allocation is that all successive records of a file are normally physically adjacent to each other. This increases the accessing speed of records. It means that if records are scattered through the disk it is accessing will be slower. For sequential access the file system remembers the disk address of the last block and when necessary reads the next block. For direct access to block B of a file that starts at location L, we can immediately access block L+B. Thus contiguous allocation supports both sequential and direct accessing. The disadvantage of contiguous allocation algorithm is, it suffers from external fragmentation. As files are allocated and deleted, the free disk space is broken into little pieces. Depending on the total amount of disk storage and the average file size, external fragmentation may be a minor or a major problem.

**Linked Allocation**

Linked allocation solves all problems of contiguous allocation. With linked allocation, each file is a linked list of disk blocks; the disk blocks may be scattered anywhere on the disk. The directory contains a pointer to the first and last blocks of the file as shown in Figure 7.12.

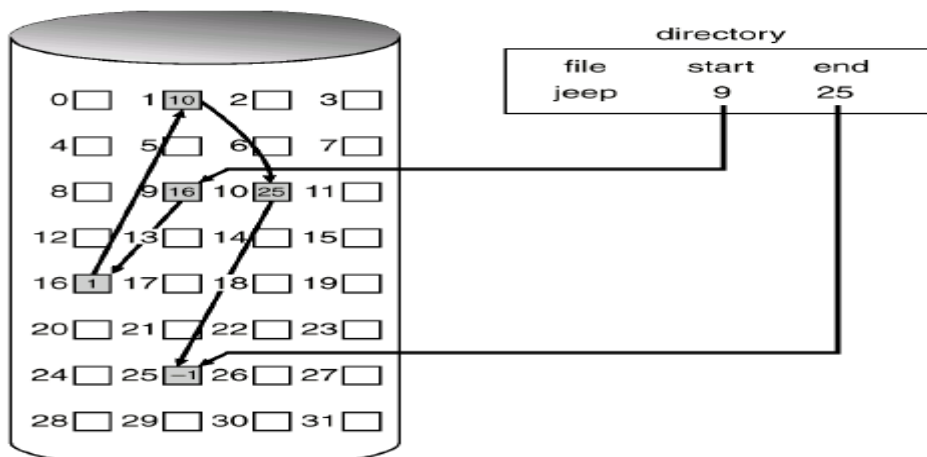


Figure 7.12 Linked Allocation of disk space

Linked allocation solves the problem of external fragmentation, which was present in contiguous allocation. But, still it has a disadvantage: Though it can be effectively used for sequential-access files,



to find *i*th file, we need to start from the first location. That is, random-access is not possible.

**Indexed Allocation**

This method allows direct access of files and hence solves the problem faced in linked allocation. Each file has its own index block, which is an array of disk-block addresses. The *i*th entry in the index block points to the *i*th block of the file. The directory contains the address of the index block as shown in Figure 7.13.

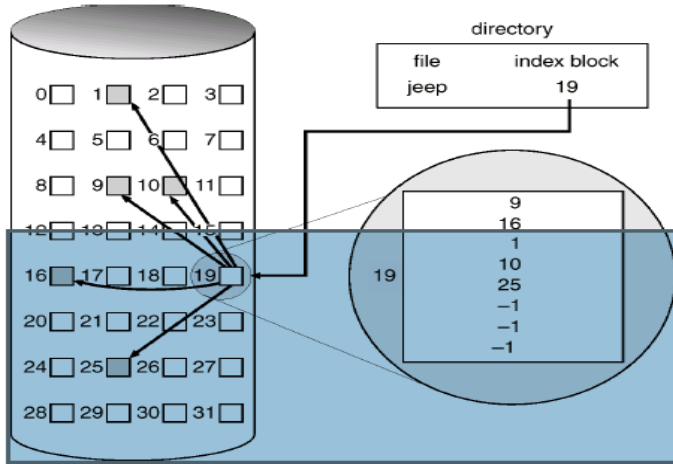


Figure 7.13 Indexed allocation of disk space

10. Describe the methods used for implementing directory structures

10

The file systems of computers can be very large. To manage all these data, we need to organize them. The disks are split into one or more partitions. Typically, each disk on a system contains at least one partition, which is a low-level structure in which files and directories reside. Sometimes, partitions are used to provide several separate areas within one disk, each treated as a separate storage device, whereas other systems allow partitions to be larger than a disk to group disks into one logical structure. Each partition contains information about files within it. This information is kept in a device directory or volume table of contents. The device directory records information—such as name, location, size, and type—for all files on that partition. Figure 7.3 shows file-system organization.

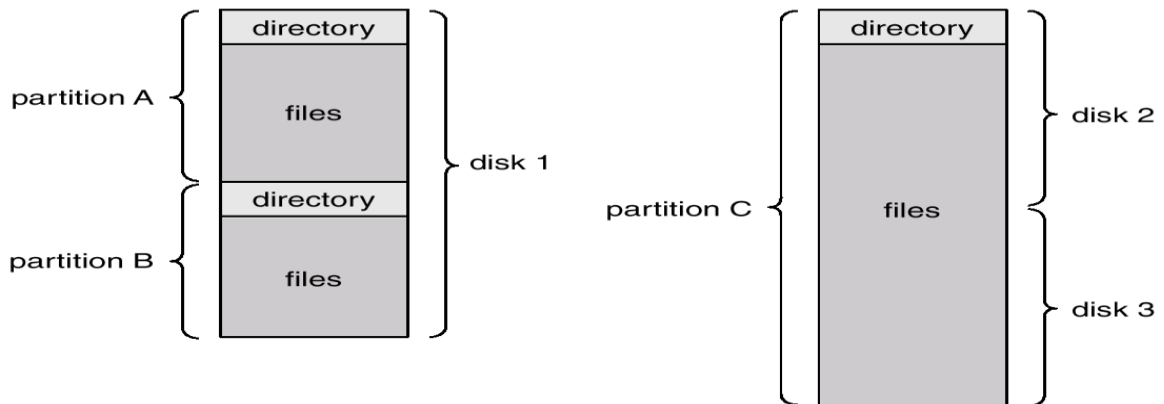


Figure 7.3 A typical file-system organization

Various operations that are to be performed on a directory:

- Search for a file:** We need to be able to search a directory structure to find the entry for

a particular file.

- Create a file:** New files need to be created and added to the directory.
- Delete a file:** When a file is no longer needed, we want to remove it from the directory.
- List a directory:** We need to be able to list the files in a directory, and the contents of the directory entry for each file in the list.
- Rename a file:** As the name of a file represents its contents to its users, the name must be changeable when the contents or use of the file changes. Renaming a file may also allow its position within the directory structure to be changed.
- Traverse the file system:** Files may need to be copied into another storage device like tape, pen-drive etc. and the disk space of that file released for reuse by another file.

### 7.3.1 Single – Level Directory

The simplest directory structure is the single-level directory. All files are contained in the same directory, which is easy to support and understand. It is shown in Figure 7.4. There are certain limitations for single-level directory:

- As all files will be in one directory, each file should have a unique name.
- If there are more users on the same system, each one of them should remember their file names – which is difficult.
- It is not possible to group the related files together.

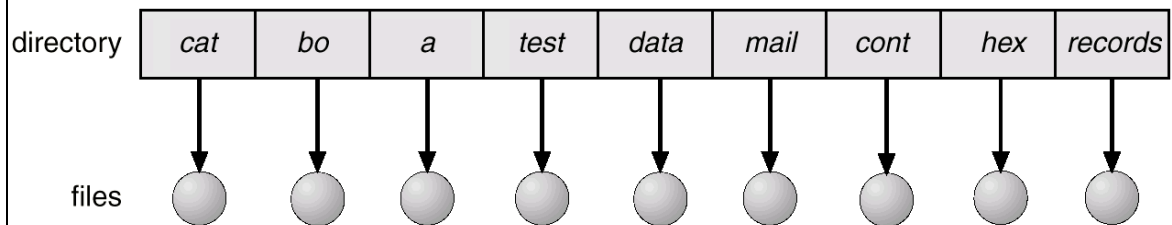


Figure 7.4 Single – level directory structure

### 7.3.2 Two – Level Directory

A solution to the problems faced with single-level directory is to create a separate directory for each user. In the two-level directory structure, each user has his/her own user file directory (UFD). Each UFD has a similar structure, but lists only the files of a single user.

When a user job starts or a user logs in, the system's master file directory (MFD) is searched. The MFD is indexed by user name or account number, and each entry points to the UFD for that user. Two – level directory is shown in Figure 7.5.

When a user refers to a particular file, only his own UFD is searched. Thus, different users may have files with the same name, as long as all the file names within each UFD are unique.

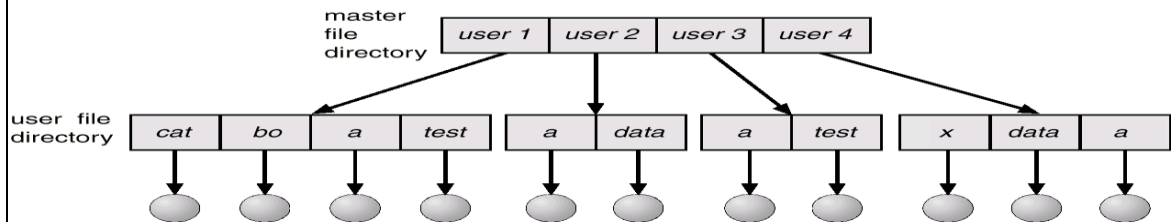


Figure 7.5 Two – level directory structure

### 7.3.3 Tree – Structured Directories

The generalization of two-level directory is a tree structure. Here, users can have any number of directories and sub-directories and files inside each directory as shown in Figure 7.6. Tree – structure is a most common directory structure applied in almost all OS.

Here, every file has a path name. Path names can be of two types: **absolute** and **relative**. An absolute path name is a sequence of directory names starting from root to file name. A relative path name defines a path from current directory. For example, in the tree-structured file system of Figure 7.6, if the current directory is *root/spell/mail*, then the relative path name *prt/first* refers to the same file as does the absolute path name *root/spell/mail/prt/first*.

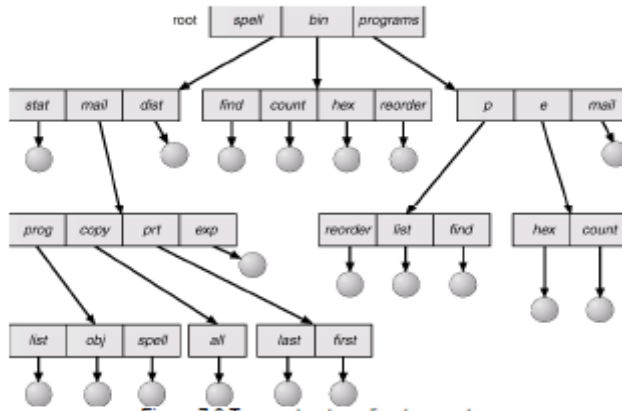


Figure 7.6 Tree – structure directory system

### 7.3.4 Acyclic – Graph Directories

A tree structure prohibits the sharing of files or directories. An acyclic graph allows directories to have shared subdirectories and files (Figure 7.7). The same file or subdirectory may be in two different directories. This is useful when more than one programmers are working on a same project and need to share the files.

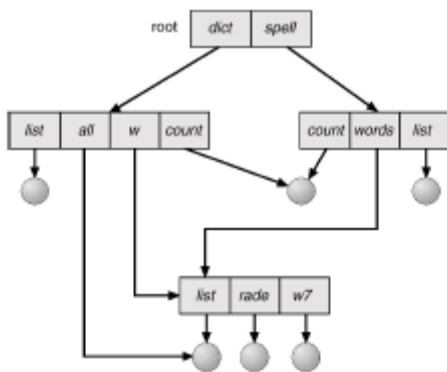


Figure 7.7 Acyclic graph directory structure