

SCHEME OF EVALUATION

Scoring Indicators

Revision : 2015		Course Code : 4134																										
Course Title : Operating Systems																												
Qst. No	Scoring Indicator	Split up Score	Sub Total	Total																								
I	PART - A																											
1.	Assembler, Loader, Compiler, Interpreter, OS (Any two)	2	2	10																								
2.	A thread is a basic unit of CPU utilization or a light weight process	2	2																									
3.	Internal and External fragmentation	2	2																									
4.	A thin client is a lightweight computer that has been optimized for establishing a remote connection with a server-based computing environment	2	2																									
5.	contiguous, linked and indexed allocations	2	2																									
	PART - B																											
II (1)	<p>COMPILERS It is a tool which translate notations from one system to another, usually from source language to target language. A compiler must handle syntax and semantic errors. Depends on its type, a compiler can generate either assembly language program or machine code from high level language program.</p> <p>INTERPRETERS: It is a translation program which converts each program statement into machine code just before the next program statement is to be executed. Translation and execution occur immediately, one after another, one statement at a time.</p> <table border="1"> <thead> <tr> <th>Basics for comparison</th> <th>Compiler</th> <th>Interpreter</th> </tr> </thead> <tbody> <tr> <td>Input</td> <td>Takes an entire program at a time</td> <td>Takes single line of code at a time</td> </tr> <tr> <td>Output</td> <td>It generates intermediate object code</td> <td>It does not produce any intermediate object code</td> </tr> <tr> <td>Working Mechanism</td> <td>Compilation done before execution</td> <td>Compilation and execution take place simultaneously</td> </tr> <tr> <td>Speed</td> <td>Comparatively faster</td> <td>Slower</td> </tr> <tr> <td>Memory</td> <td>More memory required</td> <td>Less memory required</td> </tr> <tr> <td>Errors</td> <td>Displays all errors after Compilation</td> <td>Displays one error at a time</td> </tr> <tr> <td>Error Detection</td> <td>Difficult</td> <td>Easier</td> </tr> </tbody> </table>	Basics for comparison	Compiler	Interpreter	Input	Takes an entire program at a time	Takes single line of code at a time	Output	It generates intermediate object code	It does not produce any intermediate object code	Working Mechanism	Compilation done before execution	Compilation and execution take place simultaneously	Speed	Comparatively faster	Slower	Memory	More memory required	Less memory required	Errors	Displays all errors after Compilation	Displays one error at a time	Error Detection	Difficult	Easier	6	6	6
Basics for comparison	Compiler	Interpreter																										
Input	Takes an entire program at a time	Takes single line of code at a time																										
Output	It generates intermediate object code	It does not produce any intermediate object code																										
Working Mechanism	Compilation done before execution	Compilation and execution take place simultaneously																										
Speed	Comparatively faster	Slower																										
Memory	More memory required	Less memory required																										
Errors	Displays all errors after Compilation	Displays one error at a time																										
Error Detection	Difficult	Easier																										



Figure: Compiler




Figure: Interpreter

II (2) MULTIPROGRAMMING SYSTEMS

One of the most important aspects of OS is the ability to multiprogram. A single user cannot keep either CPU or the I/O devices busy all times. Multiprogramming increases CPU utilization by organizing jobs (code and data) so that the CPU always has one to execute.

The idea is : The OS keeps several jobs in memory simultaneously. This set of jobs can be a subset of jobs kept in the job pool- which contains all jobs that enter the system- since the number of jobs that can be kept simultaneously in memory is usually smaller than the number of jobs that can be kept in the job pool. The OS picks and begins to execute one of the jobs in memory. Eventually, the job may have to wait for some task, such as an I/O operation, to complete. In a non-multiprogrammed system, the CPU would sit idle. In multiprogrammed system, the OS simply switches to, and executes another job. When that job needs to wait, the CPU is switched to another job, and so on. Eventually, the first job finishes waiting and gets the CPU back. As long as at least one job needs to execute, the CPU is never idle.



• Advantages:

- CPU is used most of time and never become idle
- The system looks fast as all the tasks runs in parallel
- Short time jobs are completed faster than long time jobs
- Multiprogramming systems support multiple users
- Response time is shorter

3

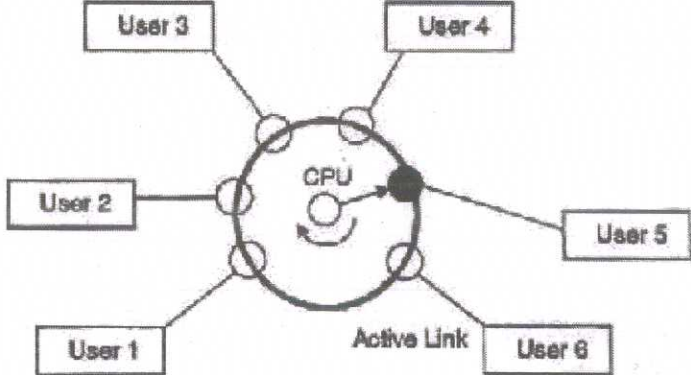
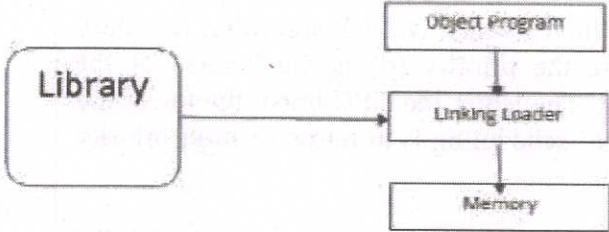
3

6

6

	<p>Disadvantages:</p> <p>It is difficult to program a system because of complicated schedule handling</p> <p>Tracking all tasks/processes is sometimes difficult to handle</p> <p>Due to high load of tasks, long time jobs have to wait long</p>													
II (3)	<p>PROCESS CONTROL BLOCK (PCB)</p> <p>Each process is represented in the OS by a PCB- also called a task control block. It contains many pieces of information associated with a specific process, including these:</p> <ul style="list-style-type: none"> • Process state: The state may be new, ready, running, waiting, halted and so on. • Program counter: The counter indicates the address of the next instruction to be executed for this process. • CPU registers: The registers vary in number and type, depending on the computer architecture. They include accumulators, index registers, stack pointers, and general- purpose registers, plus any condition code information. Along with the program counter, this state information must be saved when an interrupt occurs, to allow the process to be continued correctly afterward. • CPU scheduling information: It includes a process priority, pointers to scheduling queues, and any other scheduling parameters. • Memory management information: It includes information such as value of base and limit registers, the page tables, or the segment tables, depending on the memory system used by the OS. • Accounting information: It includes the amount of CPU and real time used, time limits, account members, job or process numbers, and so on. • I/O status information: It includes the list of I/O devices allocated to the process, a list of open files, and so on. In brief, PCB serves as the repository for any information that may vary from process to process. <table border="1" style="margin-left: 20px;"> <tr><td>process state</td></tr> <tr><td>process number</td></tr> <tr><td>program counter</td></tr> <tr><td> </td></tr> <tr><td>registers</td></tr> <tr><td>memory limits</td></tr> <tr><td>list of open files</td></tr> <tr><td> </td></tr> <tr><td>•</td></tr> <tr><td>•</td></tr> </table>	process state	process number	program counter		registers	memory limits	list of open files		•	•	3	6	6
process state														
process number														
program counter														
registers														
memory limits														
list of open files														
•														
•														
II (4)	<p>A set of processes is in a deadlock state when every process in the set is waiting for an event that can be caused only by another process in the set. A deadlock situation can arise if the following four conditions hold simultaneously in a system:</p> <p>1. Mutual exclusion. At least one resource must be held in a</p>		6	6										

	<p>First fit. Allocate the first hole that is big enough. Searching can start either at the beginning of the set of holes or at the location where the previous first-fit search ended. We can stop searching as soon as we find a free hole that is large enough.</p> <p>Best fit. Allocate the smallest hole that is big enough. We must search the entire list, unless the list is ordered by size. This strategy produces the smallest leftover hole.</p> <p>Worst fit. Allocate the largest hole. Again, we must search the entire list, unless it is sorted by size. This strategy produces the largest leftover hole, which may be more useful than the smaller leftover hole from a best-fit approach.</p>	3																				
II (7)	<p>File is an abstract data type</p> <ul style="list-style-type: none"> • Create • Write - at write pointer location • Read - at read pointer location • Reposition within file - seek • Delete • Truncate • Open(F_i) - search the directory structure on disk for entry F_i, and move the content of entry to memory • Close (F_i) - move the content of entry F_i in memory to directory structure on disk <p>Listing</p> <p>Brief Description.....</p>	3 3	6	6																		
PART - C																						
III(a)	<ul style="list-style-type: none"> • Process Management • Memory Management • Secondary Storage Management • Device Management • File Management • Security and Protection <p>Listing Explanation</p>	4 4	8																			
III(b))	<table border="0"> <tr> <td style="text-align: center;"><u>BASIS FOR COMPARISON</u></td> <td style="text-align: center;"><u>LINUX</u></td> <td style="text-align: center;"><u>WINDOWS</u></td> </tr> <tr> <td>Cost</td> <td>Free of cost</td> <td>Expensive</td> </tr> <tr> <td>Open Source</td> <td>Yes</td> <td>No</td> </tr> <tr> <td>Customizable</td> <td>Yes</td> <td>No</td> </tr> <tr> <td>Security</td> <td>More Secure</td> <td>Vulnerable to viruses</td> </tr> <tr> <td style="text-align: center;">and</td> <td></td> <td>malware attacks</td> </tr> </table>	<u>BASIS FOR COMPARISON</u>	<u>LINUX</u>	<u>WINDOWS</u>	Cost	Free of cost	Expensive	Open Source	Yes	No	Customizable	Yes	No	Security	More Secure	Vulnerable to viruses	and		malware attacks	7 7	7	15
<u>BASIS FOR COMPARISON</u>	<u>LINUX</u>	<u>WINDOWS</u>																				
Cost	Free of cost	Expensive																				
Open Source	Yes	No																				
Customizable	Yes	No																				
Security	More Secure	Vulnerable to viruses																				
and		malware attacks																				

	<p>running.</p> <p>Time sharing requires an interactive computer system, which provides direct communication between the user and the system. A time shared OS allows many users to share the computer simultaneously. This system uses CPU scheduling and multiprogramming to provide each user with a small portion of a time-shared computer.</p>  <p>Advantages: Increases performance Many applications can run at the same time</p> <p>Disadvantages: It consumes much resources so it need special operating systems Switching between tasks becomes sometimes sophisticated as there are lot of users and applications running which may hang up the system Time sharing systems should have high specifications of hardware</p>	2		
IV(b)	<p>Loader is a program which accepts the object program decks, prepares this program for execution by the computer and initializes the execution. It is actually responsible for initiating execution of process.</p> <p>Loader must perform four functions:</p> <ol style="list-style-type: none"> 1. Allocation: Allocate space in memory for the program. 2. Linking: Resolve symbolic references between object decks. 3. Relocation: Adjust all address dependent locations, such as address constants, to correspond to the allocated space. 4. Loading: Physically place the machine instructions and data into  <p>memory.</p>	2 4 7 1		
V (a)	<p>FIRST-COME, FIRST-SERVED (FCFS) SCHEDULING</p> <p>The simplest CPU-scheduling algorithm is the first-come, first-served (FCFS) scheduling algorithm. With this scheme, the process that requests the CPU first is allocated the CPU first. The implementation of the FCFS policy is easily managed with a FIFO</p>	8	15	


queue. When a process enters the ready queue, its PCB is linked onto the tail of the queue. When the CPU is free, it is allocated to the process at the head of the queue. The running process is then removed from the queue. The code for FCFS scheduling is simple to write and understand.

The average waiting time under the FCFS policy, however, is often quite long. Consider the following set of processes that arrive at time with the length of the CPU burst given in milliseconds

Process	Burst Time
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

Priority Scheduling

A priority is associated with each process, and the CPU is allocated to the process with the highest priority. Equal-priority processes are scheduled in FCFS order. An SJF algorithm is simply a priority algorithm where the priority (p) is the inverse of the (predicted) next CPU burst. The larger the CPU burst, the lower the priority, and vice versa. The scheduling is in terms of **high** priority and **low** priority.

Process	Burst Time	Priority	As an example, consider the following set of processes, assumed to have arrived at time 0 in the order P_1, P_2, P_3, P_4, P_5 , with the length of the CPU burst given in milliseconds:
P_1	10	3	
P_2	1	1	
P_3	2	4	
P_4	1	5	
P_5	5	2	

Using priority scheduling, we would schedule these processes according to the following Gantt chart:



Average waiting time = 8.2 msec

V (b)

A **cooperating process** is one that can affect or be affected by other processes executing in the system. Cooperating processes can either directly share a logical address space

Concurrent access to shared data may result in data inconsistency

Maintaining data consistency requires mechanisms to ensure the orderly execution of cooperating processes.

Illustration of the problem:

Suppose that we wanted to provide a solution to the consumer-producer problem that fills **all** the buffers. We can do so by having an integer **counter** that keeps track of the number of full buffers. Initially, **counter** is set to 0. It is incremented by the producer after it produces a new buffer and is decremented by the consumer after it consumes a buffer.

Producer

```
while(true){
    /* produce an item in next produced */

    while (counter == BUFFER_SIZE);
        /* do nothing */
    buffer[in] = next_produced;
    in = (in + 1) % BUFFER_SIZE;
    counter++;
}
```

Consumer

```
while (true) {
    while (counter == 0)
        /* do nothing */
    next_consumed = buffer[out];
    out = (out + 1) % BUFFER_SIZE;
    counter--;
    /* consume the item in next consumed */
}
```


1.5

1.5

1.5

7

	<p>The statement “counter++” may be implemented in machine language (on a typical machine) as follows:</p> <pre> register1=counter register1=register1+1 counter = register1 </pre> <p>where <i>register₁</i> is one of the local CPU registers. Similarly, the statement “counter--” is implemented as follows:</p> <pre> register2=counter register2=register2-1 counter = register2 </pre> <p>Consider this execution interleaving with “count = 5” initially:</p> <p>S0: producer execute register1 = counter {register1 = 5} S1: producer execute register1 = register1 + 1 {register1 = 6} S2: consumer execute register2 = counter {register2 = 5} S3: consumer execute register2 = register2 – 1 {register2 = 4} S4: producer execute counter = register1 {counter = 6 } S5: consumer execute counter = register2 {counter = 4}</p> <p>Notice that we have arrived at the incorrect state “counter == 4”, indicating that four buffers are full, when, in fact, five buffers are full. If we reversed the order of the statements at <i>T₄</i> and <i>T₅</i>, we would arrive at the incorrect state “counter == 6”.</p> <p>We would arrive at this incorrect state because we allowed both processes to manipulate the variable counter concurrently. A situation like this, where several processes access and manipulate the same data concurrently and the outcome of the execution depends on the particular order in which the access takes place, is called a race condition.</p>	1.5		
VI (a)	<p>METHODS FOR HANDLING DEADLOCKS</p> <ul style="list-style-type: none"> • Deadlock Prevention <ul style="list-style-type: none"> ➤ Eliminate Mutual Exclusion ➤ Eliminate Hold and wait ➤ Eliminate No Preemption ➤ Eliminate Circular Wait • Deadlock Avoidance <p>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. A resource allocation graph is generally used to avoid deadlocks. A claim edge in resource allocation graph denotes that a request may be made in future and is represented as a dashed line. Based on claim 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. In the following RAG If R2 is allocated to p2 and if P1 request for R2, there will be a deadlock.</p>	2 2	8	15



Banker's Algorithm

Banker's Algorithm is resource allocation and deadlock avoidance algorithm which test all the request made by processes for resources. It check for safe state, after granting request, if the system remains in the safe state then it allows the request and if there is no safe state it don't allow the request made bythe process.

- **Deadlock Detection**

1.If resources have single instance: In this case for deadlock detection we can run an algorithm to check for cycle in the Resource Allocation Graph. Presence of cycle in the graph is the sufficient condition for deadlock in system with resources having single instance.

2. If there are multiple instances of resources:

Detection of cycle is necessary but not sufficient condition for deadlock detection with multiple resource instances , in this case system may or may not be in deadlock varies according to different situations.

- **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.

Recovery methods

1. Killing the process.

Killing all the process involved in deadlock. Killing process one by one. After killing each process check for deadlock again keep repeating process till system recover from deadlock. disadvantage: lot of overhead; must re-run algorithm after each kill

2. Resource Preemption

Resources are preempted from the processes involved in deadlock, preempted resources areallocated to other processes, so that there is a possibility of recovering the system from deadlock. In thiscase system go into starvation.

2

2

VI(b)

Schedulers

A process migrates among the various scheduling queues throughout its lifetime. The operating system must select, for scheduling purposes, processes from these queues in some fashion. The selection process is carried out by the appropriate scheduler.

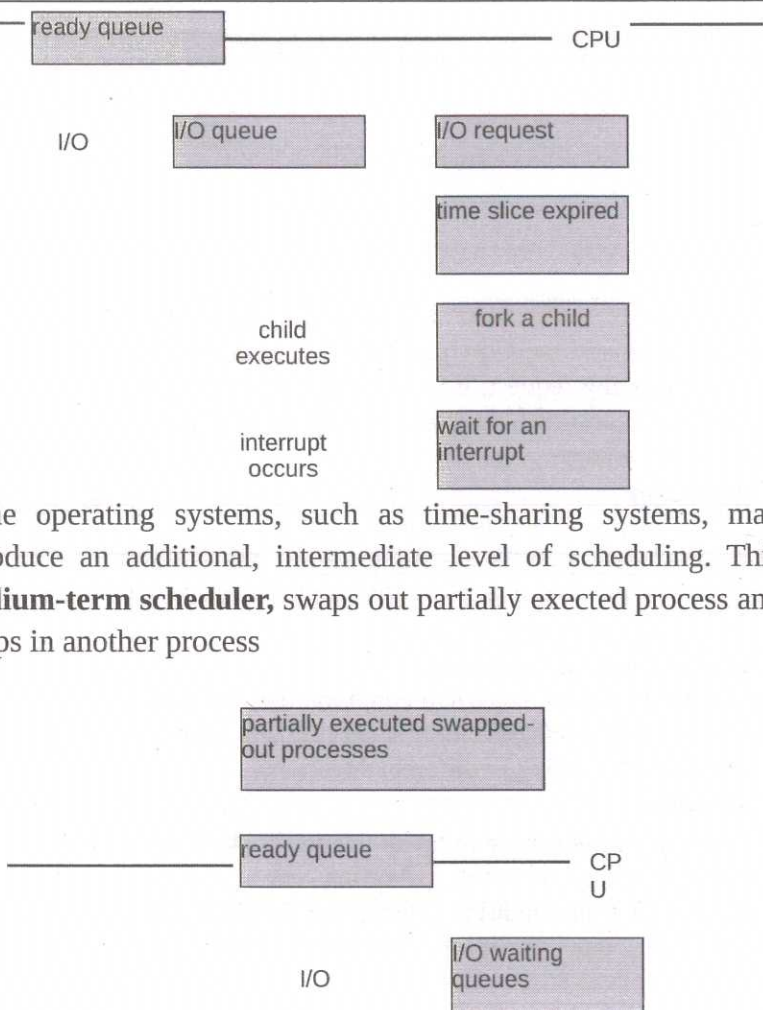
The **long-term scheduler**, or job scheduler, selects processes from this pool and loads them into memory for execution.

The **short-term scheduler**, or CPU scheduler, selects from among the processes that are ready to execute and allocates the CPU to one of them.

1

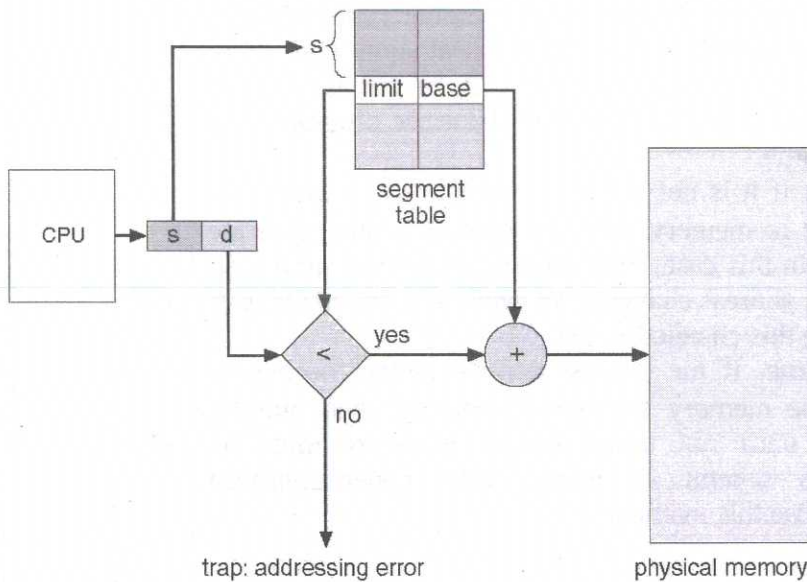
1

7

	 <p>Some operating systems, such as time-sharing systems, may introduce an additional, intermediate level of scheduling. This medium-term scheduler, swaps out partially executed process and swaps in another process</p>	2	1																																														
VII(a)	<p>Reference String</p> <p>7, 0, 1, 2, 0, 3, 0, 4, 2, 3, 0, 3, 2, 1, 2, 0, 1, 7, 0, 1</p> <p><u>FIFO Page Replacement</u></p> <p>The simplest page-replacement algorithm is a first-in, first-out (FIFO) algorithm. A FIFO replacement algorithm associates with each page the time when that page was brought into memory. When a page must be replaced, the oldest page is chosen.</p> <p>7, 0, 1, 2, 0, 3, 0, 4, 2, 3, 0, 3, 2, 1, 2, 0, 1, 7, 0, 1</p> <table border="1" data-bbox="271 1747 957 1948"> <tr> <td>7</td><td>7</td><td>7</td><td>2</td> <td>2</td><td>2</td><td>4</td><td>4</td><td>4</td><td>0</td> <td>0</td><td>0</td> <td>7</td><td>7</td><td>7</td> </tr> <tr> <td></td><td>0</td><td>0</td><td>0</td> <td>3</td><td>3</td><td>3</td><td>2</td><td>2</td><td>2</td> <td>1</td><td>1</td> <td>1</td><td>0</td><td>0</td> </tr> <tr> <td></td><td></td><td>1</td><td>1</td> <td>1</td><td>0</td><td>0</td><td>0</td><td>3</td><td>3</td> <td>3</td><td>2</td> <td>2</td><td>2</td><td>1</td> </tr> </table> <p>No of page</p> <p>faults : 15</p>	7	7	7	2	2	2	4	4	4	0	0	0	7	7	7		0	0	0	3	3	3	2	2	2	1	1	1	0	0			1	1	1	0	0	0	3	3	3	2	2	2	1	2	9	15
7	7	7	2	2	2	4	4	4	0	0	0	7	7	7																																			
	0	0	0	3	3	3	2	2	2	1	1	1	0	0																																			
		1	1	1	0	0	0	3	3	3	2	2	2	1																																			

VIII(a)
)

Segmentation is a memory-management scheme that supports the programmer view of memory. A logical address space is a collection of segments.
Segmentation Hardware



A 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

1

9

15

4

1

1

The segment number is used as an index to the segment table. The offset d of the logical address must be between 0 and the segment limit. If it is not, we trap to the operating system (logical addressing attempt beyond end of segment). When an offset is legal, it is added to the segment base to produce the address in physical memory of the desired byte. The segment table is thus essentially an array of base– limit register pairs.

2

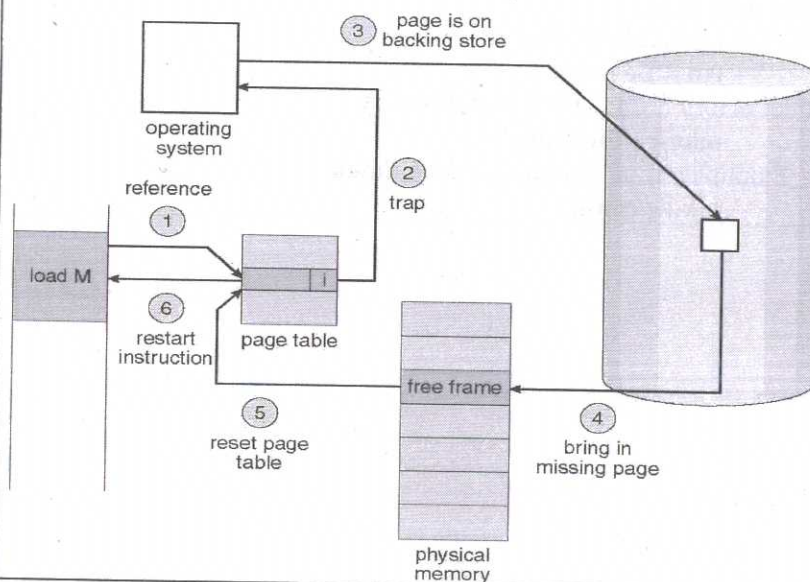
VIII(b)
)

If there is a reference to a page, first reference to that page will trap to operating system. It is called a page fault.

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 an 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, we now page it in.
3. We find a free frame (by taking one from the free-frame list, for example).
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 trap. The process can now access the page as though it had always been in memory.

3


6



3

1. Single level Directory

A single directory for all users

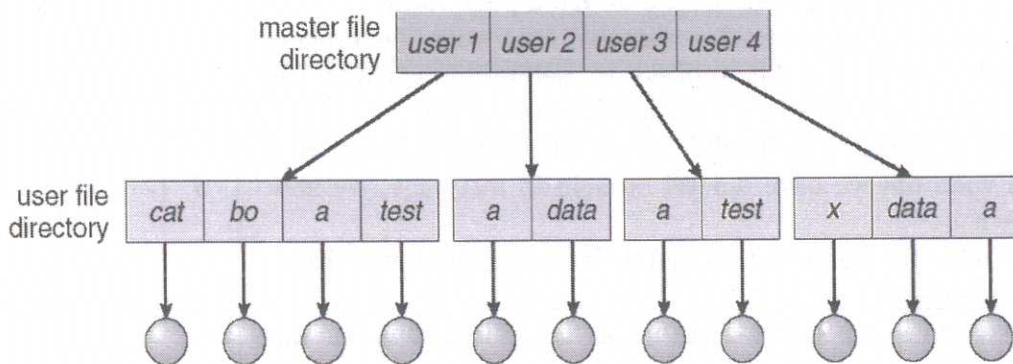


3

Naming problem
Grouping problem

2. Two level Directory

Separate directory for each user



3

Path name
Can have the same file name for different user
Efficient searching
No grouping capability

3. Tree structured Directories

Absolute or **relative** path name

Creating a new file is done in current directory

Delete a file **rm <file-name>**

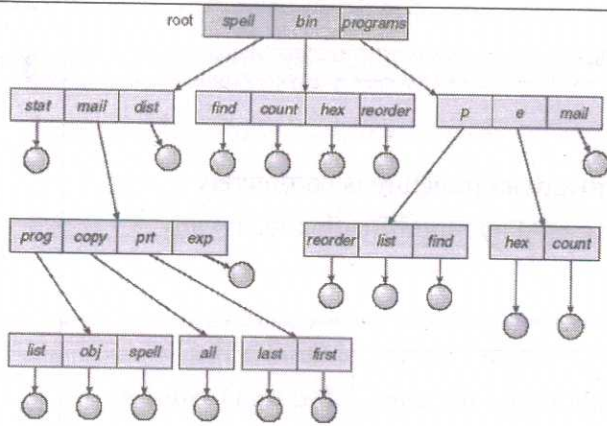
Creating a new subdirectory is done in current directory

mkdir <dir-name>

Example: if in current directory **/mail**

mkdir count

3



IX (b) Hardware virtualization

6

The most common type of virtualization today, hardware virtualization is made possible by a virtual machine manager (VM) called the “hypervisor”. The hypervisor creates virtual versions of computers and operating systems and consolidates them into one large physical server, so that all the hardware resources can be utilized more efficiently. It also enables users to run different operating systems on the same machine simultaneously.

It involves embedding virtual machine software into the server's hardware components. That software is called the hypervisor. The hypervisor manages the shared physical hardware resources between the guest OS & the host OS. The abstracted hardware is represented as actual hardware. Virtualization means abstraction & hardware virtualization is achieved by abstracting the physical hardware part using Virtual Machine Monitor (VMM) or hypervisor.

The term hardware virtualization is used when VMM or virtual machine software or any hypervisor gets directly installed on the hardware system. The primary task of the hypervisor is to process monitoring, memory & hardware controlling. After hardware virtualization is done, different operating systems can be installed, and various applications can run on it. Hardware virtualization, when done for server platforms, is also called server virtualization.

3

Advantages of hardware virtualization

- Lower Cost:
- Efficient resource utilization:
- Increase IT flexibility:
- Advanced Hardware Virtualization features:

Types of hardware virtualization

Hardware virtualization is of three kinds.

3

These are:

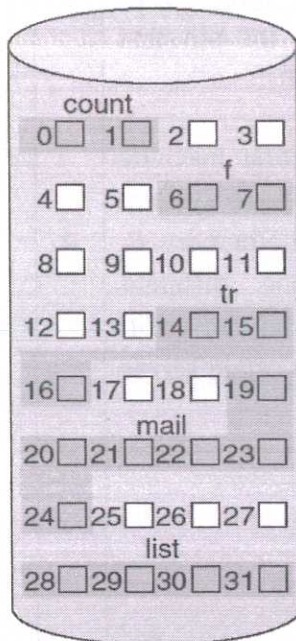
- Full Virtualization: Here the hardware architecture is completely simulated. Guest software doesn't need any modification to run any applications.
- Emulation Virtualization: Here the virtual machine simulates the hardware & is independent. Furthermore, the guest OS doesn't require any modification.
- Para-Virtualization: Here, the hardware is not simulated; instead the guest software runs its isolated system.

X File Allocation methods
(a) 1. Contiguous

2. Linked
3. Indexed

Contiguous allocation requires that each file occupy a set of contiguous blocks on the disk. Disk addresses define a linear ordering on the disk. Contiguous allocation of a file is defined by the disk address and length (in block units) of the first block.

Accessing a file that has been allocated contiguously is easy



directory

file	start	length
count	0	2
tr	14	3
mail	19	6
list	28	4
f	6	2

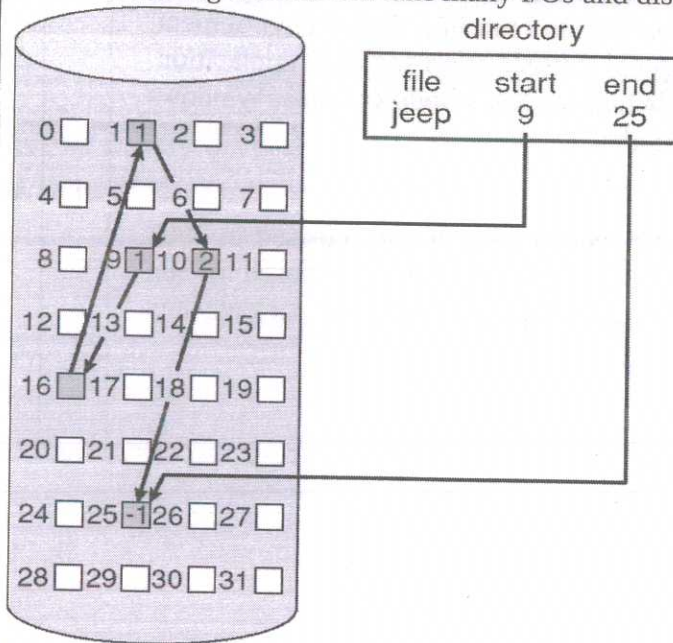
3

9 15

Linked allocation – each file a linked list of blocks

- File ends at nil pointer
- No external fragmentation
- Each block contains pointer to next block
- No compaction, external fragmentation
- Free space management system called when new block needed
- Improve efficiency by clustering blocks into groups but increases internal fragmentation

- Reliability can be a problem
- Locating a block can take many I/Os and disk seeks



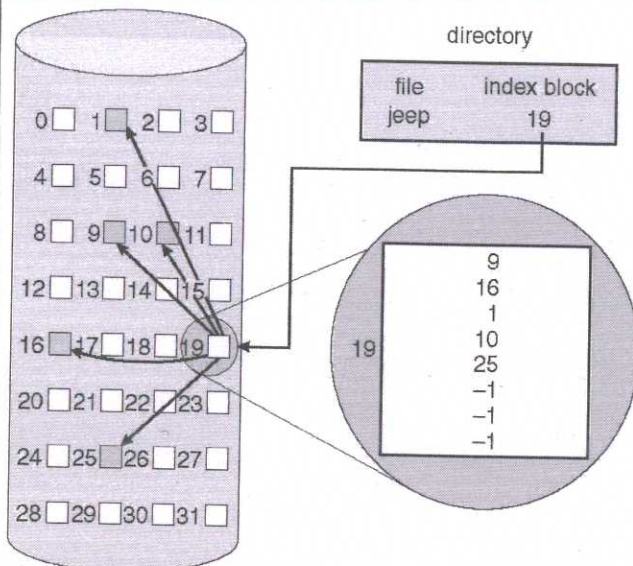
3

Each file is a linked list of disk blocks: blocks may be scattered anywhere on the disk. The directory contains a pointer to the first and last blocks of the file.

Indexed Allocation

Linked allocation solves the external-fragmentation and size-declaration problems of contiguous allocation. However, in the absence of a FAT, linked allocation cannot support efficient direct access, since the pointers to the blocks are scattered with the blocks themselves all over the disk and must be retrieved in order. **Indexed allocation** solves this problem by bringing all the pointers together into one location: the **index block**.

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. To find and read the *i*th block, we use the pointer in the *i*th index-block entry.



3

<p>X (b)</p>	<p>A VirtualBox or VB is a software virtualization package that installs on an operating system as an application. VirtualBox allows additional operating systems to be installed on it, as a Guest OS, and run in a virtual environment. In 2010, VirtualBox was the most popular virtualization software application. Supported operating systems include Windows XP, Windows Vista, Windows 7, macOS X, Linux, Solaris, and OpenSolaris.</p> <p>VirtualBox was originally developed by Innotek GmbH and released in 2007 as an open-source software package. The company was later purchased by Sun Microsystems. Oracle Corporation now develops the software package and titles it Oracle VM VirtualBox.</p>	<p>6</p>	<p>6</p>	
--------------------------------	---	----------	----------	--