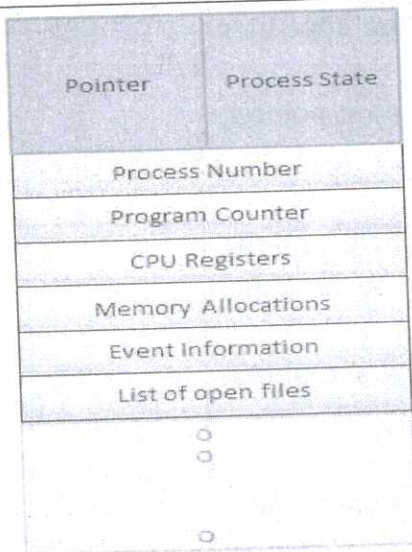


SCORING INDICATOR – (15)4134 OPERATING SYSTEMS

<p>I</p> <p>1</p> <p>2</p> <p>3</p> <p>4</p> <p>5</p>	<p><b>PART A</b></p> <p>System software is a type of computer program that is designed to run a computer's hardware and application programs.</p> <p>A thread is a single sequence stream within a process. Because threads have some of the properties of processes, they are sometimes called lightweight processes. the basic unit of CPU utilization is a thread. A thread has or consists of a program counter (PC), a register set, and a stack space</p> <p>Process Management, Memory Management, File-System Management, Mass-storage Management, Protection and Security, User-interface etc etc</p> <p>A process is thrashing if it is spending more time on paging than executing.</p> <p>A thin client is a lightweight computer that is built for remoting into a server. It depends heavily on another computer (its server) to fulfil its computational roles.</p>	<p>2</p> <p>Any 2 x 1</p> <p>Any 4 x ½</p> <p>2</p> <p>2</p>	<p>10</p>
<p>II</p> <p>1</p>	<p>Both are language translators that translate high level language to machine language. Both are machine independent.</p> <p>A compiler is a computer program (or a set of programs ) that transforms source code written in a programming language (the source language) into another computer language(the target language), with the latter often having a binary form known as object code</p> <p>Features</p> <p>Compiler Takes entire program as input</p> <p>Intermediate object code is generated</p> <p>Memory requirement : More(Since object code is generated)</p> <p>Program need not be compiled every time</p> <p>Errors are displayed after entire program is checked</p> <p>Faster execution</p>	<p>3+3</p>	<p>6</p>

	<p>Eg: C compiler</p> <p>Interpreter Interpreter takes single instruction as input          No intermediate object code is generated          Memory requirement is less          Every time higher level program is converted into lower level program          Errors are displayed for every instruction interpreted(if any)          Line by line error checking, translation and execution.          Slower in execution compared to compilers          Interpreters are mainly used for training purpose and initial stage of development, for easily locating errors. But once the development is completed, compiler can be used.          Eg: BASIC</p>		
2	<p><b>File Operations</b></p> <p>A file is an <b>abstract data type</b>.</p> <p>The operating system can provide system calls to create, write, read, reposition, delete, and truncate files.</p> <ul style="list-style-type: none"> <li>• <b>Creating a file.</b></li> <li>• Two steps are necessary to create a file.</li> <li>• First, space in the file system must be found for the file.</li> <li>• Second, an entry for the new file must be made in the directory.</li> <li>• <b>Writing a file.</b></li> <li>• To write a file, we make a system call specifying both the name of the file and the information to be written to the file.</li> <li>• Given the name of the file, the system searches the directory to find the file's location.</li> <li>• The system must keep a <b>write pointer</b> to the location in the file where the next write is to take place.</li> <li>• The write pointer must be updated whenever a write occurs.</li> <li>• <b>Reading a file.</b></li> <li>• To read from a file, we use a system call that specifies the name of the file and where (in memory) the next block of the file should be put.</li> <li>• Again, the directory is searched for the associated entry, and the system needs to keep a read pointer to the location in the file where the next read is to take place.</li> <li>• Once the read has taken place, the read pointer is updated.</li> <li>• Because a process is usually either reading from or writing</li> </ul>	6 x 1	6





Process state. Shows the process state such as 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. Used in context switching

CPU-scheduling information. This information includes a process priority, **pointers** to scheduling queues, and any other scheduling parameters.

Memory-management information. This information may include such information as the value of the base and limit registers, the page tables, or the segment tables, depending on the memory system used by the operating system

Accounting information. This information includes the amount of CPU and real time used, time limits, account numbers, job or process numbers, and so on.

I/O status information. This information includes the list of I/O devices allocated to the process, a list of open files, and so on. Useful in context switching

4

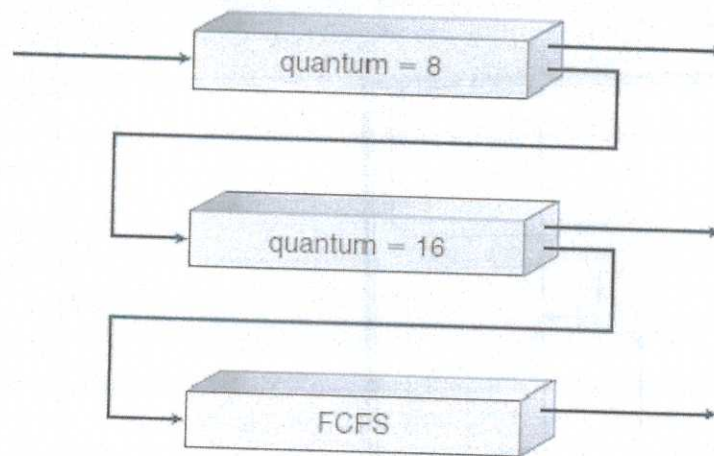
- The multilevel feedback queue scheduling algorithm partitions the ready queue into several separate queues and allows a process to move between queues.
- The idea is to separate processes according to the characteristics of their CPU bursts.
- If a process uses too much CPU time, it will be moved to a lower-priority queue.
- This scheme leaves I/O-bound and interactive processes in the higher-priority queues.
- In addition, a process that waits too long in a lower-priority queue

Fig : 2

Explanation : 4

6

may be moved to a higher-priority queue.  
 → This form of aging prevents starvation.



→ In general, a multilevel feedback queue scheduler is defined by the following parameters:

- The number of queues
- The scheduling algorithm for each queue
- The method used to determine when to upgrade a process to a higher priority queue
- The method used to determine when to demote a process to a lower priority queue
- The method used to determine which queue a process will enter when that process needs service.

5 The binding of instructions and data to memory addresses can be done at the following steps:

- **Compile time:** If it is known at compile time where the process will reside in memory, then absolute code can be generated. For example, if it is known that a user process resides starting at location R, then the generated compiler code will start at that location and extend up from there. If, at some later time, the starting location changes, then it will be necessary to recompile this code.
- **Load time:** If it is not known at compile time where the process will reside in memory, then the compiler must generate relocatable code. In this case, final binding is delayed until load time. If the starting address changes, we need only to reload the user code to incorporate this changed value.
- **Execution time:** If the process can be moved during its execution from one memory segment to another, then binding must be

2 m x 3

6

delayed until run time. Special hardware must be available for this scheme to work. Most general-purpose operating systems use this method

6

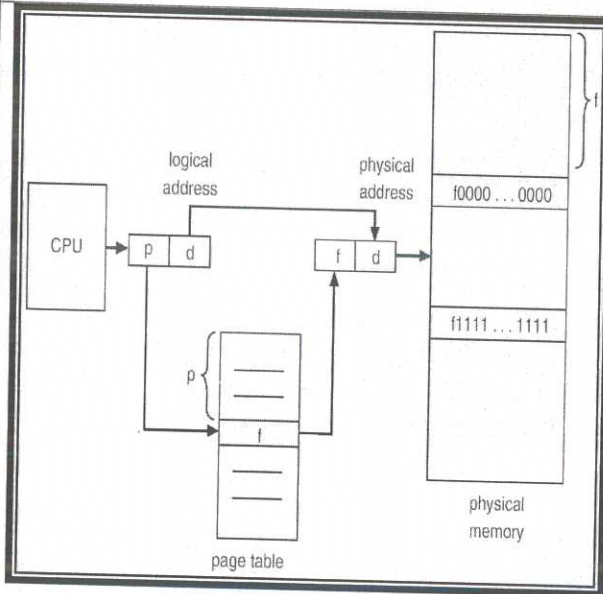


Fig : 3

Explain : 3

Every address generated by the CPU is divided into two parts: a page number (**p**) and a page offset (**d**). The page number is used as an index into a page table. The page table contains the base address of each page in physical memory. This base address is combined with the page offset to define the physical memory address that is sent to the memory unit.

7 **Need of virtualization**

✓ Virtualization can help companies maximize the value of IT investments, decreasing the server hardware footprint, energy consumption, and cost and complexity of managing IT systems while increasing the flexibility of the overall environment.

Virtualization can affect

- Cost
- Administration
- Fast Deployment
- Reduced Infrastructure Costs

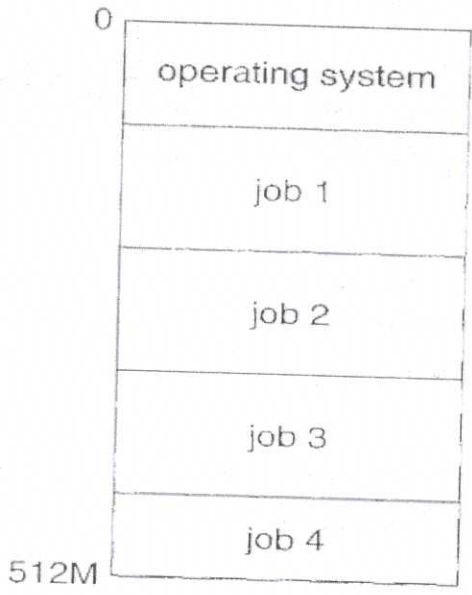
Explain each

6

1 ½ x 4

III  
a In multiprogramming OS several jobs resides in main memory. One job is brought to CPU and executes, and when this job has to wait for some task, like an I/O operation, OS switches to another job in memory and executes it in CPU, and so on. CPU utilization increases

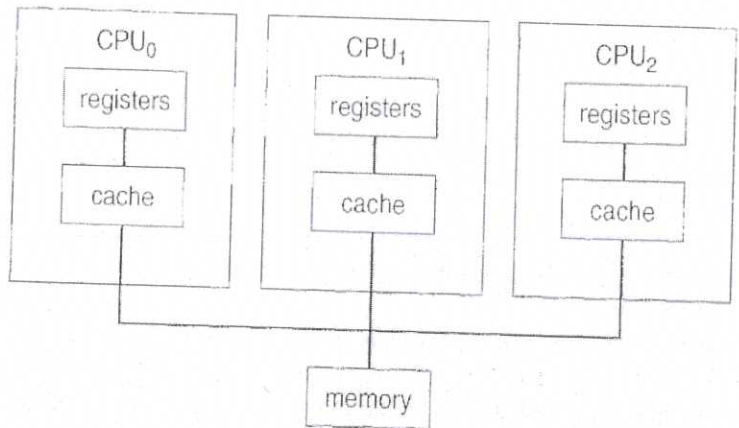
2 ½



1 ½

In multiprocessing OS, OS supports execution of multiple processes in multiple processors or different portions of a single process in multiple processors. Presence of multiple processors can increase the addressable memory

2 ½



1 ½

b Assembler is used to translate the assembly language program in to machine language so that system can understand and execute it.

Easier to understand, easy to correct and modify program instructions, same efficiency as machine language.- machine dependent

Assembler's functions

- Convert mnemonic operation codes to their machine language equivalents





Resource instances:

- o One instance of resource type  $R_1$
- o Two instances of resource type  $R_2$

- o One instance of resource type  $R_3$
- o Three instances of resource type  $R_4$

Process states:

- o Process  $P_1$  is holding an instance of resource type  $R_2$  and is waiting for an instance of resource type  $R_1$ .
- o Process  $P_2$  is holding an instance of  $R_1$  and an instance of  $R_2$  and is waiting for an instance of  $R_3$ .
- o Process  $P_3$  is holding an instance of  $R_3$ .

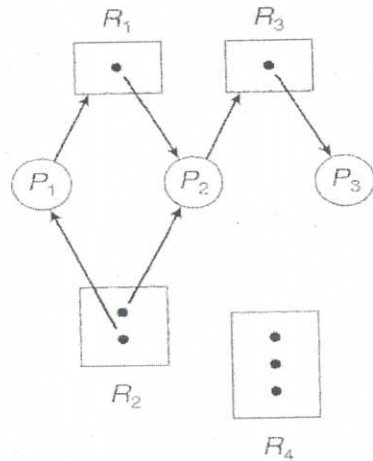


Figure 7.2 Resource-allocation graph.

If each resource type has exactly one instance, then a cycle implies that a deadlock has occurred. If the cycle involves only a set of resource types, each of which has only a single instance, then a deadlock has occurred. Each process involved in the cycle is deadlocked. In this case, a cycle in the graph is both a necessary and a sufficient condition for the existence of deadlock. If each resource type has several instances, then a cycle does not necessarily imply that a deadlock has occurred. In this case, a cycle in the graph is a necessary but not a sufficient condition for the existence of deadlock.

VI a **Critical section** is a segment of code in which the process may be changing common variables, updating tables, writing files so on. **Critical-section problem** is to design a protocol that the process can use to cooperate.

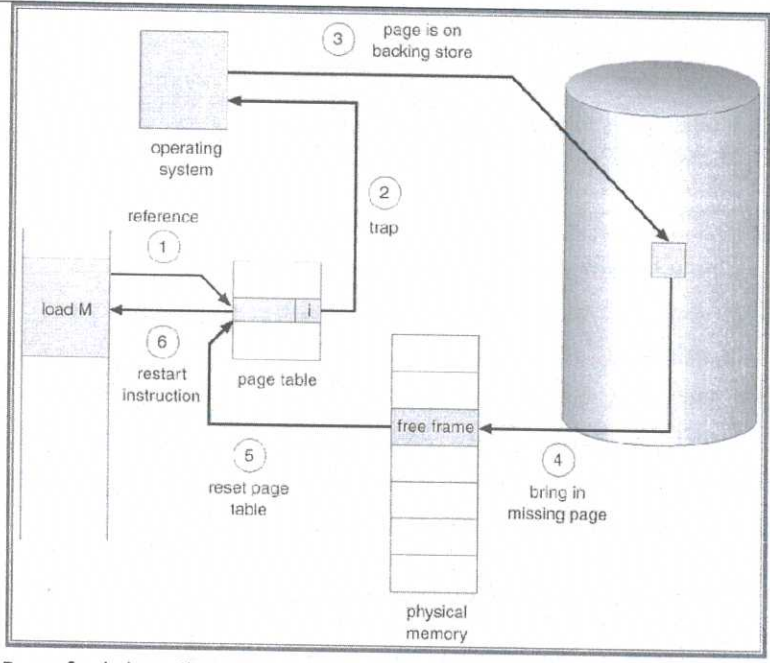
- each process must request permission to enter its critical section.
- entry section is the section of code implement this request.
- exit section follows entry section.
- remainder section contains remaining code.

5

2

	<p>Do {</p> <p style="padding-left: 40px;">[ entry section]</p> <p style="padding-left: 40px;">critical section</p> <p style="padding-left: 40px;">[ exit section]</p> <p style="padding-left: 80px;">remainder section</p> <p>} while (true);</p> <p>A solution to the critical-section problem must satisfy the following three requirements:</p> <p>1. <b>Mutual exclusion.</b> If process <math>P_i</math> is executing in its critical section, then no other processes can be executing in their critical sections.</p> <p>2. <b>Progress.</b> If no process is executing in its critical section and some processes wish to enter their critical sections, then only those processes that are not executing in their remainder sections can participate in deciding which will enter its critical section next, and this selection cannot be postponed indefinitely.</p> <p>3. <b>Bounded waiting.</b> There exists a bound, or limit, on the number of times that other processes are allowed to enter their critical sections after a process has made a request to enter its critical section and before that request is granted</p>	3										
b	<p><u>SJF</u></p> <table border="1" style="margin-left: 20px;"> <tr> <td style="width: 20px; height: 20px;">P1</td> <td style="width: 20px; height: 20px;">P2</td> <td style="width: 20px; height: 20px;">P4</td> <td style="width: 20px; height: 20px;">P3</td> </tr> </table> <p>0                      8                      12                      17</p> <p>26</p> <p>Waiting time <math>P_1 = 0</math>, Waiting time <math>P_2 = 7</math>, Waiting time <math>P_3 = (17-2) = 15</math></p> <p>Waiting time <math>P_4 = (12-3) = 9</math></p> <p>Average waiting time = <math>(0+7+15+9)/4 = 31/4 =</math></p> <p><u>SRTN</u></p> <table border="1" style="margin-left: 20px;"> <tr> <td style="width: 20px; height: 20px;">P1</td> <td style="width: 20px; height: 20px;">P2</td> <td style="width: 20px; height: 20px;">P4</td> <td style="width: 20px; height: 20px;">P1</td> <td style="width: 20px; height: 20px;">P3</td> </tr> </table> <p>0    1                      5                      10                      17</p> <p>26</p> <p>Waiting time <math>P_1 = 0 + (10-1) = 9</math>, Waiting time <math>P_2 = 0</math>, Waiting time <math>P_3 = 17 -</math></p>	P1	P2	P4	P3	P1	P2	P4	P1	P3	<p>2 ½</p> <p>2</p> <p>2 ½</p>	10
P1	P2	P4	P3									
P1	P2	P4	P1	P3								





Page fault handling

2 1/2

VIII  
a

Optimal PRA

7	0	2	03	04	230	32207	01
7	7	7	3	3	3	3	1
	0	0	0	4	0	0	0
		2	2	2	2	7	7

Page faults =8

LRU PRA

7	0	2	03	04	2	3	0	32207	01	
7	7	7	3	3	2	2	2		2	1
	0	0	0	0	0	3	3		7	7
		2	2	4	4	4	0		0	0

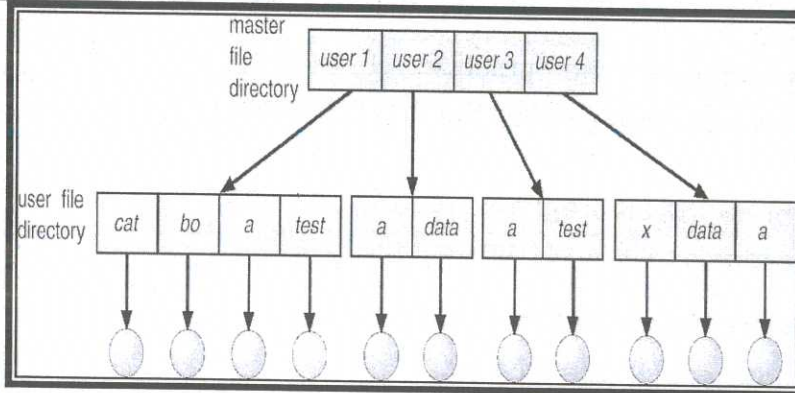
Page faults =10

10

5

5



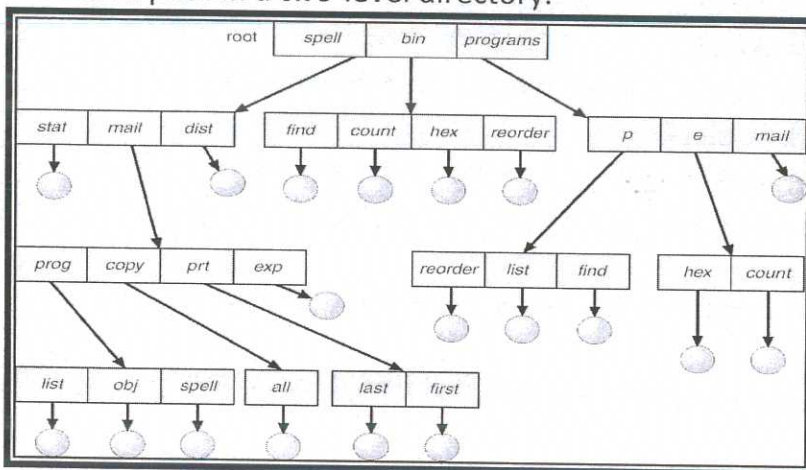


1

### Tree-Structured Directories

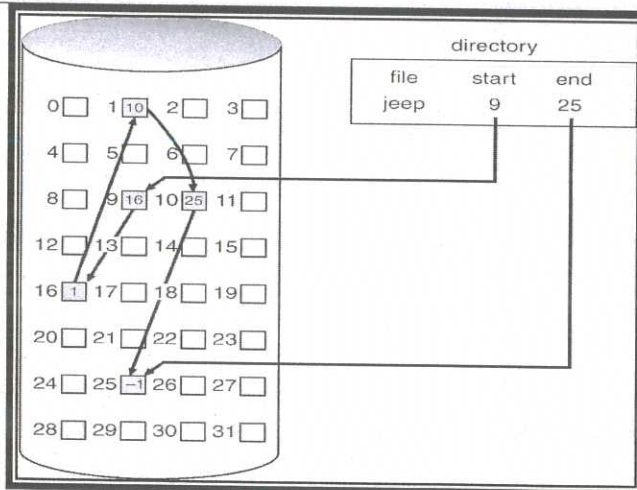
- A tree is the most common directory structure.
- The tree has a root directory, and every file in the system has a unique path name.
- A directory (or subdirectory) contains a set of files or subdirectories.
- A directory is simply another file, but it is treated in a special way.
- All directories have the same internal format. One bit in each directory entry defines the entry as a file (0) or as a subdirectory (1).
- Special system calls are used to create and delete directories.
- With a tree-structured directory system, users can be allowed to access, in addition to their files, the files of other users.
- A path to a file in a tree-structured directory can be longer than a path in a two-level directory.

2



1

b	<p><b>Software Virtualization</b></p> <ul style="list-style-type: none"> <li>✓ Software virtualization is the practice of running software from a remote server rather than on the user's computer.</li> <li>✓ Dynamic link library (DLL) programs redirect all the virtualized application's calls to the server's file system.</li> </ul> <p><b>Advantages of Software Virtualization</b></p> <ol style="list-style-type: none"> <li>1) Client Deployments Become Easier</li> <li>2) Easy to manage</li> <li>3) Software Migration is easier</li> </ol>	3	6
X a	<p><b>Linked Allocation</b></p> <ul style="list-style-type: none"> <li>• With linked allocation, each file is a linked list of disk blocks; the disk blocks may be scattered anywhere on the disk.</li> <li>• The directory contains a pointer to the first and last blocks of the file.</li> <li>• Each block contains a pointer to the next block. These pointers are not made available to the user.</li> <li>• To create a new file, we simply create a new entry in the directory.</li> <li>• With linked allocation, each directory entry has a pointer to the first disk block of the file.</li> <li>• This pointer is initialized to nil (the end-of-list pointer value) to signify an empty file. The size field is also set to 0.</li> <li>• A write to the file causes the free-space management system to find a free block, and this new block is written to and is linked to the end of the file.</li> <li>• To read a file, we simply read blocks by following the pointers from block to block.</li> <li>• There is no external fragmentation with linked allocation, and any free block on the free-space list can be used to satisfy a request.</li> <li>• The size of a file need not be declared when that file is created.</li> <li>• A file can continue to grow as long as free blocks are available.</li> </ul>	<p>Explanation : 2 ½</p> <p>Fig : 2</p>	9

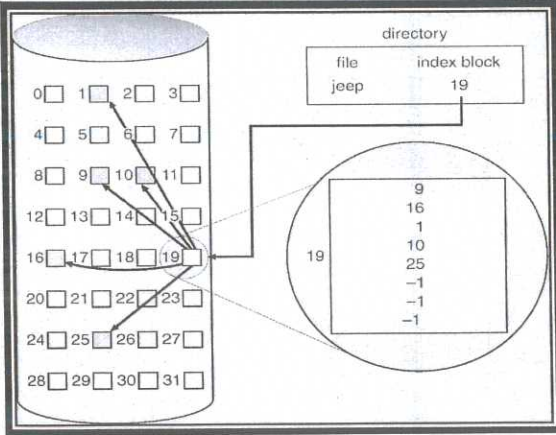


### Indexed Allocation

- Each file has its own index block, which is an array of disk-block addresses.
- The  $i^{\text{th}}$  entry in the index block points to the  $j^{\text{th}}$  block of the file.
- The directory contains the address of the index block.
- To find and read the  $i^{\text{th}}$  block, we use the pointer in the  $j^{\text{th}}$  index-block entry.
- When the file is created, all pointers in the index block are set to nil.
- When the  $i^{\text{th}}$  block is first written, a block is obtained from the free-space manager, and its address is put in the  $z^{\text{th}}$  index-block entry.
- Indexed allocation supports direct access, without suffering from external fragmentation, because any free block on the disk can satisfy a request for more space.
- Indexed allocation does suffer from wasted space, however.

Explanation  
: 2 ½

Fig : 2



b • **Full Virtualization**

- In full virtualization, the guest operating system is unaware that it is in a virtualized environment
- The hardware is virtualized by the host operating system
- The guest can issue commands to what it thinks is actual hardware, but really are just simulated hardware devices created by the host.

2

• **Para Virtualization**

- In para virtualization the guest operating system is aware that it is a guest and it has drivers for it
- Instead of issuing hardware commands, simply issues commands directly to the host operating system.

2

• **Partial Virtualization**

- The virtual machine simulates multiple instances of an underlying hardware environment.
- The entire operating systems cannot run in the virtual machine but many applications can run.
- This type of virtualization is far easier to execute than full virtualization.

2

6