

**C0-R4.B2 : OPERATING SYSTEM****NOTE :**

1. Answer question 1 and any FOUR questions from 2 to 7.
2. Parts of the same question should be answered together and in the same sequence.

**Total Time : 3 Hours****Total Marks : 100**

1. (a) What is the difference between hypervisor and virtualization ?  
 (b) What is cache ? What does L3 cache mean ?  
 (c) List four functions of Operating system as Process Management.  
 (d) A child process is created using fork(). Which states are shared by parent and child process : Heap, Stack, Shared memory segments ?  
 (e) How is Encryption different from Hashing ?  
 (f) What is the difference between cat command and more command in Unix ?  
 (g) List possible types of failures in Distributed Systems. (7x4)
  
2. (a) Differentiate between Monolithic systems and Microkernels. Give an example of each.  
 (b) A 255-GB disk has 65,536 cylinders with 255 sectors per track and 512 bytes per sector. How many platters and heads does this disk have ? Assuming an average cylinder seek time of 11 ms, average rotational delay of 7 msec and reading rate of 100 MB/sec, calculate the average time it will take to read 400 KB from one sector.  
 (c) Describe a method for process migration across different architectures running :  
 (i) The same operating system  
 (ii) Different operating systems  
 (d) Consider a system that has two CPUs, each CPU having two threads (hyperthreading). Suppose three programs, P0, P1, and P2, are started with run times of 5, 10 and 20 msec, respectively. How long will it take to complete the execution of these programs ?  
 Assume that all three programs are 100% CPU bound, do not block during execution, and do not change CPUs once assigned. (3+5+5+5)
  
3. (a) What information is stored in PCB ?  
 (b) Consider a system with two users, each of which has been promised 50% of the CPU. User 1 has four processes, A, B, C, and D, and user 2 has only one process, E. What will be the scheduling sequence that meets all the constraints if :  
 (i) round-robin scheduling is used  
 (ii) if user 1 is entitled to twice as much CPU time as user 2  
 (c) Can a thread ever be pre-empted by a clock interrupt? If so, under what circumstances ? If not, why not ?  
 (d) Suppose that a system is in an unsafe state. Show that it is possible for the processes to complete their execution without entering a deadlock state. (4+4+4+6)

4. (a) Why is rotational latency usually not considered in disk scheduling ? How would you modify SSTF, SCAN, and C-SCAN to include latency optimization ?  
 (b) Write about the techniques for structuring the page table. (9+9)
5. (a) A DMA controller has five channels. The controller is capable of requesting a 32-bit word every 40 nsec. A response takes equally long. How fast does the bus have to be to avoid being a bottleneck ?  
 (b) Could a RAID level 1 organization achieve better performance for read requests than a RAID level 0 organization (with nonredundant striping of data) ? If so, how ?  
 (c) How much cylinder skew is needed for a 7200-RPM disk with a track-to-track seek time of 1 msec ? The disk has 200 sectors of 512 bytes each on each track. Calculate the maximum data rate in bytes/sec for the disk. (5+5+8)
6. (a) Explain which memory management scheme overcomes the problem of fragmentation and how ?  
 (b) A process contains eight virtual pages on the disk and is assigned a fixed allocation of four pages frame in main memory the following page trace access.  
 1, 0, 2, 2, 1, 7, 6, 7, 0, 1, 2, 0, 3, 0, 4, 5, 1, 5, 2, 4, 5, 1, 5, 2, 4, 5, 6, 7, 2, 4, 2, 7, 3, 3, 2, 3  
 Show the successive page residing in the four-frame using the LRU replacement policy.  
 (c) Explain the solution to the Bounded Buffer and Dining philosophers Problem. (5+5+8)
7. (a) Why Access matrix used in Operating System ? Discuss with example.  
 (b) Explain the different types of shells available in Linux ?  
 (c) Consider the following set of processes with the length of CPU burst time given in milliseconds.

Process Name	Arrival Time	CPU Burst Time
A	0	3
B	1	5
C	3	2
D	9	6
E	12	5

Calculate the average waiting time and turnaround time (for each process) for the Round Robin (quantum=2ms) and SJF (pre-emptive) CPU scheduling.

(5+5+8)

- o O o -