

B34 - R4 : OPERATING SYSTEM**NOTE :**

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

Time : 3 Hours

Total Marks : 100

1. (a) Write two differences between user-level threads and kernel-level threads ? Under what circumstances is one type better than the other ?
 (b) Why does a system maintain Access Matrix ? Justify your answer by giving an example.
 (c) Define network operating system ? Write down its two main functionalities over a stand-alone system.
 (d) Why is page size always power of 2 ?
 (e) Briefly explain the paging technique used to provide virtual memory in computer system.
 (f) How does DMA increase system concurrency ? How does it complicate hardware design ?
 (g) Describe the actions taken by a kernel to context-switch between kernel level threads. (7x4)

2. (a) With respect to the RPC mechanism, consider the "exactly once" semantic. Does the algorithm for implementing this semantic execute correctly even if the ACK message sent back to the client is lost due to a network problem ? Describe the sequence of messages and discuss whether "exactly once" is still preserved.
 (b) Suppose that the following processes arrive for execution at the times indicated. Each process will run for the amount of time listed. In answering the questions, use nonpreemptive CPU scheduling, and base all decisions on the information you have at the time the decision must be made.

Process	Arrival Time	Burst Time
P1	0.0	8
P2	0.4	4
P3	1.0	1

 (i) What is the average turnaround time for these processes with the FCFS scheduling algorithm ?
 (ii) Find the average waiting time for these processes with the SJF scheduling algorithm ? (9+9)

3. (a) Why are spinlocks not appropriate for single-processor systems but are often used in multiprocessor systems. Explain briefly.
 (b) Consider the following snapshot of a system:

	Allocation	Max	Available
	A B C D	A B C D	A B C D
P0	0 0 1 2	0 0 1 2	1 5 2 0
P1	1 0 0 0	1 7 5 0	
P2	1 3 5 4	2 3 5 6	
P3	0 6 3 2	0 6 5 2	
P4	0 0 1 4	0 6 5 6	

 Answer the following using the banker's algorithm:
 (i) What is the content of the matrix Need ?
 (ii) Is the system in a safe state ?
 (iii) If a request from process P1 arrives for resources (0,4,2,0), can the request be granted immediately ? (6+12)

4. (a) What is the meaning of the term busy waiting? What other kinds of waiting are there in an operating system? Can busy waiting be avoided altogether? Explain your answer.
- (b) Assume that an operating system maps user-level threads to the kernel using the many-to-many model and the mapping is done through LWPs. Furthermore, the system allows developers to create real-time threads for use in real-time systems. Is it necessary to bind a real-time thread to a LWP? Explain. (9+9)
5. (a) Can a system detect that some of its processes are starving? Justify your answer.
- (b) An operating system supports a paged virtual memory, using a central processor with a cycle time of 1 microsecond. It costs an additional 1 microsecond to access a page other than the current one. Pages have 1000 words, and the paging device is a drum that rotates at 3000 revolution per minute and transfers 1 million words per second. The following statistical measurements were obtained from the system:
- 1 percent of the executed instructions access a page other than the current page.
 - Of the instructions that accessed another page, 80 percent accessed a page already in memory.
 - When a new page was required, the replaced page was modified 50 percent of the time.
- Calculate the effective instruction time on this system, assuming that the system is running one process only and that the processor is idle during drum transfers. (9+9)
6. (a) Consider the following page reference string:
1, 2, 3, 4, 2, 1, 5, 6, 2, 1, 2, 3, 7, 6, 3, 2, 1, 2, 3, 6.
How many page faults would occur for the following page replacement algorithms, assuming two, four, seven frames? Assume all frames are initially empty.
- LRU replacement
 - FIFO replacement
 - Optimal replacement
- (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? (12+6)
7. (a) Consider a system that supports the contiguous, linked, and indexed file allocation methods. What criteria should be used in deciding which method is best utilized for a particular file?
- (b) A distributed system has two sites, A and B. How does site A can identify the following?
- (i) B goes down.
 - (ii) The link between A and B goes down.
 - (iii) B is extremely overloaded, and its response time is 100 times longer than normal.
- How does site A recover from such situation in the distributed system?
- (c) Distinguish between the client-server and peer-to-peer models of distributed systems. (6+6+6)

- o o o -