

- (b) How is each page indexed in a page table? What is the reason for associating a valid/invalid bit for each entry in a page table?

$$6 + (3 + 3) = 12$$

7. (a) What is demand paging? What are the benefits of using it? What is pure demand paging?
- (b) Consider the following page reference string: 5, 8, 4, 0, 2, 8, 4, 9, 6, 5, 6, 4, 5, 8. How many page faults would occur for the following replacement algorithms, assuming memory consisting of 5 frames?
- Not Recently Used (NRU) page replacement strategy.
 - Optimal page replacement strategy.
- (c) Consider a paging system using TLB where it takes 30 nanoseconds to access the TLB and 150 nanoseconds to access memory. Given the hit ratio of 70%, calculate how long it takes to access a desired byte in memory.

$$(2 + 2 + 1) + 4 + 3 = 12$$

Group - E

8. (a) Compare and contrast linked allocation and contiguous allocation for file systems.
- (b) Suppose that a disk drive has 300 cylinders, numbered 0 to 299. The drive is currently serving a request at cylinder 90, and the previous request was at cylinder 50. The queue of pending requests, in FIFO order, is: 36, 79, 15, 120, 199, 270, 89 and 170. Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests for each of the following disk-scheduling algorithms?

(i) FCFS (ii) SSTF (iii) LOOK (iv) C-LOOK

$$4 + 8 = 12$$

9. (a) What is the difference between polling and interrupt? What is an interrupt vector?
- (b) Explain the difference between memory-mapped I/O and I/O through special instructions to device registers.
- (c) Explain the term Direct Memory Access (DMA). What is the use of a DMA controller?
- (d) How can an access control matrix be expanded to include domains as objects?

$$4 + 2 + 4 + 2 = 12$$

OPERATING SYSTEMS (MCAP 2104)

Time Allotted : 3 hrs

Full Marks : 70

Figures out of the right margin indicate full marks.

Candidates are required to answer Group A and any 5 (five) from Group B to E, taking at least one from each group.

Candidates are required to give answer in their own words as far as practicable.

Group - A (Multiple Choice Type Questions)

1. Choose the correct alternative for the following: **10 × 1 = 10**
- Which one of the following is not shared by threads?
 (a) Program counter (b) Stack.
 (c) Both program counter and stack (d) None of these.
 - In paged memory systems, if the page size is increased, then the internal fragmentation generally
 (a) becomes less (b) becomes more
 (c) remains constant (d) vanishes.
 - Inter process communications can be done through
 (a) Mails (b) Messages
 (c) System calls (d) Traps.
 - If all page frames are initially empty, and a process is allocated 3 page frames in real memory and references its pages in the order 1 2 3 2 4 5 2 3 2 4 1 and the page replacement is FIFO, the total number of page faults caused by the process will be
 (a) 7 (b) 8 (c) 9 (d) 10.
 - Which of the following disk scheduling techniques could lead to starvation of a process?
 (a) SCAN (b) C-LOOK (c) SSTF (d) FCFS.
 - DMA transfers are initiated by:
 (a) The process being executed (b) Processor
 (c) I/O device (d) DMA controller.

- (vii) In which of the following page replacement policies, Belady's anomaly occurs?
 (a) FIFO (b) LRU (c) LFU (d) NRU.
- (viii) A counting semaphore was initialized to 10. Then 6P (wait) operations and 4V (signal) operations were completed on this semaphore. The resulting value of the semaphore
 (a) 0 (b) 8 (c) 10 (d) 12.
- (ix) The amount of time that a job spends waiting in the ready queue is called ____ time.
 (a) Waiting time (b) Response time
 (c) Turnaround time (d) Throughput.
- (x) Thrashing
 (a) Reduces page I/O
 (b) Decreases the degree of multiprogramming
 (c) Implies excessive page I/O
 (d) Improves the system performance.

Group - B

2. (a) A CPU-scheduling algorithm determines an order for the execution of its scheduled processes. Given n processes to be scheduled on one processor, how many different schedules are possible? (Give a formula in terms of n .)
- (b) Discuss about Multilevel feedback Queue Scheduling.
- (c) Consider the following processes with the length of CPU burst time given in milliseconds.

Process	Burst Time(ms)	Arrival Time(ms)	Priority
P1	10	0	3
P2	1	1	1
P3	2	3	3
P4	1	6	4
P5	5	8	2

Lower number define higher priority.

Draw Gantt charts, illustrating the execution of these processes using:

- Preemptive SJF (SRTN) (Ignore priority column).
- Preemptive Priority Scheduling.

Also find out average waiting time in all the cases.

$$2 + 2 + 8 = 12$$

3. (a) What is process control block (PCB)? Briefly explain the different components of PCB.

- (b) What is the main advantage of the layered approach to system design? What are the disadvantages of using the layered approach?
- (c) What are two differences between user-level threads and kernel-level threads?
- (d) Describe the differences among short-term, medium-term, and long-term scheduling.

$$(2 + 3) + 2 + 2 + 3 = 12$$

Group - C

4. (a) Define Mutual-exclusion implementation with test_and_set() instruction. Explain with example.
- (b) Explain the first known correct software solution to the critical-section problem for two processes, developed by Dekker.
- (c) What is the similarity between the Dekker's Algorithm and Peterson's Algorithm?
- (d) Write the difference between Semaphore and Monitor.

$$4 + 5 + 1 + 2 = 12$$

5. (a) What is semaphore? What are the different types of semaphore?
- (b) Describe the classical Sleeping Barber Problem. Write down a solution of it using semaphore.
- (c) Explain the message passing techniques used in inter process communication.

$$2 + 7 + 3 = 12$$

Group - D

6. (a) Consider the following snapshot of a system:

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

Answer the following questions using the banker's algorithm:

- What is the content of the matrix **Need**?
- Is the system in a safe state?
- If a request from process P1 arrives for (0, 2, 1, 0), then can the request be granted immediately?

Briefly describe the methods used by operating system to *recover* from deadlock.