Paper id: 252378

Roll No: | | | | | | | | | | | | | |

MCA
(SEM II) THEORY EXAMINATION 2024-25
OPERATING SYSTEMS

TIME: 3 HRS **M.MARKS: 100**

Note: Attempt all Sections. In case of any missing data; choose suitably.

SECTION A

1. Attempt *all* questions in brief. **2 x 10 = 20**

Q No.	Question	CO	Level
a.	Define the term "Operating System" and explain its role in a computer system.	1	K_2
b.	What is the role of cache memory management in an operating system?	1	K_2
c.	What is CPU scheduling? Why is it necessary in a multiprogramming environment?	3	K_2
d.	Define semaphores. How are they used to solve synchronization issues in operating systems?	2	K_3
e.	Describe the typical elements of the process control block	3	K_2
f.	What is the use of inter process communication and context switching?	2	K_3
g.	Explain the logical address space and physical address space diagrammatically.	4	K_5
h.	What do you mean by the safe state and an unsafe state?	4	K_3
i.	Define Seek time and Latency time.	5	K_5
j.	What do you mean by the I/O Buffering?	5	K_5

SECTION B

2. Attempt any *three* of the following: **10 x 3 = 30**

a.	In what scenario would a microkernel be preferred over a monolithic kernel, considering system reliability and security?	1	K_2
b.	How do semaphores handle synchronization differently compared to busy-waiting techniques like Test-and-Set? Which is more CPU efficient?	2	K_3
c.	Why is context switching essential during CPU scheduling, and what information must be stored in the Process Control Block (PCB) during this operation?	3	K_2
d.	Explain how paging solves the problem of external fragmentation but still suffers from internal fragmentation.	4	K_4
e.	A hard disk having 2000 cylinders, numbered from 0 to 1999. the drive is currently serving the request at cylinder 143,and the previous request was at cylinder 125.The status of the queue is as follows 86, 1470, 913, 1774,948,1509,1022,1750,130 What is the total distance (in cylinders) that the disk arm moves to satisfy the entire pending request for each of the following disk-scheduling algorithms? i) SSTF ii) FCFS	5	K_5

SECTION C

3. Attempt any *one* part of the following: **10 x 1 = 10**

a.	Why are reentrant kernels crucial for systems with multiple concurrent processes? What problem do they solve?	1	K_2
b.	How would you differentiate between multithreading and multiprocessing in terms of resource usage and inter-process communication?	1	K_2

Paper id: 252378

Roll No: | | | | | | | | | | | | | |

MCA
(SEM II) THEORY EXAMINATION 2024-25
OPERATING SYSTEMS

TIME: 3 HRS **M.MARKS: 100**

4. Attempt any *one* part of the following: **10 x 1 = 10**

a.	Compare Peterson's and Dekker's solutions to the Critical Section Problem. Under what condition might one be preferred over the other?	**2**	**K_3**
b.	Explain how the Test-and-Set operation helps in solving the Critical Section Problem. What is its major drawback in modern systems?	**2**	**K_3**

5.Attempt any *one* part of the following: **10 x 1 = 10**

a. (3, K_2)

Consider the following process:

Process	Arrival Time	Burst Time
P1	0	8
P2	1	4
P3	2	9
P4	3	5

What is the average waiting and turnaround time for these processes with
i) FCFS Scheduling
ii) Preemptive SJF Scheduling

b. (3, K_2)

Consider the following snapshot of a system:

	Allocated			Maximum			Available		
Process	R1	R2	R3	R1	R2	R3	R1	R2	R3
P1	2	2	3	3	6	8	7	7	10
P2	2	0	3	4	3	3			
P3	1	2	4	3	4	4			

Answer the following questions using the banker's algorithm:
i) What is the content of the matrix need?
ii) Is the system in a safe state

6.Attempt any *one* part of the following: **10 x 1 = 10**

a.	Why is demand paging considered more efficient in terms of memory usage, and what is the risk if the page fault rate becomes too high?	4	K_3
b.	Consider a logical address space of 64 pages of 1,024 words each, mapped onto a physical memory of 32 frames. i) How many bits are there in the logical address? ii) How many bits are there in the physical address?	4	K_3

7.Attempt any *one* part of the following: **10 x 1 = 10**

a.	How does a RAID system improve data reliability and performance? What trade-offs are involved?	5	K_5
b.	How do hierarchical file directories help organize data more effectively than flat file systems? What are the benefits in terms of access and management?	5	K_5