No. of Printed Pages : 4

Sl. No.

B1.4-R5 : OPERATING SYSTEMS

DURATION : 03 Hours **MAXIMUM MARKS : 100**

Roll No. : | | | | | | |

Answer Sheet No. : | | | | | | |

Name of Candidate : _____ **; Signature of Candidate :** _____

INSTRUCTIONS FOR CANDIDATES :

- Carefully read the instructions given on Question Paper, Answer Sheet.
- Question Paper is in English language. Candidate has to answer in English Language only.
- **Question paper contains Seven questions. The Question No. 1 is compulsory. Attempt any FOUR Questions from Question No. 2 to 7.**
- **Parts of the same question should be answered together and in the same sequence.**
- **Questions are** to be answered in the **ANSWER SHEET** only, supplied with the Question Paper.
- **Candidate cannot leave the examination hall/ room without signing on the attendance sheet and handing over his/her Answer Sheet to the Invigilator. Failing in doing so, will amount to disqualification of Candidate in this Module/Paper.**
- After receiving the instruction to open the booklet and before answering the questions, the candidate should ensure that the Question Booklet is complete in all respects.

DO NOT OPEN THE QUESTION BOOKLET UNTIL YOU ARE TOLD TO DO SO.

1. (a) Distinguish between 'user mode' and 'privileged mode' of a CPU.

(b) What is an embedded operating system and how does it differ from a traditional operating system ? In what situations is a hard-real-time operating system typically used in an embedded system ?

(c) What is the difference between a stateful and a stateless distributed network operating system ?

(d) In a system using paging with a page size of 4 KB, how many bits are required for the page offset if the size of the virtual address space is 32-bit ?

(e) What is swap disk management and why is it important for computer performance ?

(f) What is a kernel in an operating system, and what are its primary functions ? How does the kernel interact with other system components, such as device drivers and user applications ?

(g) What is the maximum size of a message that can be sent between two processes using message passing ? **(7x4)**

2. (a) How can virtualization technology be improved to provide better support for real-time applications running on virtual machines ? Real-time applications often require low latency and high throughput, but virtualization can introduce performance overheads that can impact their responsiveness. What techniques can be used to minimize this overhead and provide real-time guarantees for such applications ?

(b) Describe the key features of Windows operating system, and how they differ from open source system like Linux.

(c) What are the key differences between hard, soft, and firm real-time operating systems, and how do these differences affect their use in embedded systems ? **(8+5+5)**

3. (a) Consider a disk with the following queue of disk requests, given the current head position is 50 :

Request Queue : 98, 183, 37, 122, 14, 124, 65, 67

Using the SCAN (Elevator) disk scheduling algorithm, what is the total head movement if the disk arm is moving towards the larger numbers first ?

(b) Assume that a computer has an NVM (Non-Volatile Memory) storage device with a total capacity of 1 TB. The device uses a file system that supports block sizes of 4 KB. If the device is currently 80% full and you want to save a file of size 10 MB, how many blocks will be used to store the file ?

(c) What are the different structures of an operating system ? List main structures of an operating system and explain briefly. **(7+7+4)**

4. (a) Consider a system with three processes, P1, P2, and P3, each of which is executed by a separate thread. These processes share a common resource, *R*, which can only be used by one process at a time. Each process requires access to *R* to complete its task. Initially, *R* is free. The processes are executed in the following sequence :

P1, P2, P3, P1, P2, P3, …

Assume that each process takes 2 units of time to complete its task once it has access to *R* and that context switching between processes takes 1 unit of time. If the processes are implemented using a semaphore-based solution with an initial value of 1 and acquiring and releasing the semaphore each takes 1 unit of time, what is the total time taken to complete the first 10 rounds of the sequence ?

(b) What is starvation in the context of process scheduling in operating systems ? How can starvation be prevented ?

(c) What is the difference between simple and shadow volumes in file systems ? **(7+7+4)**

5. (a) Describe the two main types of kernel architectures : monolithic kernels and microkernels. What are the advantages and disadvantages of each architecture, and in what scenarios might one be more suitable than the other ?

(b) What is the critical section problem and how does Peterson's solution address it ?

(c) Consider a system with 3 processes and 3 resources of a particular type. Each process needs one resource to complete its execution. The following table represents the current allocation and request matrix for the system.

Process	Allocation	Request
P1	1	1
P2	1	1
P3	1	1

Is the system currently in a deadlock state ? If so, what would be the safe sequence to recover from the deadlock ? **(5+5+8)**

6. (a) How does the Optimal page replacement algorithm compare to the LRU and FIFO algorithms in terms of page fault rates and memory usage in different types of workloads ?

(b) What is thrashing in operating systems, and how can it be prevented ?

(c) What are the four necessary conditions of deadlock prevention ? **(6+6+6)**

7. (a) What is the domain and goal of protection, and how does it relate to protection rings ?

(b) What is Security Enhanced Linux (SE Linux) and how does it use Mandatory Access Control (MAC) to improve system security ?

(c) What are the goals and objectives of a distributed file system, and how do they relate to transparency and accuracy ? **(6+6+6)**

- o O o -

SPACE FOR ROUGH WORK