

B3.4-R4: OPERATING SYSTEMS

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) What is sector interleaving and why it is needed?
 - b) Explain how a process is represented in system memory.
 - c) What are the main differences between capability lists and access lists?
 - d) What are interrupts? How are interrupts handled by the OS?
 - e) What is the importance of memory management in operating systems?
 - f) In Linux, shared libraries perform many operations central to the operating system. What is the advantage of keeping this functionality out of the kernel?
 - g) Differentiate between Network File System and Distributed File System.

(7×4)

2.
 - a) Given memory partitions of 100K, 500K, 200K, 300K, and 600K (in order), how would each of the First-fit, Best-fit, and Worst-fit algorithms place processes of 212K, 417K, 112K, and 426K (in order)? Which algorithm makes the most efficient use of memory?
 - b) A system has four processes P1 through P4 and two resource types R1 and R2. It has 2 units of R1 and 3 units of R2. Given that:
P1 requests 2 units of R2 and 1 unit of R1, P2 holds 2 units of R1 and 1 unit of R2, P3 holds 1 unit of R2, P4 requests 1 unit of R1
Show the resource allocation graph for this state of the system. Is the system in deadlock, and if so, which processes are involved?
 - c) Does DFS Replication support file screens created by File Server Resource Manager? Explain.

(6+6+6)

3.
 - a) Explain working flow of Peterson's solution when interrupt is occurred at the position Other = 1 – Process.
 - b) Consider following reference string: 7, 0, 1, 2, 0, 3, 0, 4, 2, 3, 0, 3, 2, 1, 2, 0, 1, 7, 0, 1. Find how many page faults occur using
 - i) FIFO
 - ii) LRU Page replacement algorithms for three page frames
 - c) What are computer worms and viruses? What are the safeguards used against worms and viruses?

(7+6+5)

4.
 - a) The Linux kernel does not allow paging out of kernel memory. What effect does this restriction have on the kernel's design? Give two advantages and two disadvantages of this design decision?
 - b) Draw and explain a process state transition diagram with one suspended state.
 - c) Why do some systems keep track of the type of a file, while others leave it to the user or simply do not implement multiple file types? Which system is "better"?
 - d) What are two main differences between a WAN and a LAN?

(6+6+4+2)

- 5.
- Describe three circumstances under which blocking I/O should be used. Describe three circumstances under which nonblocking I/O should be used. Why not just implement nonblocking I/O and have processes busy-wait until their device is ready?
 - Differentiate between monolithic kernel and microkernel with proper diagram.
 - Explain deadlock detection methods in distributed systems.

(6+6+6)

- 6.
- Consider the following set of process in order P1, P2, P3, P4 and P5 with the length of the CPU burst time given in milliseconds. Their priorities are 3,1,4,5 and 2 respectively, with 1 being the highest priority.

Process	Arrival Time	Burst Time
P1	0	10
P2	0	1
P3	0	2
P4	0	1
P5	3	5

Calculate average waiting time and turn round time using following scheduling algorithms.

- Non-preemptive priority
 - Preemptive priority
 - SJF
- What is demand paging? How it is different from demand segmentation?
 - Explain about protection domain in UNIX.

(9+5+4)

- 7.
- Write short notes on the following:
 - Access Matrix
 - Semaphore
 - Dispatcher
 - Multithreading
 - Discuss about Windows (NT) file system in detail.
 - Is the Process, before and after the swap, the same? Give reason.
 - Under what circumstances is a token-ring network more effective than an Ethernet network?

(8+5+3+2)