

C0-R4.B2: 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) What are the three main purposes of an operating system?
 - b) Why do some systems store the operating system in firmware, while others store it on disk?
 - c) Describe the actions taken by a kernel to context-switch between kernel level threads.
 - d) Name two differences between logical and physical addresses.
 - e) Why is rotational latency usually not considered in disk scheduling?
 - f) Systems that support sequential files always have an operation to rewind files. Do systems that support random access files need this too?
 - g) Explain Process Control Block.

(7x4)

2.
 - a) What resources are used when a thread is created? How do they differ from those used when a process is created?
 - b) 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?

(9+9)

3.
 - a) 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 non-preemptive 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) What is the average turnaround time for these processes with the SJF scheduling algorithm?
 - iii) The SJF algorithm is supposed to improve performance, but notice that we chose to run process P1 at time 0 because we did not know that two shorter processes would arrive soon. Compute what the average turnaround time will be if the CPU is left idle for the first 1 unit and then SJF scheduling is used.
 - b) Show that, if the wait() and signal() semaphore operations are not executed atomically, then mutual exclusion may be violated.

(12+6)

4.
 - a) Consider a paging system with the page table stored in memory. If a memory reference takes 200 nanoseconds, how long does a paged memory reference take? If we add a TLB and 75 percent of all page-table references are TLB hits, what is the effective memory reference time?
 - b) Could a RAID level 1 organization achieve better performance for read requests than a RAID level 0 organization? If so, how?

(9+9)

5.

- a) Consider a system that supports 5,000 users. Suppose that you want to allow 4,990 of these users to be able to access one file.
- i) How would you specify this protection scheme in UNIX?
 - ii) Can you suggest another protection scheme that can be used more effectively for this purpose than the scheme provided by UNIX?
- b) Explain architecture of UNIX Environment.

(9+9)

6.

- a) Consider the following segment table:

Segment	Base	Length
0	219	600
1	2300	14
2	90	100
3	1327	580
4	1952	96

What are the physical addresses for the following logical addresses?

- i) 0, 430
 - ii) 1, 10
 - iii) 2, 500
 - iv) 3, 400
- b) Suppose that a disk has 1,500 cylinders numbered 0 to 1499. The disk head is currently at cylinder 143 and previous request was at 120. The queue of pending requests in FIFO order is:
470, 913, 777, 1248, 1409, 130, 500

Calculate and show total head movements for the following disk scheduling algorithms:

- i) FCFS
- ii) SSTF
- iii) SCAN

(9+9)

7.

- a) Consider a system that supports the strategies of contiguous, linked, and indexed allocation. What criteria should be used in deciding which strategy is best utilized for a particular file?
- b) Why is deadlock detection much more expensive in a distributed environment than in a centralized environment?

(9+9)