Reg. No. : | | | | | | | | | | | | |

Question Paper Code : 10266

B.E/B.Tech. DEGREE EXAMINATION, MAY/JUNE 2012.

Fourth Semester

Computer Science and Engineering

CS2254/141404/CS45/CS1253/10144CS405/080250012 — OPERATING SYSTEMS

(Common to Information Technology)

(Regulation 2008)

Time : Three hours Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. What are the main purposes of an operating system?

2. What are the differences between user-level threads and kernel-level threads?

3. What is the difference between preemptive and nonpreemptive scheduling?

4. What are the four necessary conditions that are needed for deadlock can occur?

5. Consider a logical address space of eight pages of 1024 words each, mapped onto a physical memory of 32 frames. How many bits are there in the logical address and in the physical address?

6. What is the advantage of demand paging?

7. What are the two types of system directories?

8. What is garbage collection?

9. What is seek time?

10. What characteristics determine the disk access speed?

PART B — (5 × 16 = 80 marks)

11. (a) (i) Define the essential properties of the following types of operating systems:

(1) Batch

(2) Time sharing

(3) Real time

(4) Distributed (8)

(ii) List five services provided by an operating system. Explain how each provides convenience to the users. Explain also in which cases it would be impossible for user – level programs to provide these services. (8)

Or

(b) (i) What two advantages do threads have over multiple processes? What major disadvantages do they have? Suggest one application that would benefit from the use of threads. (8)

(ii) Explain the various issues associated with the thread in detail. (8)

12. (a) (i) What is a Gantt chart? Explain how it is used. (4)

(ii) Consider the following set of processes, with the length of the CPU – burst time given in milliseconds:

Process	Burst Time	Priority
P1	10	3
P2	1	1
P3	2	3
P4	1	4
P5	5	2

The processes are arrived in the order P1, P2, P3, P4, P5, all at time 0.

(1) Draw four Gantt charts illustrating the execution of these processes using FCFS, SJF, a nonpreemptive priority (a smaller priority number implies a higher priority), and RR (quantum = 1) scheduling.

(2) What is the turnaround time of each process for each of the scheduling algorithms in part a?

(3) What is the waiting time of each process for each of the scheduling algorithms in Part a?

(4) Which of the schedules in part a results in the minimal average waiting time (over All processes)? (12)

Or

(b) (i) What do you mean by busy waiting? What other kinds of waiting are there? Can busy waiting be avoided altogether? Explain your answer. (8)

(ii) Consider the following snapshot of a system:

	Allocation	Max	Available
	ABCD	ABCD	ABCD
P0	0012	0012	1520
P1	1000	1750	
P2	1354	2356	
P3	0632	0652	
P4	0014	0656	

Answer the following questions based on the banker's algorithm:

(1) Define safety algorithm.

(2) What is the content of the matrix Need?

(3) Is the system in a safe state?

(4) If a request from process P1 arrives for (0,4,2,0), can the request be granted immediately? (8)

13. (a) (i) Why are segmentation and paging sometimes combined into one scheme? (4)

(ii) Consider the following page reference string:

1, 2, 3, 4, 2, 1, 5, 6, 2, 1, 2, 3, 7, 6, 3, 2, 1, 2, 3, 6.

How many page faults would occur for the following replacement algorithms, assuming one, two, three, four, five, six, and seven frames? Remember all frames are initially empty, so your first unique pages will all cost one fault each.

(1) LRU replacement

(2) FIFO replacement

(3) Optimal replacement (12)

Or

(b) (i) 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?

(1) 0, 430 (2) 1, 10

(3) 2, 500 (4) 3, 400 (8)

(ii) Discuss briefly about memory management in LINUX. (8)

14. (a) (i) Explain the various attributes of a file. (4)

(ii) Consider a file currently consisting of 100 blocks. Assume that the file control block (and the index block, in the case of indexed allocation) is already in memory. Calculate how many disk I/O operations are required for contiguous, linked, and indexed (single-level) allocation strategies, if, for one block, the following conditions hold. In the contiguous allocation case, assume that there is no room to grow in the beginning, but there is room to grow in the end. Assume that the block information to be added is stored in memory.

(1) The block is added at the beginning.

(2) The block is added in the middle.

(3) The block is added at the end.

(4) The block is removed from the beginning.

(5) The block is removed from the middle.

(6) The block is removed from the end. (12)

Or

(b) (i) Explain the various schemes used for defining the logical structure of a directory. (8)

(ii) Describe the approaches used in free space management. (8)

15. (a) (i) Consider the following I/O scenarios on a single-user PC.

(1) A mouse used with a graphical user interface

(2) A tape drive on a multitasking operating system (assume no device preallocation is available)

(3) A disk drive containing user files

(4) A graphics card with direct bus connection, accessible through memory-mapped I/O

For each of these I/O scenarios, would you design the operating system to use buffering, spooling, caching, or a combination? Would you use polled I/O, or interruptdriven I/O?

Give reasons for your choices. (8)

(ii) How do you choose a optimal technique among the various disk scheduling techniques? Explain. (8)

Or

(b) (i) Describe the various disk scheduling techniques. (8)

(ii) Describe the various levels of RAID. (8)
