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

C 3225

B.E./B.Tech. DEGREE EXAMINATION, NOVEMBER/DECEMBER 2008.

Seventh Semester

(Regulation 2004)

Electronics and Communication Engineering

EC 1007 — OPERATING SYSTEMS

(Common to B.E. (Part-Time) Sixth Semester Regulation 2005)

Time : Three hours Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. Under what circumstances would a user be better off using a time-sharing system, rather than a PC or single-user workstation?

2. How does the distinction between monitor mode and user mode function as a rudimentary form of protection (security) system?

3. What resources are used when a thread is created? How do they differ from those used when a process is created?

4. A CPU scheduling algorithm determines an order for the execution of its scheduled processes. Given n processes to be scheduled on one processor, how many different schedules are possible?

5. Name two differences between logical address and physical address.

6. Under what circumstances do page faults occur?

7. Why might a system use interrupt-driven I/O to manage a single serial port, but polling I/O to manage a front-end processor, such as a terminal concentrator?

8. Why is rotational latency usually not considered in disk scheduling?

9. List two benefits of DFS when compared to a file system in a centralized system?

10. How is atomicity implemented in DFS?

PART B — ($5 \times 16 = 80$ marks)

11. (a) (i) In a multi-programming and time-sharing environment, several users share the system simultaneously. This situation can result in various security problems.

(1) Mention two such problems.

(2) Can we ensure the same degree of security in a time-shared machine as we have in a dedicated machine? Explain your answer.

(ii) Describe the differences between symmetric and asymmetric multiprocessing. What are the advantages and disadvantages of multiprocessor systems?

Or

(b) List the services provided by an operating system. Explain how each one provides convenience to the users. Explain in which cases it would be impossible for user-level programs to provide these services.

12. (a) (i) Consider a system consisting of four resources of the same type that are shared by three processes, each of which needs at most two resources. Show that the system is deadlock-free.

(ii) Describe the actions taken by a kernel to switch context between processes.

Or

(b) Consider a variant of RR scheduling algorithm where the entries in the ready queue are pointers to PCBs.

(i) What would be the effect of putting two pointers to the same process in the ready queue?

(ii) What would be the major advantages and disadvantages of this scheme?

(iii) How would you modify the basic RR algorithm to achieve the same effect without the duplicate pointers?

13. (a) Sharing segments among processes without requiring the same segment number is possible in a dynamically linked segmentation system.

(i) Define a system that allows static linking and sharing of segments without requiring that the segment numbers be the same

(ii) Describe a paging scheme that allows pages to be shared without requiring that the page numbers be the same.

Or

(b) Discuss the segment-replacement algorithms based on FIFO and LRU policies both for static relocation system (where segments cannot be relocated) and for dynamic relocation system (where segments can be relocated).

14. (a) (i) Explain why SSTF scheduling tends to favor middle cylinders over the innermost and outermost cylinders.

(ii) What are the tradeoffs involved in rereading code pages from the file system, versus using swap space to store them?

Or

(b) Discuss the relative advantages and disadvantages of sector sparing and sector slipping.

15. (a) Explain the election algorithm with neat illustrations.

Or

(b) Under what circumstances would a client prefer a location-transparent DFS? Under which would she prefer a location-independent DFS? Discuss the reasons for these preferences.
