

M.E. DEGREE EXAMINATIONS: JANUARY 2009

First Semester

COMPUTER SCIENCE AND ENGINEERING

P07CS103 Operating Systems

Time: Three Hours

Maximum Marks: 100

Answer ALL Questions:-

PART A (20 x 1 = 20 Marks)

1. Which of the following states is not a discrete process state?
A) ready state B) blocked state C) unblocked state D) running state
2. Major problem with priority scheduling is
A) Aging B) Convoy effect C) more context switching D) Starvation
3. Which technique was introduced because a single job could not keep both the CPU and the I/O devices busy?
A) Time-sharing. B) Spooling C) Preemptive scheduling D) Multiprogramming.
4. Which of the following is not a role of a typical operating system?
A) Control the allocation of the machine's resources
B) Control access to the machine
C) Maintain records regarding files stored in mass storage
D) Assist the computer user in the task of processing digital photographs
5. Multithreading is important because it can _____.
A) facilitate software design and promote good programming practices
B) improve performance and scalability
C) facilitate cooperation/synchronization of activities
D) all of the above
6. How can a semaphore s be used to solve the critical-section problem?
A) Initialize s to 0; enter section after waiting on s ; exit after signaling on s .
B) Initialize s to 1; enter section after waiting on s ; exit after signaling on s .
C) Initialize s to 0; enter section after signaling on s ; exit after waiting on s .
D) Initialize s to 1; enter section after signaling on s ; exit after waiting on s .
7. Process termination is done in
A) Dead lock prevention B) Deadlock avoidance
C) Dead lock detection D) Dead lock recovery
8. What are the issues to be considered for resource preemption?
i) Selecting a victim process to preempt.
ii) Rollback the process to a safe state.
iii) Ensure that starvation does not occur for any process.
A) ii only B) ii and iii only C) i, ii and iii D) i and ii only

9. When an OS spends much of its time paging, instead of executing application software, it is said to be _____.
- A) Thrashing B) formatting C) spooling D) booting
10. Noncontiguous memory allocation splits programs into blocks of memory called _____ that can be loaded into nonadjacent "holes" in main memory.
- A) pages B) partitions C) frames D) segments
11. If the virtual memory is 129 KB, main memory is 32 KB, page size is 8 KB, then the size of the page table is
- A) 16 B) 8 C) 4 D) 2
12. A FIFO page-replacement strategy can be implemented with relatively low overhead using a _____.
- A) queue B) hash table C) binary tree D) stack
13. Which of the following components of an operating system maintains the directory system?
- A) Device drivers B) File manager C) Memory manager D) Processor
14. Which of the following is an activity of an operating system in regard to file management?
- A) The support of primitives for manipulating files and directories
 B) Allocate and deallocate memory space as needed
 C) Creation and deletion of user and system processes
 D) Allocation of storage space when new files have to be written.
15. The amount of time required for the head to move from its current cylinder to its destination.
- A) Rotational latency time B) Transmission time
 C) Seek time D) Process time
16. The shared subdirectories and files are example of
- A) acyclic graph B) tree structure directory
 C) one level directory D) two level directory
17. Location transparency is name mapping of distributed file system is
- A) name of file does not reveal files physical storage location
 B) name of file does reveal files physical storage location
 C) name of file need not be changed when files physical storage location changes
 D) name of file need to be changed when files physical storage location changes
18. Which environment considers memory, process, and device and file management from a global viewpoint?
- A) Distributed Operating System B) Network Operating System
 C) Multiprogramming Operating System D) Realtime OS

19. Regarding to a thread,
 A) multiple threads in a process may share the same address space;
 B) a complete independent address space needs to be allocated;
 C) when it (thread) is blocked, all the threads in the same process must also be blocked;
 D) it is a kernel activity
20. By default, Linux supports _____
 A) symmetric multiprocessing (SMP)
 B) cache-only memory access (COMA)
 C) programs written for operating systems including Windows and MacOS
 D) all of the above

PART B (5 x 16 = 80 Marks)

- 21a) i) What are the desirable features of an operating system? Compare these features with Distributed systems and Real time systems. (8)
- ii) Why must the operating system be more careful when accessing input to a system call (or producing the result) when the data is in memory instead of registers? Explain the Operating System services. (8)

(OR)

- 21b) Consider the following sets of processes, with the length of the CPU- burst time given in milliseconds. Arrival time is the time at which the process is added to the ready queue.

| Process | Burst time | Arrival time |
|---------|------------|--------------|
| P1 | 9 | 0 |
| P2 | 6 | 0 |
| P3 | 3 | 0 |
| P4 | 4 | 0 |
| P5 | 2 | 4 |
| P6 | 1 | 6 |

Show in a diagram form (Gantt chart) how the methods you describe would schedule the processes.

What is the turnaround time for each job, the waiting time for each job, and the average waiting time for all jobs for each of the following scheduling algorithms? State the algorithms which are preemptive and non-preemptive and why it is considered so?

- (i) First Come First Served
 (ii) Shortest Job First
 (iii) Shortest Remaining Time First
 (iv) Round Robin with a quantum size = 2

(16)

- 22.a) i) A system has three processes (P1,P2, P3) and three resources (R1,R2, R3). There are two instances of R1 and R3. There is one instance of R2. P1 holds an R1 and an R3 and is requesting an R2. P2 holds an R1 and an R2 and is requesting an R3. P3 hold an R3 and is requesting an R2. Draw the resource allocation graph for this situation. Does a deadlock exist? Explain.

(10)

- ii) What is a test-and-set instruction? How can it be used to implement mutual exclusion? Consider using a fragment of psuedo-assembly language. (6)

(OR)

- 22b) i) A system has 3 "A" resources, 10 "B", and 5 "C." Five processes, their current allocations and their maximum allocations are shown below. Is the system in a safe state? If so, show one sequence of processes which allows the system to complete. If not, explain why.

| | Allocation | | | Max | | | Available | | |
|----|------------|---|---|-----|---|---|-----------|---|---|
| | A | B | C | A | B | C | A | B | C |
| P0 | 1 | 3 | 0 | 1 | 5 | 3 | 1 | 4 | 2 |
| P1 | 0 | 2 | 0 | 2 | 2 | 5 | | | |
| P2 | 1 | 0 | 1 | 2 | 4 | 2 | | | |
| P3 | 0 | 1 | 1 | 0 | 6 | 4 | | | |
| P4 | 0 | 0 | 1 | 1 | 4 | 4 | | | |

(8)

- ii) A set of producer processes communicate with a set of consumer processes by means of two buffers, each of which can hold N items of equal size. The producers place single items in one buffer while consumers extract single items from the other buffer. The buffers are exchanged as soon as one becomes full and the other becomes empty.

Write code to be used by the producers and consumers in depositing and extracting items, using (a) semaphores and (b) monitors. (8)

23. a) i) A process accesses the following sequence of pages:

6 3 2 4 3 5 3 1 4 5 3 5 4 2 4 3 2 6 3 2

You have three frames available. For each of the following algorithms, show the pages mapped to all frames after every page access. Also state how many page faults occur for each algorithm?

- I. FIFO
- II. Second-chance (aka Clock)
- III. Least recently used
- IV. Optimal replacement algorithm (8)

- ii) What is a page fault? Describe the actions taken by the Operating System whenever page fault occurs? Explain the Belady's anomaly. What are the techniques where Belady's anomaly does not occur? Prove using the techniques from Q.13a)i). (8)

(OR)

- b) i) State the differences between segmentation and paging? What are their advantages and disadvantages? (8)

ii) Write short notes on:

- I. Virtual memory (8)
- II. Contiguous memory allocation

24. a) i) Explain the various logical structures of file directories with their benefits and limitations. (8)

ii) How can one improve throughput to/from a file system in a computer that has several disk drives? How can one improve the reliability of a file system? Can one do both at the same time? (8)

(OR)

b) Suppose that the head of a moving-head disk with 192 tracks, numbered 0 to 191, is currently serving a request at track 80 and has just finished a request at track 62. The queue of requests is kept in the FIFO order: 119, 58, 114, 28, 111, 55, 103, 30, 75. What is the total number of tracks traversed by head movements needed to satisfy these requests for the following disk-scheduling algorithms?

- i) FCFS.
- ii) SSTF.
- iii) Elevator (SCAN).
- iv) Modified Elevator (C-SCAN). (16)

25. a) i) How some of the concurrency-control schemes in an operating system can be modified for use in a distributed environment. (8)

ii) Explain the schemes for handling deadlock prevention, deadlock avoidance, and deadlock detection in a distributed system. How many class are the Distributed deadlock detection algorithm divided? Elaborate on all of these classes. (8)

(OR)

b) i) Why are transport-level communication services often inappropriate for building distributed applications? (6)

ii) The Linux kernel does not allow paging out of kernel memory. What effect does this restriction have on the kernel's design? What are advantages and disadvantages of this design decision? (10)
