

- (b) What is deadlock? What are the necessary conditions for deadlock to occur? Explain the deadlock prevention method of handling deadlock.

Consider the following information about resources in a system.

- (i) There are two classes of allocatable resource labeled R1 and R2
 - (ii) There are two instances of each resource
 - (iii) There are four processes labeled p1 through p4
 - (iv) There are some resource instances already allocated to processes as follows:
 - One instance of R1 held by p2, another held by p3
 - One instance of R2 held by p1, another held by p4
 - (v) Some processes have requested additional resources, as follows:
 - p1 wants one instance of R1
 - p3 wants one instance of R2
- (1) Draw the resource allocation graph for this system
 - (2) What is the state (runnable, waiting) of each process? For each process that is waiting indicate what it is waiting for.
 - (3) Is this system deadlocked? If so, state which processes are involved. If not, give an execution sequence that eventually ends, showing resource acquisition and release at each step.

Reg. No. :

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Question Paper Code : 71678

B.E./B.Tech. DEGREE EXAMINATION, APRIL/MAY 2017

Fourth/Fifth/Sixth Semester

Computer Science and Engineering

CS 6401 — OPERATING SYSTEMS

(Common Electronics and Communication Engineering, Electronics and Instrumentation Engineering, Instrumentation and Control Engineering, Medical Electronics Engineering, Information Technology)

(Regulations 2013)

Time : Three hours

Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. Consider a memory system with a cache access time of 10 ns and a memory access time of 110 ns — assume the memory access time includes the time to check the cache. If the effective access time is 10% greater than the cache access time, what is the hit ratio H?
2. What are the objectives of operating systems?
3. "Priority inversion is a condition that occurs in real time systems where a low priority process is starved because higher priority processes have gained hold of the CPU" — Comment on this statement.
4. Differentiate single threaded and multi-threaded processes.
5. What is the difference between a user-level instruction and a privileged instruction? Which of the following instructions should be privileged and only allowed to execute in kernel mode?
 - (a) Load a value from a memory address to a general-purpose register.
 - (b) Set a new value in the program counter (PC) register.
 - (c) Turn off interrupts.

6. Will optimal page replacement algorithm suffer from Belady's anomaly? Justify your answer.
7. Suppose that the disk rotates at 7200 rpm. What is the average rotational latency of the disk drive?
8. Differentiate between file and directory.
9. Mention any two features of Linux file systems.
10. Enlist the advantages of using kernel modules in Linux.

PART B — (5 × 13 = 65 marks)

11. (a) (i) Explain the concept of multiprocessor and Multicore organization. (7)
(ii) Discuss about direct memory access. (6)

Or

- (b) (i) Explain the various structures of an operating system. (8)
(ii) Describe system calls and system programs in detail with neat sketch. (5)

12. (a) Consider the following set of processes, with the length of the CPU — burst time in given ms:

Process	Burst Time	Arrival time
P1	8	0.00
P2	4	1.001
P3	9	2.001
P4	5	3.001
P5	3	4.001

Draw four Gantt charts illustrating the execution of these processes using FCFS, SJF, priority and RR (quantum=2) scheduling. Also calculate waiting time and turnaround time for each scheduling algorithms. (13)

Or

- (b) What is a race condition? Explain how a critical section avoids this condition. What are the properties which a data item should possess to implement a critical section? Describe a solution to the Dining philosopher problem so that no races arise. (13)
13. (a) Discuss the given Memory Management techniques with diagrams
(i) Partition Allocation Methods. (7)
(ii) Paging and Translation Look-aside Buffer. (6)

Or

- (b) (i) Describe about free space management on I/O buffering and blocking. (7)
(ii) Discuss the concept of buddy system allocation with neat sketch. (6)

14. (a) (i) Discuss about the various file access methods. (7)
(ii) With neat sketch explain about the
(1) Directory structure
(2) File sharing (6)

Or

- (b) (i) Explain about Kernel I/O subsystem and transforming I/O to hardware operations. (7)
(ii) On a disk with 1000 cylinders, numbers 0 to 999, compute the number of tracks the disk arm must move to satisfy the entire request in the disk queue. Assume the last received was at track 345 and the head is moving towards track 0. The queue in FIFO order contains requests for the following tracks. 123, 874, 692, 475, 105 and 376. Find the seek length for the following scheduling algorithm.

(1) SSTF (2) LOOK (3) CSCAN (6)

15. (a) Explain the concepts of domain name system and multifunction server. (13)

Or

- (b) Write short notes on LINUX kernel and virtualization with neat sketch. (13)

PART C — (1 × 15 = 15 marks)

16. (a) What do you mean by term synchronization? What is Semaphore? Explain how semaphore can be used as synchronization tool. Consider a coke machine that has 10 slots. The producer is the delivery person and the consumer is the student using the machine. It uses the following three semaphores:

semaphore mutex

semaphore fullBuffer /* Number of filled slots */

semaphore emptyBuffer /* Number of empty slots */

- (i) Write pseudo code for delivery_person() and student()
(ii) What will be the initial values of the semaphores?
(iii) Write a solution that guarantees the mutual exclusion and has no deadlocks.

Or