PART – C (1×15=15 Marks)

16. a) Consider the following resource allocation graph :

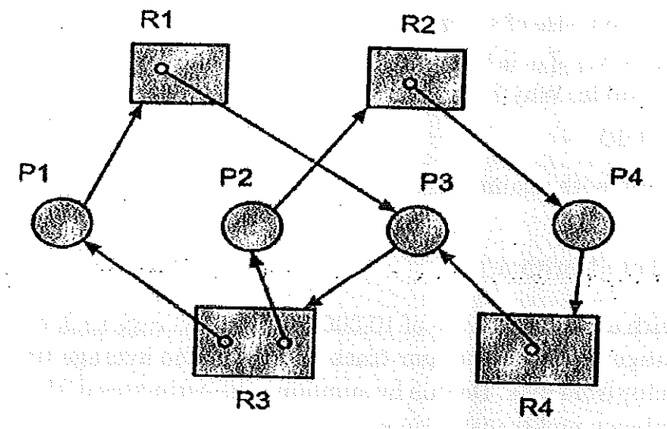

Determine if there is a deadlock. If so, indicate the processes and resources involved. Show how the deadlock can be resolved through addition of resources. If not, argue why this is the case, i.e. there is no deadlock. In either case, provide a feasible sequence of processes to show completion.

(OR)

b) Explain in detail how UNIX Virtual File System has been Implemented.

Reg. No. :

Question Paper Code : 91397

B.E./B.Tech. DEGREE EXAMINATIONS, NOVEMBER/DECEMBER 2019
Fourth/Fifth/Sixth Semester
CS6401 – OPERATING SYSTEMS
(Common to : Computer Science and Engineering/Electronics and Communication Engineering/Electronics and Instrumentation Engineering/Instrumentation and Control Engineering/Medical Electronics/Information Technology)
(Regulations 2013)
(Also Common to PTCS 6401 – Operating Systems – for Third Semester Computer Science and Engineering – Regulations 2014)

Time : Three Hours Maximum : 100 Marks

Answer ALL questions

PART – A (10×2=20 Marks)

1. What are the pros and cons of Microkernals Operating system Structures ?
2. Why it is important for scheduler to distinguish I/O bound programs from CPU bound programs ?
3. List out the various process states available.
4. What is semaphore ? Explain the two primitive operations of a semaphore.
5. State the differences between static and dynamic memory allocation.
6. When does a page fault occur ? Explain various page replacement strategies/ algorithms.
7. Suppose that a disk drive has 5000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 143 and the previous was at cylinder 125. The queue of pending requests, in FIFO order is: 86, 1470, 913, 1774, 948, 1509, 1022, 1750, 130. Starting from the current head position, what is the total distance (in cylinders) that the disk arm moves to satisfy all the pending requests for SCAN.
8. What type of file accessing method can be used for batch and payroll applications ?
9. Define latency, transfer and seek time with respect to disk I/O.
10. List out the features of Linux OS.

PART – B (5×13=65 Marks)

11. a) Describe the features of the following types of operating systems structures
    i) MS-DOS Layer Structure
    ii) Layered Approach
    iii) Microkernel System Structure
    iv) Modules.

(OR)

b) What are system calls ? How do system calls help user programs to interact with the OS ? Explain.

12. a) Assume some OS needs to schedule four processes using different scheduling algorithms. For each process, the following information shows its burst time (processing time), priority (The lower the number, the higher the priority) and arrival time.

Processes	Burst Time	Priority	Arrival Time
P1	12	3	0
P2	6	4	2
P3	4	1	4
P4	18	2	6

What is the Average Waiting Time of those processes for each of the following schedule algorithms ? (Draw a Gantt Chart for each algorithm.)
    i) First Come First Serve (FCFS)
    ii) Non-preemptive Shortest Job First (NP-SJF)
    iii) Preemptive Shortest Job First (P-SJF)
    iv) Priority Scheduling
    v) Round-Robin (scheduling time quantum is 5 time units)

(OR)

b) What is a critical section ? Explain readers and writers problem with semaphore.

13. a) A computer system has a 36-bit virtual address space with a page size of 8 K, and 4 bytes per page table entry.
    i) How many pages are in the virtual address space ?
    ii) What is the maximum size of addressable physical memory in this system ?
    iii) If the average process size is 8 GB, would you use a one-level, two-level or three-level page table. Why ?

(OR)

b) Explain the following with example :
    i) Thrashing (6)
    ii) Page replacement algorithm. (7)

14. a) Consider a disk with a rotational rate of 10,000 RPM, an average seek time of 8 ms, and an average of 500 sectors per track. Estimate the average time to read a random sector from disk. Do this by summing the estimates of the seek time, rotational latency and transfer time.

(OR)

b) Elaborate on the various File allocation methods.

15. a) A multicore processor is said to be sequentially consistent if all loads and stores appear to occur in some global total order that is consistent with program order in every core. Sadly, most modern processors are not sequentially consistent: memory accesses can appear to occur in different orders from the perspective of different cores or even in circular order
    i) Give a possible reason why different cores might see stores in different orders.
    ii) All non-sequentially consistent machines provide special (expensive) instructions that can, when desired, be used to force a memory access to be seen everywhere at once, after all previous accesses of the same core and before all subsequent accesses of the same core. Suggest how the programmer, language, and/or compiler might use such instructions to achieve the illusion of sequential consistency.

(OR)

b) Discuss how the following pairs of scheduling criteria conflict in certain settings.
    i) CPU utilization and response time
    ii) Average turnaround time and maximum waiting time
    iii) I/O device utilization and CPU utilization.