[Total No. of Questions - 9] [Total No. of Printed Pages - 4]
(2123)

1414

B. Tech 5th Semester Examination

Principles of Operating System (O.S.)

CS-5001

Time : 3 Hours **Max. Marks : 100**

The candidates shall limit their answers precisely within the answer-book (40 pages) issued to them and no supplementary/continuation sheet will be issued.

Note : Attempt five questions in all selecting one from each of the Sections A, B, C & D. Section E is compulsory.

SECTION - A

1. (a) What is the main advantage of the microkernel approach to system design? How do user programs and system services interact in microkernel architecture? What are the disadvantages of using the microkernel approach? **(10)**

   (b) Assume you have the following jobs to execute with one processor, with the jobs arriving in the order listed here:

i	0	1	2	3	4
T(pi)	80	20	10	20	50

   a. Suppose a system uses SJF scheduling. Create a Gantt chart illustrating the execution of these processes?

   b. What is the turnaround time for process p4?

   c. What is the average wait time for the processes? **(10)**

2. (a) What is the purpose of the command interpreter? Why is it usually separate from the kernel? Would it be possible for the user to develop a new command interpreter using the system-call interface provided by the operating system? **(10)**

   (b) What is the purpose of interrupts? What are the differences between a trap and an interrupt? Can traps be generated intentionally by a user program? If so, for what purpose? **(10)**

SECTION - B

3. (a) How are the access-matrix facility and the role-based access-control facility similar? How do they differ? Discuss the strengths and weaknesses of implementing an access matrix using access lists that are associated with objects. **(10)**

   (b) 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 using the banker's algorithm:

   a. What is the content of the matrix Need?

   b. Is the system in a safe state?

   c. If a request from process P1 arrives for (0, 4, 2, 0), can the request be granted immediately? **(10)**

4. (a) Explain the Banker's Algorithm. Explain with the help of an example, how can you prevent deadlock by preventing circular wait condition. **(10)**

   (b) What are the advantages and disadvantages of a system providing mandatory locks instead of providing advisory locks whose usage is left to the users' discretion? **(10)**

SECTION - C

5. (a) Compare the main memory organization schemes of contiguous-memory allocation, pure segmentation, and pure paging with respect to the following issues:

   a. external fragmentation

   b. internal fragmentation

   c. ability to share code across processes **(10)**

   (b) Explain the concept of paging and segmentation. Why are segmentation and paging sometimes combined into one scheme? **(10)**

6. (a) Consider a demand-paging system with the following time-measured utilizations:

CPU utilization 20%, Paging disk 97.7% Other I/O devices 5%. Which (if any) of the following will (probably) improve CPU utilization? Explain your answer.

(i) Install a faster CPU.

(ii) Install a bigger paging disk,

(iii) Increase the degree of multiprogramming,

(iv) Decrease the degree of multiprogramming,

(v) Install more main memory.

(vi) Install a faster hard disk or multiple controllers with multiple hard disks.

(vii) Add prepaging to the page fetch algorithms.

(viii) Increase the page size. **(16)**

(b) What is the cause of thrashing? How does the system detect thrashing? Once it detects thrashing, what can the system do to eliminate this problem? **(4)**

SECTION - D

7. (a) Compare the performance of C-SCAN and SCAN scheduling, assuming a uniform distribution of requests. Consider the average response time (the time between the arrival of a request and the completion of that request's service), the variation in response time, and the effective bandwidth, How does performance depend on the relative sizes of seek time and rotational latency? **(10)**

(b) Explain the following with necessary example.

(i) synchronization hardware (ii) Semaphore. **(10)**

8. (a) 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 request 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 each of the following disk-scheduling algorithms?

i. SSTF

ii. C-SCAN

iii. LOOK **(10)**

(b) Explain critical section problem. A file is to be shared among different processes, each of which has a unique number. The file can be accessed simultaneously by several processes, subject to the following constraint: The sum of all unique numbers associated with all the processes currently accessing the file must be less than n. Write a monitor to coordinate access to the file. **(10)**

SECTION - E

9. (a) When a process requests for I/O, how many process switches take place?

(b) State three disadvantages of placing functionality in a device controller, rather than in the kernel.

(c) What is advantage of pipes?

(d) What is the purpose of paging the page tables?

(e) Discuss the reasons why the operating system might require accurate information on how blocks are stored on a disk.

(f) What are the advantages of supporting memory mapped I/O to device control registers?

(g) What are different block based file organizations?

(h) Explain Inode Structure in UNIX?

(i) Discuss the strengths and weaknesses of implementing an access matrix using access lists that are associated with objects.

(j) What is internal fragmentation? **(10×2=20)**