Reg. No.										

MANIPAL INSTITUTE OF TECHNOLOGY

MANIPAL

A Constituent Institution of Manipal University

n

V SEMESTER B.TECH. (COMPUTER SCIENCE AND ENGINEERING)

END SEMESTER EXAMINATIONS, NOV/DEC 2016

SUBJECT: OPERATING SYSTEMS [CSE 3102]

REVISED CREDIT SYSTEM

(03/12/2016)

Time: 3 Hours MAX. MARKS: 50

Instructions to Candidates:

- Answer **ALL** the questions.
- Missing data may be suitable assumed.

1A. 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? **2 M**

1B. (i) Explain the **six** different purpose of system call in an operating system? For each of the system calls: **fork** and **exec**, give a condition that causes it to fail.
(ii) Can the

count = write(fd, buffer, nbytes)

call return any value in count other than **nbytes**? If so, why? **5 M**

1C. (i) How many processes will be created when the following program is executed? Explain your answer.
(ii) Assume that all fork system calls are successful. What will be printed?

```
main()
{
    int i, pid;
    for (i = 1; i ≤ 3; i++)
        pid = fork( );
    if (pid > 0)
        printf("In parent %d \n", i );
    else
        printf("In child %d \n", i );
}
```

3M

2A. What are the three differences between user-level threads and kernel-level threads? Under what circumstances is one type is better than the other? What is the essential cause of the difference in cost between a context switch for kernel-level threads and a switch that occurs between user-level threads? **3M**

2B. (i) Explain **five** different scheduling and optimization criteria.
(ii) 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 possible different schedules are there? Give a formula in terms of ***n***. **3M**

2C. Consider the following set of four process in Table 2C, with arrival time, length of CPU burst given in msec. Draw the grant chart that illustrate the execution of these processes using RR scheduling algorithm (quantum = 10)

Table 2C

Process	Arrival Time	Burst time
P_1	0	60
P_2	3	30
P_3	4	40
P_4	9	10

Calculate average waiting time and turnaround time for these processes. **4M**

3A. State first readers-writers problem. Write pseudo-code semaphore solution for the first readers-writers problem with required data structures. Is this solution leads to starvation? Justify your answer. **4M**

3B. Write an algorithm for bounded-waiting mutual exclusion with **TestAndSet()** instruction. Does it satisfy all the three requirements of the critical section problem? Justify your answer. **3M**

3C. Consider the following snapshot of a system in Table 3C.

Table 3C

	Allocation	Max	Available
	A B C D E	A B C D E	A B C D E
P_0	1 0 2 1 1	1 1 2 1 3	0 0 **x** 1 1
P_1	2 0 1 1 0	2 2 2 1 0	
P_2	1 1 0 1 0	2 1 3 1 0	
P_3	1 1 1 1 0	1 1 2 2 1	

Answer the following questions using the banker algorithm:
(i) What is the content of the matrix Need?
(ii) Find the smallest value of **x** for which the system in a safe state? Justify your answer. **3M**

4A. Define Hit ratio. Write the formula for effective access time (EAT). Consider a paging system with the page table stored in memory.
(i) If a memory reference takes 200 nanoseconds, how long does a paged memory reference take?
(ii) If we add associative registers, and 75 percent of all page-table references are found in the associative registers, what is the EAT? (Assume that finding a page-table entry in the associative registers takes zero time, if the entry is there.) **3M**

4B. What is Belady's anomaly? Suppose there are 8 virtual pages and 4 page frames. Determine the number of page faults that will occur with the reference string

0, 2, 3, 1, 4, 2, 5, 6, 0, 1, 3, 2, 4, 7, 1, 2,

if the page frames are initially empty, using each of the following page replacement algorithms: (i) LRU (ii) Optimal. Show the page frame at each step. Does LRU and Optimal suffer from Belady's anomaly? Explain your answer. **4M**

4C. Mention six comparisons of demand paging with segmentation. **3M**

5A. What are the advantages and disadvantages of providing mandatory locks instead of advisory locks whose usage is left to users' discretion? **2M**

5B. (i) Mention two disadvantages of SSTF disk scheduling. Is SSTF optimal? Explain your answer.
(ii) Suppose that a disk drive has 5,000 cylinders, numbered 0 to 4999. The drive is currently serving a request at cylinder 143, and previous request was at cylinder 125. The queue of pending requests, in FIFO order, is:

86, 1470, 913, 1774, 948, 1509, 1022, 1750, 130 **5M**

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? Depict the sequence in graph.

(i) SCAN (initially moving upward) (ii) C-SCAN (initially moving upward)

5C. (i) Discuss the strengths and weaknesses of implementing an access matrix using capabilities that are associated with domains.

(ii) What are the primary goals of the conflict resolution mechanism used by the Linux kernel for loading kernel modules? **3M**

****