Reg. No.										

MANIPAL INSTITUTE OF TECHNOLOGY

MANIPAL

A Constituent Institution of Manipal University

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

END SEMESTER EXAMINATIONS, NOV/DEC 2017

SUBJECT: OPERATING SYSTEMS [CSE 3102]

REVISED CREDIT SYSTEM
(24 /11/2017)

Time: 3 Hours MAX. MARKS: 50

Instructions to Candidates:

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

1A. (a) The issue of resource utilization shows up in different forms in different types of operating systems. List what resources must be managed carefully in the following settings:
(i) Mainframe or minicomputer systems
(ii) Workstations connected to servers
(iii) Handheld computers
(b) Mention three main advantages of the microkernel approach to system design? **3M**

1B. What are the two models of inter-process communication? Mention one strength and one weakness of the two approaches? **3M**

1C. (a)
```
int counter = 0;
int main( ) {
        int i;
        for (i = 0; i < 2; i ++){
              fork();
              counter++;
              printf("counter = %d\n", counter);
        }
        printf("counter = %d\n", counter);
        return 0;
}
```
(i) How many times is the value of counter printed?
(ii) What is the value of counter printed in the first line?
(iii) What is the value of counter printed in the last line?
(b) Describe the difference between the fork() and clone() Linux system calls. **4M**

2A. What are the two different ways in which a thread library could be implemented? Consider two scenarios: (i) two user threads are mapped into one kernel thread, and (ii) each of the two user threads is mapped into a unique kernel thread. Which achieves better concurrency in execution? Why? **2M**

2B. What is meant when saying that a scheduling algorithm could *result in starvation?* Clearly mention the algorithms listed below that could result in starvation. You can assume that each process will use the CPU for a finite burst before performing I/O.
(i) First-come, First-Served (ii) Round Robin (iii) Shortest Job First
(iv) Shortest Remaining Time First (v) Priority (vi) Preemptive Priority **2M**
(vii) Multilevel Feedback Queue Scheduling

2C. With the help of a Gantt Chart (Use the data given below in Table 2C) calculate waiting time and turnaround time for each processes, hence compute average waiting time and average turnaround time for a **SRTF** scheduling algorithm (in case of tie use FCFS). Assume the time in msec.

Table 2C

Process	Arrival Time	Burst Time
P1	8	1
P2	5	1
P3	2	7
P4	4	3
P5	2	8
P6	4	2
P7	3	5

6M

3A. Consider the following algorithm in Table 3A, that provides 2-process solution to the critical section problem.

Table 3A

```
flag[0] = false;
flag[1] = false;
```

P0:	P1:
0: **while** (true) {	0: **while** (true) {
1: flag[0] = true;	1: flag[1] = true;
2: **while** (flag[1]) {	2: **while** (flag[0]) {
3: flag[0] = false;	3: flag[1] = false;
4: **while** (flag[1]) {	4: **while** (flag[0]) {
5: no-op;	5: no-op;
6: }	6: }
7: flag[0] = true;	7: flag[1] = true;
8: }	8: }
9: critical section	9: critical section
10: flag[0] = false;	10: flag[1] = false;
11: remainder section	11: remainder section
12: }	12: }

Specify which of the following requirements are satisfied or not by this algorithm. Explain why or why not.
(i) Mutual Exclusion (ii) Progress (iii) Bounded Wait **4M**

3B. Consider a system has P processes. Each process need a maximum of m resources and a total of r resources available. What condition must hold to make the system deadlock free? **2M**

3C. Consider the following snapshot of a system in Table 3C. **4M**

Table 3C

	Allocation	Max	Available
	A B C D	A B C D	A B C D
P_0	0 0 1 2	0 0 1 2	1 5 2 0
P_1	1 0 0 0	1 7 5 0	
P_2	1 3 5 4	2 3 5 6	
P_3	0 6 3 2	0 6 5 2	
P_4	0 0 1 4	0 6 5 6	

Answer the following questions using the banker's algorithm:

(i) What is the content of the matrix Need?

(ii) Is the system in the safe state?

(iii) If a request from process P_1 arrives for (0, 4, 2, 0), can the request be granted immediately?

4A. What is the working set model? 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? **2M**

4B. Consider a logical address space of 8 pages; each page is 2048 byte long, mapped onto a physical memory of 64 frames.

(i) How many bits are there in the logical address and how many bits are there in the physical address?

(ii) A 6284 bytes program is to be loaded in some of the available frames= {10, 8, 40, 25, 3, 15, 56, 18, 12, 35}. Show the contents of the program's page table.

(iii) What is the size of the internal fragmentation?

(iv) Convert the following logical addresses 2249 and 5245 to physical addresses. **4M**

4C. Suppose there are 5 virtual pages and 4 page frames. Determine the number of page faults that will occur with the reference string

1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5

if the page frames are initially empty, using second-chance algorithm. Show the reference bit and page frames at each step. **4M**

5A. Suppose that a disk drive has 1000 cylinders, numbered 0 to 999. The drive is currently serving a request at cylinder 100, and the previous request was at cylinder 50. The queue of pending requests, in FIFO order, is: 75, 400, 200, 800, 900, 500, 25. 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 a graph.

(i) FCFS (ii) SSTF **4M**

5B. (a) Mention two consistency semantics of UNIX file system.

(b) Mention one strength and one weaknesses of implementing an access matrix using access lists that are associated with objects. **3M**

5C. Mention the algorithm used by Linux scheduler. With neat diagram briefly explain various features of Linux scheduler. **3M**