BE 5th Semester(CST) End-Semester Examination 2012

Bengal Engineering and Science University Shibpur,
Howrah-711103, India.
Operating system (CS-501)

Time: 3hrs. Full Marks: 70

Use separate answer sheet for each half. Credit will be given to precise answer.

1st Half

Answer question no 1 *and any three from the rest.*

1.

(a) Differentiate between process and thread.

(b) What do you understand by degree of multiprogramming? Who actually controls it?

(c) Show that improper use of `semaphore` can lead a system to enter deadlock. [1+2+2]

2.

(a) What happens to the parent process ID of a child process when its parent process terminates before it(i.e. the child process)?

(b) How does the sharing of files occur across a `fork` system call? Explain it with a diagram.

(c) Write down the tests that a Unix `kernel` goes through to determine if a process can access a file.

(d) What are the disadvantages of pipe? How is it removed using a named pipe? [2+3+3+2]

3.

(a) How many different types of solution exist to solve critical section problem. Write down the assumptions taken by these solutions.

(b) Show that if *wait* and *signal* operations are not executed atomically then mutual exclusion may be violated.

(c) Implement the following scenario using `semaphore`.

There are two processes, P_1 and P_2 and they share a critical section(say C). I want P_1 to execute C only after P_2 executes it 5 times consecutively and they may continue to repeat the pattern.

Note that any of P_1 and P_2 can execute C first. So, if P_1 execute it first, P_2 will execute 5 times next and if P_2 execute C first, it has to execute it 5 times before P_1 gets its turn and then they may continue to repeat the pattern.

So the pattern could be $P_1, P_2, P_2, P_2, P_2, P_2, P_1, P_2, P_2, P_2, P_2, P_2, P_1 \ldots$
or $P_2, P_2, P_2, P_2, P_2, P_1, P_2, P_2, P_2, P_2, P_2, P_1 \ldots$ [3+3+4]

4.

(a) What are the necessary conditions for deadlock? How do you prevent deadlock to occur?

(b) Consider a system consisting of N resources of the same type that are shared by $N-1$ processes, each of which needs at most two resources, show that the system is deadlock-free. In the above scenario if the number process becomes N, will it still be deadlock-free?

(c) What is `starvation`? Can a system detect that some of its process are `starving`? Explain how the system deal with `starvation` problem. [3+3+4]

5.

(a) Write down the difference between preemptive and nonpreemptive CPU-scheduling?

(b) Consider the following set of processes, with the length of the CPU-burst time given in milliseconds:
The processes are assumed to have arrived in the order P_1, P_2, P_3, P_4, P_5 all at time 0.

Process	Burst Time	Priority
P_1	10	3
P_2	1	1
P_3	2	3
P_4	1	4
P_5	5	2

(i) Draw Gantt charts illustrating the execution of these processes using SJF, and a nonpreemptive priority(smaller number implies higher priority) scheduling.

(ii) What are the average waiting time in each of the algorithm mentioned above?

(c) Consider the following preemtive priority scheduling algorithm based on dynamically changing priorities. Larger number implies higher priority. When a process is waiting for the CPU(in the ready queue but not running), its priority changes at a rate α; when it is running, its priority changes at a rate β. All processes are given priority 0 when they enter the ready queue. The parameters α and β can be set to give many different algorithms.

(i) What is the algorithm that results from $\alpha < \beta < 0$?

(ii) What is the algorithm that results from $\beta > \alpha > 0$? [2+5+3]

2nd Half

Answer question no 6 *and any two from the rest. One mark is reserved for neatness.*

6.

(a) A moving hard disk has 200 tracks (0 to 199). The head of the disk is at track 143 and has just finished a request at track 125. The queue of requests is kept in the FIFO order: 86, 147, 91, 177, 94, 150, 102, 175, 130.
What is the total number of head movement to satisfy these requests for the following disk scheduling algorithms? (i)FCFS (ii)SSTF (iii)SCAN (iv)LOOK (v)C-SCAN

(b) Consider the following segment table:

Segment	Base	Length
0	219	600
1	2300	14
2	90	100
3	1327	580
4	1952	96

What physical addresses are generated for the following logical addresses? (i) 0, 430 (ii) 1, 10 (iii) 2, 80 (iv) 3, 400 (v) 4, 90

[6+6]

7.

(a) Why starvation may occur in SSTF scheduling algorithm?

(b) Why fragmentation occurs in case of contiguous allocation? How it can be removed?

(c) In what situation the use of memory as a RAM disk is more useful than using it as a disk cache?

(d) How audit log helps to improve security of a system?

[2+(3+2)+2+2]

8.

(a) How stack can be used to implement LRU page replacement algorithm?

(b) What is working set? How it helps to prevent thrashing?

(c) Why sometimes segmentation and paging are combined into one scheme?

(d) What is need-to-know principle? [3+(2+2)+3+1]

9.

(a) What are the main differences between capability list and access list?

(b) How global table can be used to implement access matrix?

(c) How the information transmitting over unreliable links can be protected?

(d) Explain system threat and program threat. [3+3+2+3]

10.

Discuss briefly about any two of the following:

(a) Internal fragmentation

(b) Optimal page replacement algorithm

(c) Swap space management

(d) Acyclic graph directory

[5.5x2]