Homework 3

Due Date : May 18, 1999 **Points** : 100

Short-Answer Questions

These can be answered in a sentence or two, and are intended to reinforce important points.

1. (*4 points*) Describe the differences between short term, medium term, and long term scheduling.
2. (*3 points*) State the difference between preemptive and non-preemptive scheduling algorithms.
3. (*3 points*) What is "device independence"? (Tanenbaum, problem 3-7)

Long-Answer Questions

These questions require some thought and longer answers than the short-answer questions. They are intended to have you use the concepts discussed in class, to be sure you understand them and can work with them.

4. (*30 points*) Consider the following process arrival list:

job name	arrival time	service time
A	0	10
B	2	1
C	3	2
D	7	7
E	8	5

   For each process, give the time(s) that the process acquires the CPU, relinquishes the CPU, and the turnaround time, the waiting time, and the response ratio for each of the following job scheduling algorithms:
   a. First Come, First Serve
   b. Shortest Job Next
   c. Shortest Remaining Time Next
   d. Highest Response Ratio Next
   e. Round Robin with a quantum of 1
   f. Multilevel Feedback Queue with 3 queues, the first two being round robin with quanta of 2 and 4, respectively, and the lowest being first come first serve.
5. (*20 points*) An operating system designer has proposed a multilevel feedback queueing scheme in which there are five levels. The quantum at the first level is 0.5 seconds, and each lower level has a quantum twice the size of the quantum at the previous level. A process cannot be pre-empted until its quantum expires, and then it is moved to the next lower queue (unless it is in the lowest queue, of course). If the process does I/O, upon completion it is returned to the queue from which it left. The system runs both batch and interactive jobs, and these, in turn, consist of both CPU-bound and I/O-bound jobs.
   a. Why is this scheme deficient?
   b. What minimal changes would you propose to make the scheme more acceptable for its intended job mix?
6. (*20 points*) This question asks you to compare different disk scheduling policies.
   a. Under very light loads, all the disk scheduling policies we have discussed degenerate into which policy? Why?
   b. Will requests scheduled by a FCFS disk scheduling policy ever have a lower mean waiting time than those scheduled by a SCAN policy? Than those scheduled by a SSTF policy? Justify your answers. (*Hint*: consider a system on which a seek takes $0.5 + 0.4T$ msec, where T is the number of cylinders moved. Then assume the arm is initially at cylinder 100, the disk has 200 cylinders, and the arm is moving inward.)
7. (*20 points*) To what extent does the theory of disk-head scheduling apply to magnetic tapes?

Extra Credit

8. (*10 points*) Tanenbaum, and many other authors, refer to the LOOK and SCAN algorithms as "elevator algorithms." What is the major conceptual difference between disk scheduling and "elevator scheduling"? (*Hint*: are we trying to minimize elevator movement?)