

Homework #3

Due Date: Monday, February 21, 2000 at 11:59PM

Points: 80 regular, 10 extra credit

Short-Answer Questions

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

- (6 points) Are disk scheduling policies other than first come, first serve (FCFS) useful when at most one user at a time will be using the computer? Give a one-sentence justification for your answer.
- (6 points) In section 2.2.4 of the text, a situation with a high-priority process, H , and a low-priority process L was described, which led to H looping forever. Does the same problem occur if round robin scheduling is used rather than priority scheduling? Why or why not? (text, chapter 2, problem 12)
- (6 points) Of the policies we have discussed in class, which requires the least amount of overhead time (the time needed to compute which process runs next)? Justify your answer.

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.

- (25 points) Measurements of a certain system have shown that the average process runs for a time T before blocking on I/O. A process switch requires a time S , which is effectively wasted (*i.e.*, it is overhead time). For round robin scheduling with a quantum Q , give a formula for the processor utilization for each of the following:
 - $Q = \infty$
 - $Q > T$
 - $S < Q < T$
 - $Q = S$
 - Q nearly 0
 (text, chapter 2, problem 21)
- (12 points) In section 2.6.8, Tanenbaum writes that “MINIX uses a multilevel scheduling algorithm ...”
 - Is the MINIX processor scheduling algorithm a multilevel feedback queue algorithm? If not, what distinguishes it from a MLFQ algorithm?
 - Are the quanta fixed *within* each level, or do they vary on a per-process basis? If they are fixed, are the quanta fixed between levels, or does each level have a different quantum? If not, does the calculation of the quanta depend on the level of the process? Why do you think Tanenbaum made these choices?

- (25 points) Disk requests come into the disk drive for cylinders

23 89 132 42 187 165 21 34 101 102

in that order. A seek takes $0.5 + 0.4T$ msec, where T is the number of cylinders moved. Calculate the mean and variance of the waiting time for each of the following disk scheduling policies, assuming the arm is initially at cylinder 100, the disk has 200 cylinders, and the arm is moving inward:

- first come, first serve (FCFS)
- shortest seek time next (SSTF)
- SCAN
- LOOK
- C-SCAN

Extra Credit

- (10 points) After receiving a DEL (SIGINT) character, the MINIX device driver discards all output currently queued for that terminal. Why? (text, chapter 3, problem 30)