CSCI 5105: Foundations of Modern Operating Systems
Spring 2008
Homework Assignment 0

Note: This assignment will not be graded. However, you are **strongly** advised to solve this assignment to get an idea of the level of background expected in the course. If you are unable to solve this assignment or are unfamiliar with the concepts, you must talk to the instructor.

1. Answer the following questions briefly:
   (a) How does a system call work?
   (b) What kinds of information does a Process Control Block (PCB) contain?
   (c) Which process scheduling algorithm is optimal in terms of average waiting time? Why is it not typically used in conventional operating systems?
   (d) What is a semaphore and what operations can be performed on it?
   (e) Describe one approach to avoid deadlocks when using locks for synchronization between multiple processes.
   (f) What is the LRU Page Replacement Algorithm?
   (g) What is thrashing and why does it occur?
   (h) Most OS's maintain a file buffer in main memory. What is the benefit?
   (i) What is an inode and what kind of information does it contain?
   (j) What is the difference between a soft and hard link in a file system?

2. Consider the following set of processes, with the length of the CPU burst for each: P1 (10), P2 (1), P3 (2), P4 (1), and P5 (5). Assume that the processes arrive in the order P1, P2, P3, P4, P5, all at time 0.
   (a) What is the turnaround and waiting time for each process using FCFS?
   (b) What is the turnaround and waiting time for each process using Round Robin scheduling (quantum=1)?

3. A barbershop consists of a waiting room with n chairs and a barber room with one barber chair. If there are no customers to be served,

the barber goes to sleep. If a customer enters the barbershop and all chairs are occupied, then the customer leaves the shop. If the barber is busy but chairs are available, then the customer sits in one of the free chairs. If the barber is asleep, the customer wakes up the barber.

Write a program (pseudocode) to coordinate the barber and the customers.

4. Consider a disk with 100 tracks numbered 1 to 100. Suppose the current disk head position is on track 49 and is moving towards track 100. Suppose it has pending disk I/O requests to the following tracks: 42 78 3 51 99. In what order will it visit the tracks using the following disk scheduling algorithms:

   (a) Shortest seek time first (SSTF)

   (b) SCAN (Elevator)

   (c) What is the total number of tracks seeked in each case?

   (d) If using SSTF, should more frequently accessed data be stored towards the middle tracks or the outermost/innermost tracks?