CIS 25 Test 1 Fall, 2009 Name_____

1. a. What does it mean to say that an OS is interrupt-driven? (2 pts)
b. On a PC, how does the system find the appropriate interrupt handler and switch to it? Include all tables and how they are used. (7 pts)
c. What does it mean to "ignore" an interrupt? Which interrupts can be ignored? (4 pts)

2. a. What is a thread? (3 pts)
b. Identify and explain the 3 threading models. (6 pts)
c. What is an advantage of system-level threads over user-level threads? (4 pts)
d. What is an advantage of user-level threads over system-level threads (4 pts)

3. a. What is the problem with using the Shortest Job Next algorithm for scheduling? What is the solution to the problem in a batch system? (5 points)
b. Suppose you used a non-preemptive scheduling algorithm in a system that had both batch and interactive jobs. How would this algorithm affect response time? Why? (4 pts)
c. Draw the process state diagram (from the book, not the project), labeling all states. Show all the arcs. When do processes move from one state to another? (7 pts)

4. a. What is the job of the long-term scheduler? (3 pts)
b. Explain the difference between CPU-bound and I/O bound jobs. (4 pts)

5. a. There are four jobs on the ready queue. Job A requires 8 time units, job B requires 12 time units, and job C requires 4 time units. Assume the jobs come in at time zero, in the order B, A, C.
Draw a Gantt diagram for each of the following algorithms, illustrating when each of the jobs is running: (9 points)
i. First Come First Served
ii. Shortest Job Next
iii. Round-Robin (time quantum = 2)
b. What is the dilation for job B under SJN? (3 pts)
c. What is the waiting time for job C under Round Robin? (3 pts)
d. Assume that you are using Shortest Remaining Time Next. In addition to jobs A, B, and C that entered at time 0, job D enters the system at time 6. Job D requires 5 time units. Draw a Gantt diagram (starting from time 0) to show how to schedule the four jobs using SRTN. (4 pts)

6. Assume you have an MVP system with free space as follows: at address 48K a free space of size 12K, and at address 22K a free space of size 19K.
a. Show the First Fit free space table. (2 pts)
b. Show the Best Fit free space table. (2 pts)
c. The following sequence of processes enters: 12K, 8K, 11K.
i. Using the First Fit free space table, at what address will each job be placed using First Fit? (Your answer will be addresses, not a table.) (3 pts)
ii. Using the Best Fit free space table, at what address will each job be placed using Best Fit? (Your answer will be addresses, not a table.) (3 pts)
d. What problem does coälescing holes solve? Why isn't it done often? (4 pts)

7. In a paged system with 1000-unit pages, explain the process by which logical address 5780 is translated into a physical address, if the page is stored in frame 9 of memory. Include an explanation or diagram of all tables and registers involved and how they are used, including all validity checks. (8 pts)

8. a. What is the major disadvantage of segmentation? (3 pts)
b. Explain how paged segmentation solves the problem of part b. (3 pts)