HW1::CMSC421::Spring 2009::Due 10 March 2009::Total Pts 100

0) (12 pts) When a process performs a blocking send, it must receive an acknowledgement message to unblock. What problem might result from this scheme, and how can it be avoided?

1) (12 pts) Define the difference between preemptive and non-preemptive scheduling. State why strict non-preemptive scheduling is unlikely to be used in a computer center.

2) (12 pts) Suppose that a scheduling algorithm (at the level of short-term cpu scheduling) favors those programs which have used little processor time in the recent past. Why will this algorithm favor I/O bound programs and yet not permanently starve cpu-bound programs?

3) (12 pts) Can a system detect that some of its processes are starving? If the answer is yes, explain how. If the answer is no, explain how the system may deal with the starvation problem.

4)(12 pts) A producer P and a consumer C have a buffer of capacity 2 between them. The following sequence of actions occurs:

Step	**Actor**	**Item**
1	P	1
2	P	2
3	P	3
4	C	1
5	C	2
6	P	4
7	P	5
7	C	3
9	C	4
10	C	5
11	P	6
12	C	6

After which steps does P and/or C sleep? What is in the buffer after each step (draw a table with the first column having the heading “After Step” and the second column having “Buffer Contents”)?

5) (20 pts) Consider a computing system with 13 tape drives. All jobs running on this system require a maximum of 5 tape drives to complete, but they each run for long periods of time with just 4 drives and request the fifth one only at the very end of the run. The job stream is endless.
If your operating system supports a very conservative device allocation policy no job will be started unless all tapes required have been allocated to it for the duration of its run.
i) What is the maximum number of jobs that can be active at once?
ii) What are the minimum and maximum number of tape drives that may be idle as a result of this policy?

6) (20 pts) Our textbook claims (page 190) that SJF is "provably *optimal*, in that it gives the mimimum average waiting time for a given set of processes." Prove this statement in its most general form.