

CS4410 - Fall 2008
Homework 2
Due September 23, 11:59PM

Q1. Explain what goes wrong in the following version of Dekker's Algorithm:

```
CSEnter(int i)
{
    inside[i] = true;
    while(inside[j])
    {
        inside[i] = false;
        while(turn == j) continue;
        inside[i] = true;
    }
}

CSExit(int i)
{
    turn = j;
    inside[i] = false;
}
```

Q2. Round-robin schedulers normally maintain a list of all runnable processes, with each process occurring exactly once in the list. What would happen if a process occurred twice in the list? Can you think of any reason for allowing this?

Q3. Five batch jobs A through E, arrive at a computer center at almost the same time. They have estimated running times of 11, 6, 2, 4, and 8 minutes. Their (externally determined) priorities are 3, 5, 2, 1, and 4, respectively, with 5 being the highest priority. For each of the following scheduling algorithms, determine the mean process turnaround time.

Ignore process switching overhead.

- (a) Round-robin
- (b) Priority scheduling
- (c) First come, First served (run in order 11, 6, 2, 4, 8)
- (d) Shortest job first

For (a), assume that the system is multiprogrammed, and that each job gets its fair share of the CPU. For (b) through (d) assume that only one job runs at a time, until it finishes.

All jobs are completely CPU bound. Assume that the time quantum of the scheduler is 1 minute.

Q4. Can we use the test-and-set instruction to implement a lock on a multiprocessor environment? Explain why or why not.

Q5. Consider the environment where many user-level threads are mapped to a single kernel thread (i.e., the many-to-one model). Describe a situation the one-to-one model outperforms this model.