CS 377 – Operating System

Discussion Session 3 Questions

Name: _____

Write your answers individually without consulting your notes, the textbook, or the Internet. Be succinct (complete sentences not necessary).

1. The lifecycle of a process consists of five execution states, which are (in no particular order): running, terminated, new, ready, and waiting. Say what each of these states means in a few words and fill in the state labels in the sequence graph shown below.

   Solution:

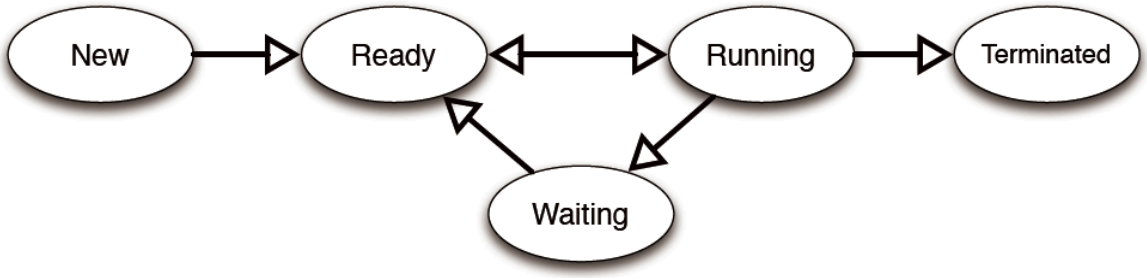

Figure 1: A state sequence graph depicting the lifecycle of a process.

new - *the process state is being created and initialized by the OS*
ready - *the process is ready to run and waiting to be executed by the CPU*
running - *the process is executing instructions on the CPU*
waiting - *the process is waiting for an event to complete (e.g., an I/O operation)*
terminated - *the process is finished and is being destroyed by the OS*

2. Several types of scheduling policies were discussed in class {first-come-first-served (FCFS), round robin (RR), shortest job first (SJF), multilevel feedback queues (MLFQ), and lottery scheduling (LS)}. Suppose you want to optimize your scheduler for certain types of workloads. For each type, state and briefly justify which type of scheduler you would use: **(i)** multiuser workloads in which no individual user should be favored, **(ii)** workloads with many mixed CPU and I/O jobs, and **(iii)** workloads with frequent I/O bound jobs and some very long-running, CPU-heavy jobs.

   Solution:

For (i), we are interested primarily in fairness, which makes round robin a good choice - starvation is impossible and every job will get equal time on the CPU. For (ii), we want to get I/O work off-loaded quickly so that the average waiting time isn't too bad, so SJF/MLFQ would be a good choice - since I/O bound jobs are prioritized, processes will avoid being delayed by I/O work. For (iii), we want to avoid starvation of the long-running CPU-intensive jobs, which makes lottery scheduling a good choice - this will trade off the waiting time benefits of SJF with the fairness and starvation avoidance of round robin.

3. Suppose there are 3 jobs: A, B, and C, of length 30, 20, and 10 seconds, and with start time at the 0th, 2nd, and 4th second respectively. The last 4 seconds of the 10-seconds job (job C) has 1 second of I/O every other second. There are 3 queues, and initial time slice is 1 second. Assume that the context-switch time is 0. Under the strategy of Multilevel Feedback Queues,

   a. Please indicate that the 1st time for job B entering "*Ready*" state is _2_ second, and the 1st time for job C entering "*Wait*" state is _21_ second.
   b. Sketch the scheduling of the jobs below. Remember the notation $Job_{time}^{workDone}$; for example B_6^2 means that job *B* has completed 2 seconds of work at time *t* = 6. You are required to show the progress to which the job C is accomplished.

Queue	Time Slice	Jobs
1	1	A_1^1 B_4^1 C_5^1 C_{31}^9 C_{36}^{10}
2	2	A_3^3 B_7^3 C_9^3 C_{26}^8
3	4	A_{13}^7 B_{17}^7 C_{21}^7 A_{25}^{11} B_{30}^{11} A_{35}^{15}