CSE 332 Midterm #1

Güneş ERÇAL, Olcay Taner YILDIZ

I. Question

Given the following C program, draw the process tree.

```
void main(void)
{
 int p0, p1, p2, p3, p4;
 int p5, p6;

 p0 = fork();
 if (p0 > 0) {
   p1 = fork();
   if (p1 > 0) {
      p2 = fork();
      if (p2 == 0) {
         p3 = fork();
         if (p3 == 0) {
            p4 = fork();
         } // endif p3 == 0
       } // endif p2 == 0
   } // endif p1 > 0
 } // end if p0 > 0
 else if (p0 == 0) {
   p5 = fork();
   if (p5 > 0) {
      p6 = fork();
    } // endif p5 > 0
 } // end elseif p0 == 0
}
```

II. Question

Consider the following set of processes, with the length of the CPU-burst time given in milliseconds: P1 bursttime 4, P2 bursttime 3, P3 bursttime 5, P4 bursttime 2, P5 bursttime 3. The processes are assumed to have arrived in the order P1, P2, P3, P4, P5, all at time 0.

a) Draw five Gantt charts illustrating the execution of these processes using FCFS, SJF, RR with quantum = 1, RR with quantum = 2, and RR with quantum = 3.
b) Compute the average turnaround time, response time, and waiting time for each of the algorithms of part (a).

III. Question

Consider the following Gantt Chart (with millisecond marks being given):

```
0   3    7    12  14  17  21  22 24
|---|----|-----|--|---|----|- |--|
  Px   Py  Pz   Pa  Pb   Pz  Pc Pz
```

Give the unique scheduling algorithm that is consistent with such a Gantt chart, along with consistent arrival times for each process such that the number of simultaneous (but in-order) arrivals is maximal.

IV. Question

We are given four processes, P1, P2, P3, and P4, with statements S1, S2, S3, and S4 respectively corresponding to each process (such that statement Si is in the code of process Pi) with the following constraints:

- There are no constraints on S1.
- S2 must execute after the execution of S1 is complete.
- S3 must execute after the execution of S1 is complete.
- S4 must execute after both of the statements S2 and S3 have completed execution.

Using semaphores, write each process in such a way that enforces the above precedence constraints.

V. True or False Questions

For each of the following questions, 1 through 10, write either “TRUE” if the statement is True, or “FALSE” if the statement is False.

1. Message-passing is the main bottleneck (limitation) of the micro-kernel architecture.
2. A process in the Running state can go immediately to the Terminated state upon an Interrupt.
3. A main failure of RR scheduling is in the long response time.
4. RR with minimal time quantum is the same as FCFS.
5. RR with maximal time quantum is the same as FCFS.
6. A pair of processes communicating over a network employ a pair of sockets, where each socket is defined by an IP address concatenated by a port number.
7. An advantage of using threads is that it is more efficient to create and context-switch threads than it is to create and context-switch processes.
8. It is a good idea to maximize the number of context switches to give the illusion of greater concurrency.
9. The fundamental problem of SJF scheduling is in determining the actual burst time.
10. Starvation is the situation that occurs when there exists a process Pi will not ever get out of waiting.