NAME: _____

CSCI 315 Operating Systems Design
Fall 2014 - Prof. Felipe Perrone
Activity 9

Source: Operating Systems Concepts}, Silberschatz, Galvin, and Gagne, 9th edition.
Consider the following set of processes, with the length of the CPU burst given in milliseconds:

Process	Burst Time	Priority
P_1	10	3
P_2	1	1
P_3	2	3
P_4	1	4
P_5	5	2

The processes are assumed to have arrived in the order {P_1, P_2, P_3, P_4, P_5} all at time 0. Consider the following scheduling algorithms: first-come, first-serve (FCFS), shortest-job first (SJF), nonpreemptive priority (a smaller number implies a higher priority), round-robin (RR) with *quantum*=1, and RR with *quantum*=2.

1) For each of the algorithms above:
- Draw a Gantt chart to illustrate the execution of these processes
- Calculate individual process turnaround times and the average turnaround time
- Calculate individual waiting times and average waiting time

2) Considering just the data that you were provided, which is *the best* scheduling algorithm? Justify your choice be explaining what criteria you used to identify what you called *the best*.

3) In a general sense, can any of these scheduling algorithms result in process *starvation*?

4) How realistic is it to assume that the scheduling algorithm will have the length of CPU bursts for the processes that it will need to schedule? How/where can one get information on the length of CPU bursts for a CPU scheduling algorithm to use?

5) How would your Gantt charts and metrics change if we assumed the following arrival times:
- P_1 arrives at time 0
- P_2 arrives at time 8
- P_3 arrives at time 8
- P_5 arrives at time 9
- P_4 arrives at time 10

6) How would your Gantt charts and metrics above change for preemptive priority scheduling?