Operating Systems

1. 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. Consider the following set of processes, with the length of the CPU-burst time 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.

   a. Draw four Gantt charts illustrating the execution of these processes using FCFS, SJF, a non-preemptive priority (a smaller priority number implies a higher priority), and RR (quantum = 1) scheduling.
   b. What is the turnaround time of each process for each of the scheduling algorithms in part a?
   c. What is the waiting time of each process for each of the scheduling algorithms in part a?
   d. Which of the schedules in part a results in the minimal average waiting time (over all processes)?

3. What advantage is there in having different time-quantum sizes on different levels of a multilevel queueing system?

4. Many CPU-scheduling algorithms are parameterized. For example, the RR algorithm requires a parameter to indicate the time slice. Multilevel feedback queues require parameters to define the number of queues, the scheduling algorithms for each queue, the criteria used to move processes between queues, and so on.

   These algorithms are thus really sets of algorithms (for example, the set of RR algorithms for all time slices, and so on). One set of algorithms may include another (for example, the FCFS algorithm is the RR algorithm with an infinite time quantum). What (if any) relation holds between the following pairs of algorithms?

   a. Priority and SJF
   b. Multilevel feedback queues and FCFS
   c. Priority and FCFS
   d. RR and SJF