

Question 1. (30 points) For the following mix of jobs, lengths, and arrival times, determine the scheduling order and average wait time for the First Come First Served (FCFS), Round Robin (RR), and non-preemptive Shortest Job First (SJF) schedulers. Assume all times are in msec. For the RR scheduler, use a 5 msec time slice and a 0 msec context switch cost; assume that when new tasks arrive they are placed at the head of the queue for jobs waiting to be scheduled. Fill in the table and use the reported average wait times to check your answers.

			Completion Time			Waiting Time		
Job	Arrival Time	Length	FCFS	RR	SJF	FCFS	RR	SJF
A	0	40						
B	5	15						
C	20	25						
D	40	20						
Average Wait Time:						27.5	36.25	26.25

Question 2. (10 points) State whether user or kernel level threads would be better suited to each of the following applications and explain why.

1. A file sharing application that performs network and disk I/O activities simultaneously.
2. A weather modeling system in which large numbers of threads perform computation and exchange messages with each other.

Question 3. (10 points) Recall from class how we addressed the problem of busy waiting when implementing locks by introducing the idea of a 'guard' variable (lecture 7, slide 19). Does this approach eliminate busy waiting? If it doesn't, explain why it preferable to the earlier, simpler approach where we implemented locks by calling test&set on the lock value itself.