This homework is due at the beginning of class on October 1, 2014 and is worth 1.5% of your grade.

Name: _____

CCIS Username: _____

Problem	Possible	Score
1	10	
2	10	
3	20	
4	35	
5	15	
Total	90	

1. What are the differences between user-level threads and kernel threads? Under what circumstances is one type better than the other? (10 pts)

.

2a. Describe the actions taken by the kernel to context-switch between user-level threads. (5 pts)

2b. Describe the actions taken by the kernel to context-switch between kernel-level threads. (5 pts)

3. Given the following mix of job, job lengths, and arrival times, assume a time slice of 15ms and compute the completion and average turnaround time of each job for the (non-preemptive) FCFS, RR, and SJF algorithms. Please use the following table for your solution.   (20 pts)

Job	Length (ms)	Arrival time (ms)	Completion time		
			FCFS	RR	SJF
0	45	0			
1	30	10			
2	15	11			
3	60	80			
4	45	85			
Average turnaround time					

4. Consider the following set of processes, which arrive in order at time 0, with the length of the CPU burst given in milliseconds.

Process	Burst time (ms)	Priority
P_1	10	2
P_2	1	1
P_3	2	3
P_4	1	4
P_5	5	5

4a. Draw four Gantt charts that illustrate the execution of these processes using the following scheduling algorithms: FCFS, SJF, non-preemptive priority (smaller numbers are higher priority), and RR (with quantum of 1).   (15 pts)

4b. What is the turnaround time of each process for each of the scheduling algorithms?   (10 pts)

Process	Turnaround time			
	FCFS	SJF	Priority	RR
P_1				
P_2				
P_3				
P_4				
P_5				

4c. What is the waiting time of each process for each of the scheduling algorithms?   (10 pts)

Process	Waiting time			
	FCFS	SJF	Priority	RR
P_1				
P_2				
P_3				
P_4				
P_5				

5a. Consider a system using multi-level feedback queue scheduling, and a user who has a large CPU-bound job to run. What strategy can the user employ in order to maximize the amount of CPU time allocated to his process? In this case, the user does not have any control over the queue, but does know the queue's policies (e.g., the top-level queue has an 8ms timeout, the middle-level queue has a 16ms timeout, etc). (10 pts)

5b. What is the effect of implementing the round-robin scheduler with arbitrarily large quantum values? (5 pts)