Full Name: ____________________

CS 450 Summer 2019

Midterm Exam

June 12, 2019

Instructions:

- This exam is closed-book, closed-notes.
- Calculators are not permitted. Endeavor to leave your computed answers in fraction form.
- For computed answers, you must show your work for credit.

Problem 1 (/10) :	
Problem 2 (/12) :	
Problem 3 (/8) :	
Problem 4 (/8) :	
Problem 5 (/9) :	
Problem 6 (/6) :	
Problem 7 (/6) :	
TOTAL (/59) :	

Problem 1. (10 points):

Multiple choice. For each of the following multiple choice problems, choose the *single best* answer by **circling**.

1. What statistic is most likely to be used to help identify interactive jobs for scheduling purposes (e.g., as in a MLFQ)?

   a. average CPU burst length

   b. total process turnaround time

   c. total number of I/O bursts

   d. total CPU execution time

2. Which of the following state transitions will never take place in a FCFS scheduler?

   a. ready $\rightarrow$ running

   b. running $\rightarrow$ ready

   c. running $\rightarrow$ blocked

   d. blocked $\rightarrow$ ready

3. Which of the following scheduling policies may exhibit starvation?

   a. FCFS

   b. Round-robin

   c. Non-preemptive SJF

   d. Highest penalty ratio next

4. Which of the following scheduling policies requires the use of a predictive mechanism such as the exponential moving average (EMA)?

   a. Round-robin

   b. Non-preemptive SJF

   c. Selfish round-robin

   d. Highest penalty ratio next

5. Which of the following is *not* necessarily true of a stable M/G/1 queueing system with average arrival and service rates λ and μ, and average number of waiting customers $E(L_q)$?

   a. the server utilization (ρ) is the ratio of λ and μ

   b. new arrivals will see an average of $E(L_q)$ waiting customers

   c. the average residual service time (T_r) equals the average service time ($\frac{1}{\mu}$)

   d. Little's law can predict average wait time ($E(T_q)$) directly from λ and $E(L_q)$

Problem 2. (12 points):

Briefly answer each of the following questions in the space provided.

1. List and briefly describe two advantages and two disadvantages a microkernel architecture is likely to have when compared to a monolithic kernel architecture.

2. How might an OS kernel leverage the x86 architecture to implement a secure handoff between user processes and privileged kernel code (as in a system call)?

3. Describe a scenario that demonstrates *priority inversion*, and explain how it might be averted.

Problem 3. (8 points):

Complete the diagram below by adding and labeling directed edges between the nodes, indicating possible process state transitions. While you need not label all the edges you draw, at a minimum you should indicate those transitions that correspond to:

a. Arrival

b. Preemption

c. I/O request / Syscall

d. I/O request completion

e. Scheduler selection

f. Completion

You may simply write the corresponding letter for each label given above next to the appropriate edge. Labels may be used more than once.

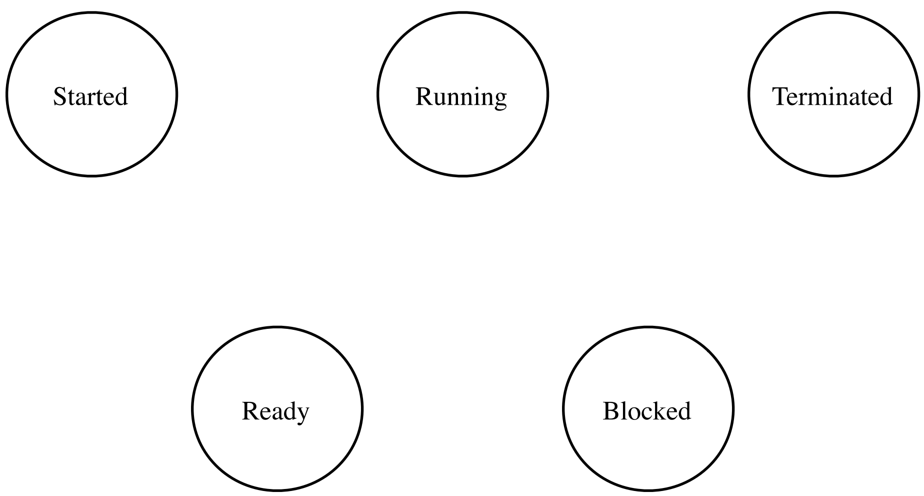

Problem 4. (8 points):

Consider a scheduling simulation involving 20 processes, each running for a total duration of 1000ms spread over separate CPU bursts ranging uniformly in duration from 25ms to 75ms.
The following four scheduling policies are used in the simulation:

- FCFS
- RR, q=50
- SJF (non-preemptive)
- PSJF

The following are the results of the simulation (policies not in order):

Policy	Total CST	Avg wait time	Std dev of wait time
#1	1212	17772	21.16
#2	816	12106	235.68
#3	816	17152	19.59
#4	980	12143	240.09

Which policies correspond to #1, #2, #3 and #4? Justify your answers.

Problem 5. (9 points):

The following are arrival times and CPU burst durations for five processes.

Process	Arrival Time	CPU Burst
P_0	0	8
P_1	3	4
P_2	8	2
P_3	12	5
P_4	13	3

Complete the following table with the individual process and average waiting times for each of the indicated scheduling policies. Ignore context switch overhead.

	Wait Times					
Scheduling policy	P_0	P_1	P_2	P_3	P_4	Avg wait time
First-Come First-Served						
Non-preemptive Shortest Job First						
Round-Robin with quantum=3						

Use the space below (or on the back of this page) for your work. You may wish to draw a Gantt chart for each scheduling algorithm — if your results above are incorrect, the charts may be evaluated for partial credit.

These equations may be useful for the following queueing theory problems. Please show your work!

Little's Law:	$E(L) = \lambda E(T)$	
M/M/1:	$E(T_q) = \dfrac{\rho}{\mu(1-\rho)}$	$E(T) = E(T_q) + E(T_s)$
M/G/1:	$E(T_q) = \dfrac{\rho E(T_r)}{1-\rho}$	$E(T_r) = \dfrac{C_{T_s}^2 + 1}{2} \cdot E(T_s)$

Problem 6. (6 points):

Consider a web server with an exponentially distributed average service time of 1ms. Assuming arrivals follow a Poisson process:

1. If we require that the average turnaround time of a request sent to the server be no more than 3ms, what is the maximum average arrival rate the server can tolerate?

2. When the average arrival rate is equal to your answer in part (1), how many requests on average are waiting to be serviced?

Problem 7. (6 points):

The average amount of time customers wait in line at the campus bookstore is 6 minutes. Monitoring the checkout clerk reveals that he is only busy $\frac{1}{3}$ of the time, and that service times are non-exponentially distributed with $C^2 = 2$. Assuming arrivals follow a Poisson process:

1. What is the average service time of the checkout clerk?

2. After a campus-wide marketing effort, the clerk finds that he is now busy $\frac{2}{3}$ of the time. What is the new average waiting time (assuming the service time distribution doesn't change)?