Midterm

CS 431
Winter 2006
Craig A. Rich

1 What is the most significant limitation of reading/writing devices using an I/O port? How does memory-mapped I/O overcome the limitation? How does direct memory access (DMA) overcome the limitation?

2 What causes a process control block (PCB) to be moved from a wait queue (other than the ready queue) to the ready queue? What code actually moves the PCB? What determines if the process "awoken" is the next process to run?

3 Recall that a child process forked from a parent process has the same initial memory contents/bounds, executable code, open files/devices and program counter as its parent. How can the child process behave differently than the parent process under those circumstances?

4 Why do two threads in the same process easily interfere with each other?

5 Consider the following queueing network with routing frequencies. Assume the following devices and visit service times:

queue number	device	visit service time (sec/visit)
1	A	0.06
2	B	0.05
3	C	0.10
4	D	0.02

a) Write the process flow balance equations.

b) Solve the process flow balance equations for the device throughputs, assuming the device D throughput X_4 is 3.4 visits/sec.

c) for each device, compute the device utilization, visit ratio, and process service time.

d) Which device is the bottleneck device?

e) What is the percent useful improvement in this device?

6 Consider processes arriving at a ready queue,

pid	arrival time (ms)	burst duration (ms)
1	0	15
2	5	15
3	10	10
4	20	5
5	35	10

dispatched to a single CPU by one of the following scheduling algorithms:

first-come first-served (FCFS)
nonpreemptive shortest job first (NSJF)
preemptive shortest job first (PSJF)
round robin with 10ms quantum (RR10)

For each CPU scheduling algorithm...

a) Give a Gantt chart showing the arrival time and running time slice(s) for each process.
b) Compute the waiting time for each process and the average waiting time.
c) Compute the process turnaround time for each process and the average turnaround time.

7 Why is average waiting time a better measure of CPU scheduling effectiveness than average turnaround time?