COSC 625 FINAL EXAM PRACTICE FALL 2010 Name: _____

Open book, notes, internet. Show work for partial credit.

1. There are 3 ISRs, each containing 5 instructions. Execution time for one instruction is one tick. Interrupts are disabled during the entire execution of an ISR.

All interrupts are at the same priority. Context switching and other kernel overhead take zero time.

Suppose at t = 1, interrupt 1 occurs
t = 2, interrupt 2 occurs
t = 3, interrupt 3 occurs
t = 4, interrupt 1 occurs
t = 5, interrupt 2 occurs
t = 6, interrupt 1 occurs

1. a. For each of these interrupts, how long before the corresponding ISR begins execution?

time	interrupt#	begin ISR execution
1		
2		
3		
4		
5		
6		

1.b. For each of these interrupts, how long before the corresponding ISR completes execution?

2. Use the line numbers given on the tinyexec source code at
http://www.emunix.emich.edu/~haynes/625/fa07/tinyexec/tinyexec.c
or in the hardcopy handed out at class on 11/29/2010.

"Tiny Exec hard codes the number of tasks and timers along with their execution addresses."
Suppose you have 5 tasks: task0, task1, task2, task3, and the background loop, and 4 timers: tmr0, tmr1, tmr2, tmr3.

The functions task0, task1, task2, task3, tmr0, tmr1, tmr2, tmr3 and main can be assumed to be correctly coded and correctly prototyped.

Give all the lines in tinyexec.c that you must change to accommodate the 5 tasks and 4 timers (hint: there are six lines you must change). Include the line numbers and give just the changed lines here.

3. Give the EDF schedule (by sketching the timeline of CPU assignment to t=10) for this job set. Jobs are pre-emptable.

#	**r** release time	**d** absolute deadline	**e** execution time
1	0	9	3
2	1	5	2
3	2	4	2
4	6	7	1

4. Consider the following set of periodic tasks, τ_i.

#	C_i	T_i
τ_1	1	3
τ_2	2	4

a. What is the hyperperiod (AKA major cycle)?

b. Give the RM schedule for one hyperperiod.

c. Over the length of the hyperperiod, what is the average response time of task τ_1?

d. Over the length of the hyperperiod, what is the average response time of task τ_2?

5. Here is a set of periodic tasks, τ_i

#	C_i	T_i
τ_1	1	3
τ_2	2	6

and a polled server:

PS	1	4

There is also a set of aperiodic tasks, Ai with the following computation requirements and release times:

#	C_i	r_i
A1	2	5
A2	1	6

a. Give the schedule using polled server scheduling for the aperiodic requests. The aperiodic task queue is a FIFO queue.

b. For the above answer to a (polled server scheduling of the aperiodic tasks), fill in the following table:

#	s_i	f_i	R_i (Response time)
A1			
A2			

c. What is the average response time of the aperiodic tasks?

d. What is the percent idle time over the hyperiod of the periodic task set (including polled server)?

6. Short answers:

a. When a running task is pre-empted, what is the task's state change (start -> finish)?

b. What is the meaning of re-entrant?

c. In order to make a function re-entrant, is it sufficient to do the following in order:

1. disable interrupts
2. save registers
3. restore interrupts

d. Aperiodic task performance is typically judged using latency (the smaller the latency, the better). Is this hard or soft real-time?

e. You use a semaphore to control access to a critical region of code. Is this a binary semaphore or a counting semaphore?

f. Round robin task scheduling is bad for hard real time. Give a reason why.