

Real-Time Systems / Real-Time Operating Systems
EE445M/EE380L.12, Spring 2019

Final Exam

Date: May 20, 2019

UT EID: _____

Printed Name: _____
Last, First

Your signature is your promise that you have not cheated and will not cheat on this exam, nor will you help others to cheat on this exam:

Signature: _____

Instructions:

- Open book and open notes.
- No calculators or any electronic devices (turn cell phones off).
- Please be sure that your answers to all questions (and all supporting work that is required) are contained in the space (boxes) provided.
- *Anything outside the boxes will be ignored in grading.*
- For all questions, unless otherwise stated, find the most efficient (time, resources) solution.

Problem 1	10	
Problem 2	15	
Problem 3	15	
Problem 4	15	
Problem 5	25	
Problem 6	20	
Total	100	

Name: _____

Problem 1 (10 points): Synchronization

Consider an OS that manages three resources. Threads can require two out of three resources, but never all three. Given below is the code for an agent thread running as part of the operating system handing out resources to different user threads (where multiple instances of each type can be running) according to an OS-internal policy. User and agent threads synchronize using semaphores $s1$, $s2$ and $s3$ associated with each resource as well as a *done* semaphore indicating that user threads are finished using resources. All semaphores are initialized to zero.

```

OS_Agent() {
  while(1) {
    switch(policy()) {
      case 1:
        OS_bSignal(&s1);
        OS_bSignal(&s2);
        break;
      case 2:
        OS_bSignal(&s2);
        OS_bSignal(&s3);
        break;
      case 3:
        OS_bSignal(&s1);
        OS_bSignal(&s3);
        break;
    }
    OS_bWait(&done);
  }
}

```

<pre> Thread1() { while(1) { OS_bWait(&s1); OS_bWait(&s2); ... OS_bSignal(&done); } } </pre>	<pre> Thread2() { while(1) { OS_bWait(&s2); OS_bWait(&s3); ... OS_bSignal(&done); } } </pre>	<pre> Thread3() { while(1) { OS_bWait(&s1); OS_bWait(&s3); ... OS_bSignal(&done); } } </pre>
--	--	--

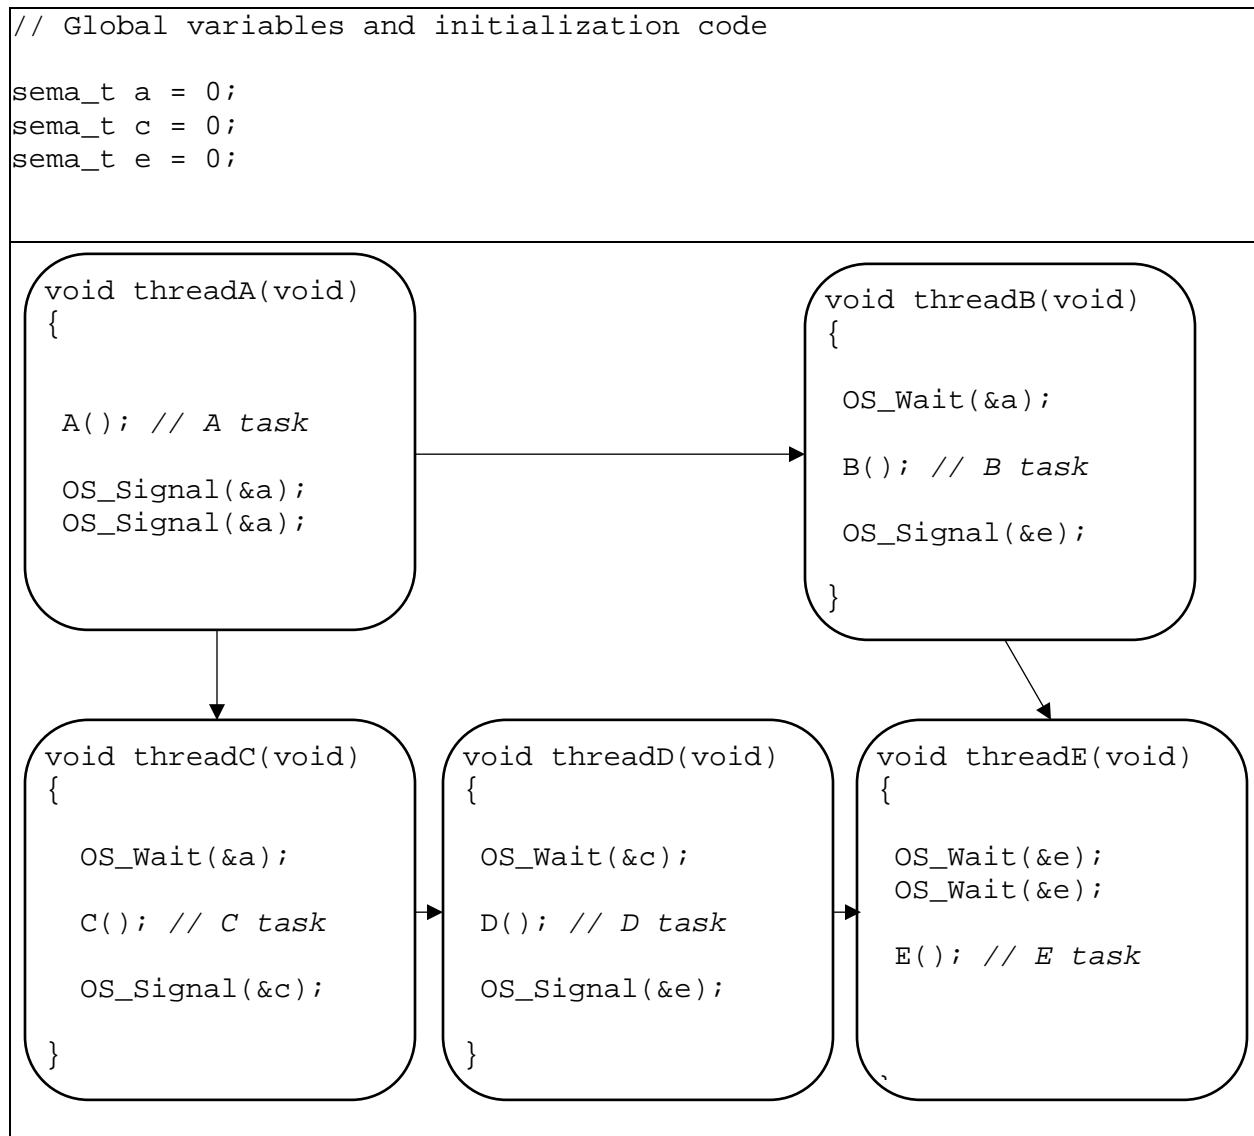
a) What is wrong with this code?

b) Can the code be modified to fix the problem while keeping the OS general, i.e. without hardcoding the number of instances running of each user thread type and without adding new threads? If so, modify the code accordingly. If not, why not?

Name: _____

Problem 2 (15 points): Scheduling

Given the task graph from Problem 3 in the midterm:




Now assume that we run this task graph on your perfectly functioning OS from Lab 3 using a priority scheduler. Further assume that the tasks have the following execution times and deadlines (relative to time 0).

Task	Execution Time	Deadline	Priority
A	1ms	5ms	
B	1ms	6ms	
C	1ms	7ms	
D	2ms	4ms	
E	2ms	8ms	

Name: _____

- a) Assign priorities to threads to realize a strict earliest deadline first (EDF) scheduling. Assuming all threads are launched at time 0, show the schedule of task executions.



- b) Are any of the deadlines violated? If so, is there an alternative priority assignment of priorities that would result in a valid schedule without deadline violations (show the priorities and change in the schedule)? If not, will EDF always be able to find a valid scheduler for a task graph with dependencies, if one exists (show why)?

- c) Is there any priority inversion in your schedule from a)? If so, would a priority inheritance or priority ceiling protocol fix the inversion including any deadline violations (if any)?

Name: _____

Problem 3 (15 points): Heap Management

Consider a heap of size 2048 bytes. Assuming an initially empty heap, given the following sequence of heap allocations and de-allocations:

```




p1 = Heap_Malloc(1024);
p2 = Heap_Malloc(512);
Heap_Free(p1);
p3 = Heap_Malloc(256);
p4 = Heap_Malloc(512);
p5 = Heap_Malloc(256);
p6 = Heap_Malloc(512);
Heap_Free(p4);
Heap_Free(p5);
p7 = Heap_Malloc(768);
Heap_Free(p2);
Heap_Free(p3);
    
```

Show the state of the heap at the end of this sequence using a first fit, best fit and worst fit allocation scheme. Assume that the heap manager does not require any overhead for extra meta-data, and that the heap is allocated from bottom to top, i.e. a block always ends up being placed at the bottom (lowest address) of its chosen free space. If an allocation fails, assume that the function returns a null pointer and the rest of the sequence continues. For each allocation scheme, answer the following questions: (1) Does the sequence complete without error? If not, show the allocations that fail. (2) What is the amount of free space and what is the largest block size that can be allocated after this sequence?

a) First fit

b) Best fit

c) Worst fit



--	--	--

Name: _____

Problem 4 (15 points): Process Management

Recall that in Lab 5, you were given an ELF loader to load a process from the SD card. The ELF loader's *exec_elf()* routine parses the ELF file, loads its code and data segments into dynamically allocated memory regions on the heap and finally calls *OS_AddProcess()* as shown below:

```
int exec_elf(const char *path, const ELFEnv_t *env)
{
    ...
    f_open(&f, path, FA_READ);          // open & read ELF header
    ...

    text = Heap_Malloc(code_size);     // allocate & load code segment
    f_read(f, text, code_size);

    data = Heap_Malloc(data_size);     // allocate & load data segment
    f_read(f, data, data_size);

    ...                                // relocation using 'env'

    f_close(f);
    return OS_AddProcess(entry, text, data); // add OS process
}
```

- a) Given basic (but incomplete) implementations of *OS_AddProcess()* and *OS_Kill()* below, complete the code to properly manage process memory and reclaim it when the process exits.

```
struct PCB {
    unsigned long Id;
    unsigned long numThreads;

};

struct PCB pcbs[NUMPCB];

struct TCB {
    struct TCB* next;
    unsigned int* sp;
    unsigned long Id;
    struct PCB *process;

};

struct TCB tcbs[NUMTCB];
struct TCB *RunPt;
```

Name: _____

```

int OS_AddProcess(void(*entry)(void), void *text, void *data){
    int sr, n;
    EnterCritical()
    for(n = 0; n < NUMPCB; n++) {
        if(pcbs[n].Id == 0) break; // found it; Id = 0 implies free PCB
    }
    if(n == NUMPCB) { ExitCritical(); return 0; }

    pcbs[n].Id = ++processId;
    pcbs[n].numThreads = 1;

    ExitCritical()
    return OS_AddThread(entry, STACK_SIZE, &pcb[n]);
} // sets TCB.process = &pcb[n]

void OS_Kill(void) {
    struct TCB *p;
    DisableInterrupts();

    if(RunPt->process) {
        RunPt->process->numThreads--;
    }

    // Remove thread and return to TCB pool
    p = RunPt;
    while((p->Next) != RunPt) p = p->Next;
    p->Next = RunPt->next;
    RunPt->Id = 0; // mark as free

    ContextSwitch();
    EnableInterrupts();
}

```

Name: _____

- b) Assuming a program on disk with a code segment of size 128 bytes and a data segment of size 384 bytes. Given a heap size of 4096 bytes and assuming that the heap is not used for anything else, what is the maximum number of concurrent instances of this program that you can load and run?

- c) Is it possible to increase the number of concurrent instances of the same program that can be loaded and running? If not, why not? If yes, describe how this can be done and sketch the changes that would need to be made to the ELF loader or OS code shown above. What is the maximum number of concurrent instances in this case?

Name: _____

Problem 5 (25 points): File System

Assume two separate disks of identical size 4kB with 16 blocks each, where one disk uses a filesystem with indexed and the other a filesystem with linked allocation. Further assume that the exact same two files ‘A’ and ‘G’ (but nothing else) have been copied onto both disks.

- a) Given the partial disk states (blocks, directory and file system tables) as shown below, complete all the missing information on each disk to represent correct filesystems. Mark each block with the name of the file it belongs to (or ‘*’ if empty space). Assume that files ‘A’ and ‘G’ fit into the same number of blocks on both disks.

<u>Indexed</u>			<u>Linked</u>				
Directory			Index Table	Disk			
File	Start	Blocks			File	Start	Blocks
A	0		0	7	0	Dir	
G		7	1	4	1	Idx	
*			2	9	2		
			3	11	3		
			4	12	4	A	
			5	5	5		
			6	14	6		
			7	13	7		
			8		8		
			9	3	9		
			10	2	10		
			11	10	11		
			12	8	12		
			13	6	13		
			14		14		
			15		15		

Directory			Disk		
File	Start	Blocks	Dir		
A	4		0	-	
G			1	5	
*			2	8	
			3	11	
			4	6	
			5	12	
			6	9	
			7	0	
			8	G	0
			9	0	
			10		
			11	2	
			12	13	
			13	7	
			14	15	
			15	10	

* Free space

- b) What is the space overhead of the two filesystems, i.e. how many out of the 4096 bytes on each of the two disks are usable for file data storage?

Indexed	Linked

Name: _____

- c) Given the sequence of file system calls below, show the order of disk accesses that are made by the each of the two filesystems. Assume that nothing is initially in memory. List the disk block numbers accessed (read or written). How many disk accesses are made in each case?

```
f_open(&fa, "A", FA_READ);    // open 'A'
f_seek(fa, 306);             // read integer from byte position 306
f_read(fa, &x, 4);
f_seek(fa, 545);             // read integer from byte position 545
f_read(fa, &y, 4);
f_seek(fa, 319);             // read integer from byte position 319
f_read(fa, &z, 4);
f_close(fa);
gpa = (x + y + z) / 3;
f_open(&fg, "G", FA_WRITE);
f_seek(fg, 21);              // write result to byte position 21
f_write(fg, &gpa, 4);
f_close(fg);
```

Indexed	Linked

- d) Can the file system performance be improved, i.e. can the number of disk accesses be reduced in either case? If so, how, and how many disk accesses are made in that case?

Name: _____

Problem 6 (20 points): Remote Procedure Call

You are asked to implement a remote procedure call (RPC) between two micro-controllers over the CAN bus. With RPCs, the OS provides the capability for a thread on one machine to call a function implemented on another machine while making it look to the programmer as if the thread is just calling a local function, where the OSs on both machines transparently handle all the necessary communication and synchronization to execute the call remotely. Given the high-level CAN driver (`can0.h`) from the lab with driver primitives as listed below, complete the code such that on the RPC client machine can call function on the RPC server. Show any changes to the driver code that you need to make.

Client:

```

// *** can0.h ***
#define CAN_BITRATE 1000000
#define RCV_ID 2
#define XMT_ID 4

int CAN0_CheckMail(void);
int CAN0_GetMailNonBlock(uint8_t data[4]);
void CAN0_GetMail(uint8_t data[4]);
void CAN0_Open(void);
void CAN0_SendData(uint8_t data[4]);

// local function stub to call remote function
uint8_t PWM_Duty(uint16_t left, unit16_t right) {

}

void speedThread(void) {
    ...
    if(!PWM_Duty(l, r)) goto error;
    ...
}

void main(void) {
    PLL_Init();
    OS_Init();

    OS_AddThread(&speedThread, 128);
    ...

    OS_Launch(TIME_1MS);
}

```

Name: _____

Server:

```
// *** can0.h ***
#define CAN_BITRATE 1000000
#define RCV_ID 4
#define XMT_ID 2

int CAN0_CheckMail(void);
int CAN0_GetMailNonBlock(uint8_t data[4]);
void CAN0_GetMail(uint8_t data[4]);
void CAN0_Open(void);
void CAN0_SendData(uint8_t data[4]);

// local function
uint8_t PWM_Duty(uint16_t left, uint16_t right) {
    PWM_0_CMPA_R = left - 1;
    PWM_1_CMPA_R = right - 1;
    return PB0;
}

void main(void)
{
    PLL_Init();
    OS_Init();
    PWM_Init();

    ...

    OS_Launch(TIME_1MS);
}
```