

REAL-TIME SYSTEMS

Solutions to final exam March 6, 2012

PROBLEM 1

- a) FALSE: The utilization guarantee bound for RM-US converges towards 33.3% as the number of processors become very large.
 - b) TRUE: Too large over-estimation of the execution time of a task will leave less slack for the execution of other tasks in the schedulability analysis.
 - c) FALSE: Sporadic tasks can be used in a real-time system with hard timing constraints since such tasks have a guaranteed minimum interarrival time, thus giving an upper bound of these tasks' utilization of the processor.
 - d) FALSE: The critical instant refers to a point in time when the response time of an analyzed task is maximized. In single-processor system the critical instant occurs when the task arrives at the same time as all tasks with higher priority.
 - e) FALSE: By following certain guidelines of how to request resources it is possible to avoid deadlock.
 - f) FALSE: By deadline inversion we mean a situation where a task with high priority (deadline close in time) is blocked from executing because another task with lower priority (deadline later in time) is holding an exclusive resource needed by the high-priority task.
-

PROBLEM 2

- a) Waiting for a corresponding time slot (TTP/C), Waiting for a transmission token (FDDI), Waiting for a contention-free transmission (Ethernet), Waiting for network priority negotiation (CAN).
- b) The message transmission delay consists of:
 - Frame delay:** $t_{frame} = N_{frame}/R$, where N_{frame} is the message length (in bits) and R is the data rate (in bits/s).
 - Propagation delay:** $t_{prop} = L/v$, where L is the communication distance (in m) and v is the signal propagation velocity (in m/s).
- c) Message queuing delay for Ethernet can in general not be bounded because a re-transmission will happen every time a collision on the communication medium occurs.
- d) The binary countdown protocol:
 1. Each node with a pending message waits until bus is idle.
 2. The node begins transmitting the highest-priority message pending on the node. Identifier is transmitted first, in the order of most-significant bit to least-significant bit.
 3. If a node transmits a recessive bit (1) but sees a dominant bit (0) on the bus, then it stops transmitting since it is not transmitting the highest-priority message in the system.
 4. The node that transmits the last bit of its identifier without detecting a bus inconsistency has the highest priority and can start transmitting the body of the message.

PROBLEM 3

- a) The WCET of main is dependent on the WCET of functions “methA”, “methB” and “multiply”. So, we first calculate the WCET of functions “methA”, “methB” and “multiply”.

WCET of “multiply”

$$\text{WCET}(\text{multiply}(a,b)) = \{\text{multiplication}\} + \{\text{return}\} = 5 + 2 = 7$$

WCET of “methA”

There are two cases for calculating the WCET of methA: case (i) $b == 0$ and case(ii) $b > 0$

Case(i)

$$\begin{aligned} \text{WCET}(\text{methA}(a,b==0)) &= \{\text{dec, p}\} + \{\text{dec, i}\} + \{\text{assign, p}\} + \{\text{assign, i}\} + \\ &\{\text{compare, } b==0\} + \{\text{return}\} = 1+1+1+1+2+2=8 \end{aligned}$$

Case(ii)

The WCET of methA for this case depends on the number of times the while loop executes. The WCET of the while loop for parameter b is given as “WCET(whileLoop, b)”. Therefore, we have

$$\begin{aligned} \text{WCET}(\text{methA}(a,b>0)) &= \{\text{dec, p}\} + \{\text{dec, i}\} + \{\text{assign, p}\} + \{\text{assign, i}\} + \\ &\{\text{compare, } b==0\} + \text{WCET}(\text{whileLoop, b}) + \{\text{return}\} \\ &= 1+1+1+1+2+(16b-14)+2=16b-6 \end{aligned}$$

where **WCET(whileLoop, b) = (16b-14)**. The value of WCET(whileLoop, b) is calculated as follows. The body of the while loop executes $(b - 1)$ times and the comparison statement of the while is checked b times.

$$\begin{aligned} \text{WCET}(\text{whileLoop, b}) &= b * \{\text{comparison}\} + (b-1)* (\{\text{call, multiply}(p,a)\} \\ &+ \{\text{WCET}(\text{multiply}(p,a))\} + \{\text{assign, p}\} + \{\text{add, i+1}\} + \{\text{assign, i}\}) \\ &= b * 2 + (b-1)*(2+7+1+3+1) = 16b-14 \end{aligned}$$

WCET of “methB”

There are two cases for calculating the WCET of methB: case (i) $b == 1$ and case(ii) $b > 1$

Case(i)

$$\text{WCET}(\text{methB}(a,b==1)) = \{\text{compare, } b==1\} + \{\text{return}\} = 2 + 2 = 4$$

Case(ii)

$$\begin{aligned} \text{WCET}(\text{methB}(a,b>1)) &= \{\text{compare, } b==1\} + \{\text{sub, } b-1\} + \{\text{call, methB}(a,b-1)\} \\ &+ \text{WCET}(\text{methB}(a,b-1)) + \{\text{call, multiply}(a,\text{methB}(a,b-1))\} \\ &+ \text{WCET}(\text{multiply}(a,\text{methB}(a,b-1))) + \{\text{return}\} \\ &= 2+3+2 + \text{WCET}(\text{methB}(a,b-1)) + 2+7+2 = 18 + \text{WCET}(\text{methB}(a,b-1)) \end{aligned}$$

WCET of “main”

Now we can calculate the WCET of main

$$\begin{aligned} \text{WCET}(\text{main}) &= \{\text{dec, ans}\} + \{\text{dec, x}\} + \{\text{dec, y}\} + \{\text{assign,x}\} + \{\text{assign,y}\} \\ &+ \{\text{call, methA}(2,3)\} + \text{WCET}(\text{methA}(2,3)) + \{\text{call, methB}(2,3)\} \\ &+ \text{WCET}(\text{methB}(2,3)) + \{\text{compare, methA}(2,3) > \text{methB}(2,3)\} + \{\text{assign, ans}\} + \{\text{return}\} \\ &= 1+1+1+1+1+2 + \text{WCET}(\text{methA}(2,3)) + 2 + \text{WCET}(\text{methB}(2,3)) + 2+1+2 \\ &= 14 + \text{WCET}(\text{methA}(2,3)) + \text{WCET}(\text{methB}(2,3)) \\ &= 14+42+40=96 \text{ (deadline is missed!!)} \end{aligned}$$

where

$$\text{WCET}(\text{methA}(2,3)) = 16*3-6=42$$

and,

$$\text{WCET}(\text{methB}(2,3)) = 18 + \text{WCET}(\text{methB}(2,2)) = 18 + 18 + \text{WCET}(\text{methB}(2,1)) = 18 + 18 + 4 = 40$$

b) The false paths are:

In methA() the statement “return 1” is never executed since the value of b is not 0 since the value for formal parameter b in methA(a,b) is 3.

The condition “if (methA(x,y) > methB(x,y))” in the main() function is never true since both methA and methB compute *x raised to the power y*. Thus, the two statements “ans='T;” and “x=x+y;” are never executed.

PROBLEM 4

a), b) The final code should look similar to this:

```
#include TinyTimber.h

typedef struct {
    Object super;
    char *id;
} PeriodicTask;

Object app = initObject();
PeriodicTask ptask1 = { initObject(), "Task 1" };
PeriodicTask ptask2 = { initObject(), "Task 2" };
```

```

Time max_wcet = 0;

void T1(PeriodicTask *self, int u) {
    Time start;
    Time diff;

    start = CURRENT_OFFSET();

    Action1(); // procedure doing time-critical work

    diff = CURRENT_OFFSET() - start;
    if (diff > max_wcet)
        max_wcet = diff;

    SEND(MSEC(85), MSEC(30), self, T1, 0);
}

void T2(PeriodicTask *self, int u) {
    Action2(); // procedure doing time-critical work

    SEND(MSEC(105), MSEC(55), self, T2, 0);
}

void kickoff(PeriodicTask *self, int u) {
    SEND(MSEC(20), MSEC(30), &ptask1, T1, 0);
    SEND(MSEC(0), MSEC(55), &ptask2, T2, 0);
}

main() {
    return TINYTIMBER(&app, kickoff, 0);
}

```

- c) The priorities for the periodic activities are given by the deadlines in the SEND() calls, since the TinyTimber kernel uses earliest-deadline-first scheduling.

PROBLEM 5

- a) **Priority Inheritance Protocol:** When a task τ_i blocks one or more higher-priority tasks, it temporarily assumes (inherits) the highest priority of the blocked tasks.
- a) **Priority Ceiling Protocol:** Each resource is assigned a priority ceiling equal to the priority of the highest-priority task that can lock it. Then, a task τ_i is allowed to enter a critical section only if its priority is higher than all priority ceilings of the resources currently locked by tasks other than τ_i . When the task τ_i blocks one or more higher-priority tasks, it temporarily inherits the highest priority of the blocked tasks.
- Advantages of PCP over PIP: (i) No deadlock: priority ceilings prevent deadlocks, (ii) No chained blocking: with PCP a task can be blocked at most the duration of one critical section (PIP can have chained blocking).
- c) Define the priorities as follows: H = highest priority, M_H = second highest priority, M_L = third highest priority and L = lowest priority. Since rate-monotonic scheduling is used, the task priorities are as follows: $prio(\tau_1) = H$, $prio(\tau_2) = L$, $prio(\tau_3) = M_L$, and $prio(\tau_4) = M_H$.

We first determine the ceiling priority for each semaphore:

$$\text{ceil}\{R_a\} = \max\{H, L\} = H \quad (\text{since } \tau_1 \text{ och } \tau_2 \text{ may lock the semaphore})$$

$$\text{ceil}\{R_b\} = \max\{M_H, M_L\} = M_H \quad (\text{since } \tau_3 \text{ och } \tau_4 \text{ may lock the semaphore})$$

$$\text{ceil}\{R_c\} = \max\{H, M_H\} = H \quad (\text{since } \tau_1 \text{ och } \tau_4 \text{ may lock the semaphore})$$

We then identify, for each task τ_i , what tasks with lower priority may block τ_i and thereby cause the corresponding blocking factor B_i :

$$B_1 = \max\{H_{2,a}, H_{4,c}\} = \max\{2, 3\} = 3 \quad (\text{since } \tau_1 \text{ may be blocked by } \tau_2 \text{ and } \tau_4 \text{ who lock semaphores whose ceiling priorities are higher than or equal to the priority of } \tau_1)$$

$$B_2 = 0 \quad (\text{since } \tau_2 \text{ has lowest priority of all tasks, and thereby per definition cannot be subject to blocking})$$

$$B_3 = \max\{H_{2,a}\} = 2 \quad (\text{since } \tau_3 \text{ may be blocked by } \tau_2 \text{ who locks a semaphore whose ceiling priority is higher than or equal to the priority of } \tau_3)$$


$$B_4 = \max\{H_{2,a}, H_{3,b}\} = \max\{2, 1\} = 2 \quad (\text{since } \tau_4 \text{ may be blocked by } \tau_2 \text{ and } \tau_3 \text{ who lock semaphores whose ceiling priorities are higher than or equal to the priority of } \tau_2)$$

PROBLEM 6

a) Time-table scheduling works as follows:

- Create a circular queue that corresponds to the time table: Each element in the queue contains start and finish times for a certain task (or task segment in case of preemptive scheduling). The elements in the queue are sorted by the start time
- Use clock interrupts: When a task starts executing, a real-time clock is programmed to generate an interrupt at the tasks expected finish time. When the interrupt occurs, the next task (the one whose start time is closest in time) in the circular queue is fetched and the system waits until that tasks given start time is due.

b) The length of the schedule is: $LCM(T_1, T_2, T_3) = LCM(4, 8, 16) = 16$. By simulating earliest-deadline-first scheduling we get the following schedule:




c) The generated schedule is the best possible (optimal). In a schedule generated by, for example, the deadline monotonic policy task τ_3 will miss its deadline.

PROBLEM 7

Since deadline-monotonic scheduling is used, the task priorities are as follows: $prio(\tau_1) = H$, $prio(\tau_2) = M$, $prio(\tau_3) = L$.

We generate a multiprocessor schedule with tasks τ_1 and τ_2 (having the highest priorities) running on one processor each. Task τ_3 is scheduled in the remaining time slots according to the following diagram (covering $t = 0$ to $t = 47$):



We observe that there is a critical instant at $t = 30$ in the diagram, maximizing the response time of τ_3 ($R_3 = 8$). Since $R_3 > D_3$ the schedule is not feasible.

Note that $t = 10$ is not a critical instant (which it would have been for a single-processor system)!