EINDHOVEN UNIVERSITY OF TECHNOLOGY

Department of Mathematics and Computer Science

Examination Real-time Architectures (2IN20)
on Wednesday, June 27^{th} 2007, 9.00h-12.00h.

First read the entire examination. There are 5 exercises in total. Grades are included between parentheses at all parts and sum up to 11 points. Good luck!

1. A piece of technical equipment has to remain in an exact temperature range given by an interval $[L, H]$. To that end, a real-time system controls a heater and a cooler via a simple on/off mechanism.

   (a) (0.5) Draw a system sketch
   - indicate controlling and controlled system;
   - explain the functionality of actuators and sensors you use.

   (b) (0.5) Specify and design the controlling system, assuming a maximum rate of temperature change of r degrees per second.

   (c) (0.5) Give a schedulability condition using r.

   Answer: The example has been taken from slide 22 of the the lecture 'RTA.B3-Specification concepts'. The solutions for (b) and (c) are similar to those of the water vessel problem; see lecture 'RTA.D0-Water-Vessel'.

2. Consider a set $\mathcal{T}$ of n tasks $\tau_1, \tau_2, \ldots \tau_n$ and the the following two equations

$$U \leq n(2^{1/n} - 1) \quad (1)$$

$$x = C_i + \sum_{j<i} \left\lceil \frac{x}{T_j} \right\rceil C_j \quad (2)$$

   (a) (0.5) Give four conditions that *must* hold for *both* equations to make them applicable.
   Answer Tasks are (strictly) periodic (or sporadic), pre-emptive, and independent, can immediately start upon release, and do not suspend themselves. Tasks are scheduled by means of fixed-priority pre-emptive scheduling.

   (b) (0.5) Give two conditions that must hold for Equation (1), but need not hold for Equation (2).
   Answer Priorities are assigned inversely proportional to periods, e.g. the task with the shortest period gets the highest priority. Deadlines are equal to periods.

   (c) (0.5) Give a condition that must hold for Equation (2), but need not hold for Equation (1).
   Answer Tasks are given in order of decreasing priority, i.e. task τ_1 has highest and task τ_n has lowest priority.

3. Consider the set $\mathcal{T}$ of three tasks with characteristics as given in Table 1, where τ_1 has highest priority and τ_3 has lowest priority.

   (a) (1.0) Determine the worst-case and best-case response times of the tasks under Fixed-Priority Pre-Emptive Scheduling (FPPS) and arbitrary phasing.
   Answer The solution can be determined by either drawing timelines (see Figure 1) or applying the analytic approach using equations. The solutions have been added to Table 1.

	$T = D$	C	WR	BR
τ_1	3	1	1	1
τ_2	5	2	3	2
τ_3	11	2	9	2

Table 1: Characteristics of the tasks of $\mathcal{T}$.

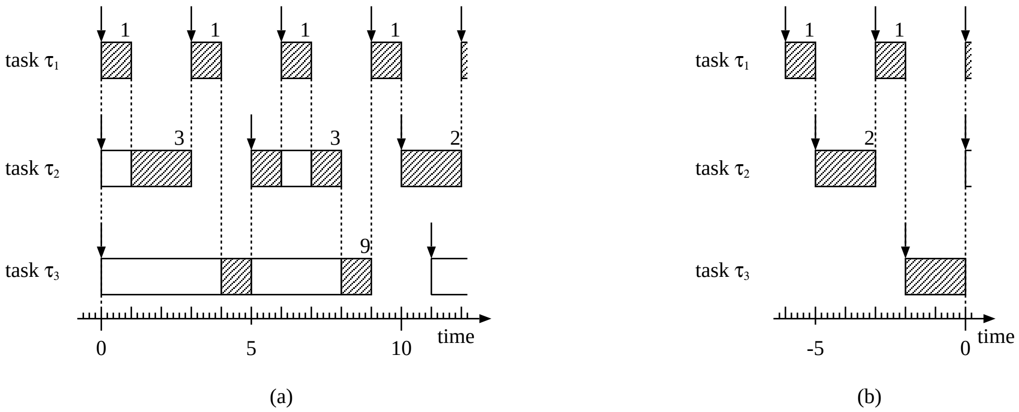

Figure 1: Timelines of $\mathcal{T}$ with (a) a critical instant for all three tasks at time zero and (b) an optimal instant for task τ_3 at time zero.

(b) (0.5) Determine the maximum activation jitter AJ_2 of task τ_2 which still guarantees schedulability of all tasks.

Answer Activation jitter of task τ_2 can only influence the response times of the jobs of tasks τ_2 and τ_3. Let's consider the influence of activation jitter of task τ_2 on task τ_3. Considering Figure 1, we can move the activations of the second and subsequent jobs of task τ_2 at most 1 back in time without causing the first job of task τ_3 to miss its deadline; see Figure 2. Moving the activations further back in time will cause the first job of task τ_3 to miss its deadline. The same result can be found by using the recursive equation with activation jitter to determine the maximum AJ_2 such that $WR_3 \leq D_3$.

An activation jitter $AJ_2 = 1$ has no influence on the worst-case response time of task τ_2 (see Figure 2). As a result, the activation jitter AJ_2 can be at most equal to 1.

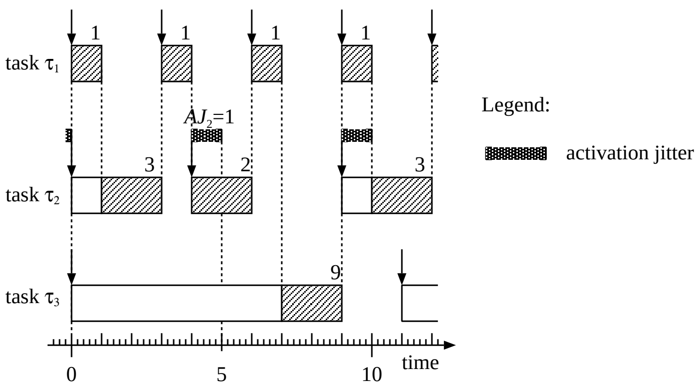

Figure 2: Timeline of $\mathcal{T}$ with activation jitter $AJ_2 = 1$ and a critical instant for task τ_3 at time zero.

(c) (0.5) Determine the maximum activation jitter AJ_3 of task τ_3 which still guarantees schedulability of all tasks (where $AJ_2 = 0$ again).

Answer Activation jitter of task τ_3 can only influence the response times of the jobs of tasks τ_3. In order to use the analysis for task τ_3, we have to make sure that $WR_3 \le T_3 - AJ_3$. Hence, $AJ_3 \le T_3 - WR_3 = 11 - 9 = 2$; see Figure 3(a). Using a

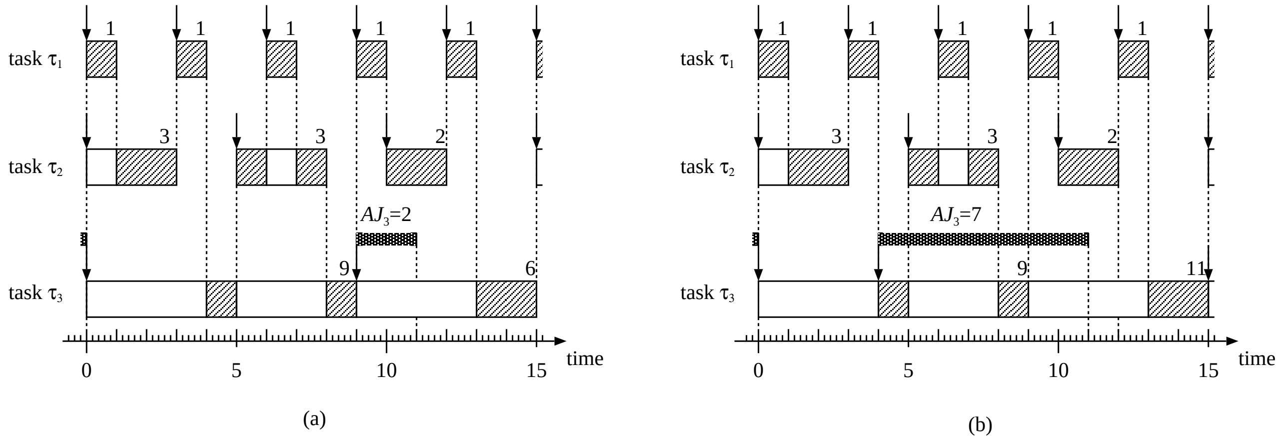

Figure 3: Timeline of $\mathcal{T}$ with a critical instant for task τ_3 at time zero and activation jitter (a) $AJ_3 = 2$ and (b) $AJ_3 = 7$.

timeline, we derive that AJ_3 can even be increased to $AJ_3 = 7$ without causing the second (or a subsequent) job of τ_3 to miss its deadline; see Figure 3(b).

(d) (1.5) Now assume task τ_3 to be a *server* σ rather than a task (with $AJ_2 = AJ_3 = 0$). Moreover, assume that σ is meant to serve an a-periodic task τ_a with a computation time $C_a = 2$, which arrives only once. Determine an upper bound for the response time of the arrival of τ_a when σ is

- a polling server;
  Answer Note that the computation time of τ_a is equal to the capacity C^σ of the the polling server σ^{PS}. Consider two consecutive periods of σ^{PS}. Assume that σ^{PS} is allowed to execute immediately upon its activation in the first period and experiences a worst-case situation in the second period. Moreover, assume that τ_a arrives just after the activation of σ^{PS} in the first period. Together, these conditions constitute a worst-case situation for τ_a. The response time R_a of τ_a is now bounded by $R_a < T^\sigma + WR^\sigma = 11 + 9 = 20$.
- a periodic server;
  Answer Consider two consecutive periods of the periodic server σ^{PS}. Assume that σ^{PS} experiences a best-case situation the in first period and a worst-case situation in the second period. Next, assume that τ_a arrives just when the capacity of σ^{PS} is depleted in the first period. The response time R_a of τ_a is now bounded by $R_a \le T^\sigma - C^\sigma + WR^\sigma = 11 - 2 + 9 = 18$.
- a sporadic server.
  Answer A sporadic server σ^{SS} 'keeps' its capacity when it is not used. Hence, the response time R_a of task τ_a is bounded by the worst-case response time WR^σ of σ^{SS}, i.e. $R_a \le WR^\sigma = 9$.

4. The *system ceiling* of a job is an essential notion for the Priority Ceiling Protocol (PCP).

(a) (0.5) Describe the notion system ceiling in your own words.
Answer See lecture 'RTA.B4-Policies-3'.

(b) (0.5) Give an example illustrating that the system ceiling of a job changes during the execution of that job, and draw the system ceiling of that job as a function of time.
Answer Fig. 7.12 is an example showing that the system ceiling of J_0 changes during

its execution; J_0 is blocked at t_5 (because the system ceiling is equal to p_1, being the priority of J_0 itself) and allowed to continue at t_6 (because the system ceiling is now reduced to p_2).

(c) (0.5) Give two advantages of the PCP when compared to the Priority Inheritance Protocol (PIP).
Answer See lecture 'RTA.B4.Policies-3'.

(d) (0.5) Give an advantage of PIP when compared to PCP.
Answer See lecture 'RTA.B4.Policies-3'.

5. One of the lectures concerned *Behavioural Analysis of Real-Time Systems with Interdependent Tasks.*

   (a) (1.0) Explain which real-time problems were addressed.

   (b) (1.0) Explain how these problems have been solved.

   (c) (0.5) Explain why these approaches were taken.

   Answers See slides of that lecture.
   Note that these questions were described as goals on sheet 6 of the lecture 'RTA.A1-Overview'.