EINDHOVEN UNIVERSITY OF TECHNOLOGY

Department of Mathematics and Computer Science

Examination Real-time Architectures (2IN60)
on Wednesday, November 4^{th} *2009, 9.00h-12.00h.*

First read the entire examination. There are 7 exercises in total. Grades are included between parentheses at all parts and sum up to 11 points. Good luck!

1. The hyperbolic bound $HB(n) : \prod_{1 \leq i \leq n}(U_i^{\mathcal{T}} + 1) \leq 2$, where $U_i^{\mathcal{T}}$ denotes the processor utilization of the task τ_i, is an example of a *sufficient* schedulability test.

   (a) (0.5) Give at least four conditions that need to hold to use this bound.

   (b) (0.5) What can be concluded concerning schedulability of a set $\mathcal{T}$ of n tasks when the hyperbolic bound does not hold for $\mathcal{T}$?

   (c) (1.0) Construct an example set of three tasks for which the left-hand side of $HB(3)$ is equal to 2.

   (d) (0.5) Is the task set that you constructed schedulable according to the Liu & Layland bound? Motivate your answer.
   Reminder: $LL(n) : U^{\mathcal{T}} \leq n(2^{1/n} - 1)$, where $U^{\mathcal{T}}$ denotes the processor utilization of the set $\mathcal{T}$.

   Answer: See exercise 1 of the examination for 2IN25 of January 6^{th}, 2009 for (b) - (d).

2. Consider two periodic tasks τ_1 and τ_2 with characteristics as given in the following table.

	$T = D$	C	φ
τ_1	3	1	0
τ_2	5	3	0

   (a) (0.5) What is a 'reasonable' length for a time line to illustrate the executions of the tasks? Motivate your answer.
   Answer: The activation pattern repeats itself after the hyperperiod H, which is equal to the least-common multiple of the tasks. Given a simultaneous release at time $t = 0$ and a utilization $U = \frac{1}{3} + \frac{3}{5} = \frac{14}{15} < 1$, the execution pattern will also repeat itself after H. A reasonable length is therefore $H = \text{lcm}(3, 5) = 15$.

   (b) (1.0) Draw a time line for the tasks under EDF (Earliest Deadline First). In case you make a specific decision, motivate it.
   Answer: See Figure 1. Because both tasks have a job with a deadline at $t = 15$, a decision has to be made which of those two jobs is allowed to run first. EDF does not prescribe which job has precedence in such a situation. In Figure 1, task τ_2 is allowed to continue.

3. Consider two periodic tasks τ_1 and τ_2 and a deferrable server S_{DS} with characteristics as given in the following table.

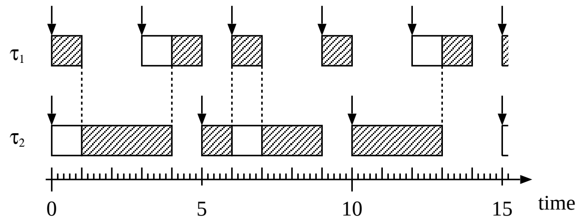

Figure 1: Example of a timeline under EDF.

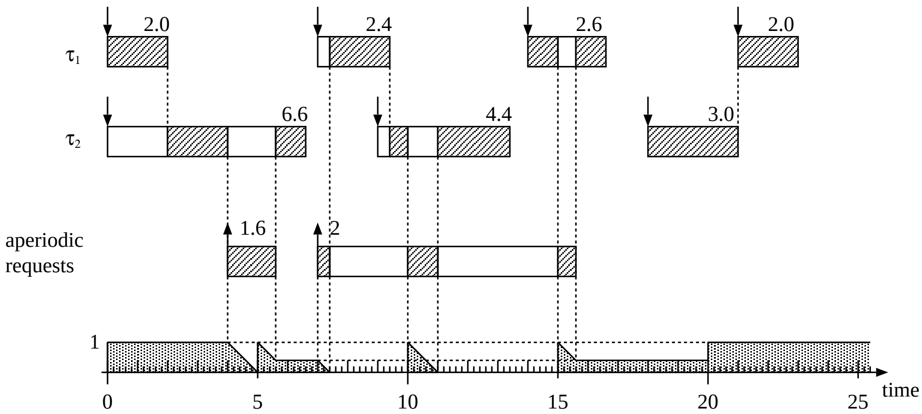

Figure 2: Example of high-priority Deferrable Server.

	$T = D$	C
S_{DS}	5	1
τ_1	7	2
τ_2	9	3

Assume scheduling based on FPPS and a rate monotonic priority assignment.

(a) (1.0) Assuming arrivals of aperiodic requests at time $t = 4$ for an amount of 1.6 and at time $t = 7$ for an amount of 2, draw time-lines illustrating the execution of the tasks and the remaining capacity of the deferrable server in an interval of length 25.
Answer: See Figure 2.

(b) (0.5) Are the tasks and the deferrable server schedulable under arbitrary phasing? Motivate your answer by means of *calculations*. *Hint*: Determine the worst-case response times of the tasks by means of the following recursive equation.

$$x = WC_i + \sum_{j<i} \left\lceil \frac{x + AJ_j}{WT_j} \right\rceil WC_j \tag{1}$$

Answer: Note that $AJ_{\text{DS}} = T_{\text{DS}} - C_{\text{DS}} = 4$. Based on (1), we find $WR_1 = 4 < D_1 = 7$ and $WR_2 = 10 > D_2 = 9$. Hence, task τ_2 misses its deadline and task τ_2 is therefore not schedulable.

4. Consider three tasks with arbitrary phasing and the following characteristics:

- an elastic task τ_1 with a minimum inter-arrival time of 5 and a maximum inter-arrival time of 6, a worst-case deadline equal to 6 and a best-case deadline equal to 2, and a fixed computation time of 2;

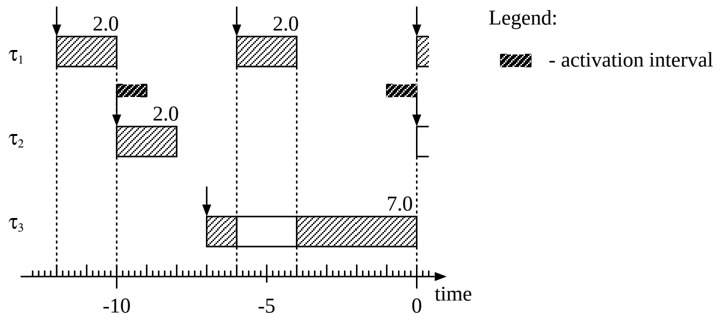

Figure 3: Time line with an optimal instant for task τ_3.

- a periodic task τ_2 with jitter with a fixed period of 9, activation jitter of 1, a worst-case deadline equal to 9 and a best-case deadline equal to 2, a best-case computation time of 2 and worst-case computation time of 3;
- a sporadic task τ_3 with a minimum inter-arrival time of 29, a worst-case deadline equal to 26 and a best-case deadline equal to 7, and a fixed computation time of 5.

Preparation towards an answer: The characteristics of the tasks are given below.

	WT	WD	BT	BD	AJ	WC	BC
τ_1	5	6	6	2	0	2	2
τ_2	9	9	9	2	1	3	2
τ_3	29	26	∞	7	0	5	5

(a) (0.5) Give a necessary and sufficient condition for the schedulability of the three tasks.
Answer: A necessary and sufficient, i.e. exact, condition is given by

$$\underset{1 \leq i \leq 3}{\forall} (BD_i \leq BR_i \wedge WR_i \leq WD_i), \tag{2}$$

where BD_i and WD_i are the best-case and worst-case deadline, respectively, and BR_i and WR_i are the best-case and worst-case response time, respectively

(b) Assume fixed-priority pre-emptive scheduling, where τ_1 has highest and τ_3 has lowest priority.

i. (0.5) Draw a time line with an optimal instant for task τ_3.
Answer: See Figure 3. Note that $BR_3 = 7$.

ii. (1.0) Determine the best-case response time of task τ_3 using the following recursive equation.

$$x = BC_i + \sum_{j<i} \left(\left\lceil \frac{x - AJ_j}{BT_j} \right\rceil - 1 \right)^+ BC_j \tag{3}$$

Answer: Using (1) we first determine the worst-case response time $WR_3 = 24$ as initial value for the iterative procedure to determine the best-case response time BR_3. Using (3), we subsequently find $BR_3 = 7$.

5. Semaphores have been conceived by Edgar W. Dijkstra in 1965. A general assumption of semaphores is the arbitrary selection of a process to proceed, i.e. if more than one process is waiting inside a P operation on the same semaphore and the semaphore becomes positive

(because of the execution of a V), one of the waiting processes is selected arbitrary to complete the P operation. This assumption has been changed for real-time systems.

(a) (1.0) What has been changed and why?
Answer: See slide 5 of RTA.B4-Policies-3.

(b) (0.5) Give two (different) examples for which semaphores can be used. Illustrate the examples by means of pseudo-code.
Answer: Synchronization and coordination of processes; see slide 5 of RTA.B4-Policies-3.

6. Explain the following concepts in your own words.

(a) (0.5) Priority inversion.
Answer: See slide 6 of RTA.B4-Policies-3 and book.

(b) (0.5) Push-through blocking.
Answer: See slide 8 of RTA.B4-Policies-3 and book.

7. (1.0) Explain the commonalities and differences between HLP (Highest Locker Protocol) and SRP (Stack Resource Protocol).
Answer: See slides RTA.B4-Policies-3 and book.