EINDHOVEN UNIVERSITY OF TECHNOLOGY

Department of Mathematics and Computer Science

Examination 2XN26 of Real-time Systems (2IN26)
on Thursday, April 16^{th} 2015, 13.30h-16.30h.

First read the entire examination. There are 8 exercises in total. Grades are included between parentheses at all parts and sum up to 10 points. *Motivate all your answers.* Good luck!

1. (1.0) Cyclic executives have as *advantages* that they are effective, have no "preemption costs" and are simple. Give at least four *disadvantages* of cyclic executives.
   Answer: see slides RTS.B3 - Cyclic Executives.

2. (0.5) The specification of a system may not include timing requirements. Where do the (hard real time) timeliness constraints in such a case come from? Give at least two examples.
   Answer: *Timeliness constraints* on a *system* are typically *derived* from the functional requirements and *determined* by (the environment and) design choices in the system; see slides RTS.B3 - Specification concepts.

3. (1.0) Describe the notions *system latency* and *event latency* in your own words.
   Answer: see slides RTS.B3 - Reference Model.

4. This question concerns *arrival patterns* of tasks.

   (a) (0.5) Let a task τ have a period T and activation jitter AJ. Describe the minimum inter-arrival time $T^{\min}$ and the maximum inter-arrival time $T^{\max}$ of τ in terms of T and AJ.
   Answer: $T^{\min} = T - AJ$ and $T^{\max} = T + AJ$.

   (b) (0.5) Describe the difference between the two arrival patterns *periodic with activation jitter* and *elastic* in your own words.
   Answer: see slides RTS.B3 - Reference Model.

   (c) (0.5) Sporadic tasks were introduced by Al Mok in his PhD-thesis in 1983. Describe a motivation for sporadic tasks.
   Answer: see slides RTS.B3 - Reference Model.

5. (1.0) Is the set of tasks shown in Table 1 schedulable by any scheduling algorithm under arbitrary phasing ? If yes, by which algorithm? If no, why not?

	T	D	C
τ_1	20	15	8
τ_2	10	5	3
τ_3	10	3	2

Table 1: Task characteristics

Answer: The set is not schedulable, irrespective of the scheduling algorithm. Consider a simultaneous release at time $t = 0$. The amount of work that needs to be executed in $[0, 15)$ is $2 \times C_3 + 2 \times C_2 + C_1 = 4 + 6 + 8 = 18$, i.e. *larger* than the length of the interval. Hence, the set is unschedulable.

6. In the following two fixed-priority scheduling (FPS) situations, the response times of all jobs in a so-called level-i active period have to be considered to determine the worst-case response time of task τ_i. Explain in your own words why looking at the first job alone is not sufficient for

   (a) (0.5) fixed-priority preemptive scheduling (FPPS) and *arbitrary deadlines*;

   (b) (0.5) fixed-priority scheduling with deferred preemption (FPDS) and *deadlines at most equal to periods*.

   Answer: See slides of the lectures.

7. This exercise concerns resource access protocols (RAPs).

   (a) (1.0) Transitive priority adjustment may occur when applying the Priority Inheritance Protocol (PIP). Describe at least three conditions that need to hold for transitive priority adjustment to occur.
   Answer: We essentially need to describe conditions that may give rise to transitive blocking: (1) at least three tasks, e.g. τ_h, τ_m, and τ_l, with distinct priorities, where τ_h has highest and τ_l has lowest priority; (2) at least two mutual exclusive resources, e.g. R_1 and R_2; (3) τ_h uses a resource R_1, τ_l uses R_2, and τ_m uses R_1 and R_2 in a nested fashion, i.e. it first locks R_1 and than locks R_2 before releasing R_1.

   (b) Assume fixed-priority pre-emptive scheduling of a set $\mathcal{T}$ of n tasks, $\tau_1, \ldots, \tau_n$. Let the index of a task represents its priority, where a higher priority is given by a lower value, i.e. τ_1 has highest and τ_n has lowest priority. Let the tasks share a set $\mathcal{R}$ of m resources. Let $C_{i,k}$ denote the worst-case computation time of task τ_i to access resource R_k, where $C_{i,k} \in \mathbb{R}^+ \cup \{0\}$.

      i. (0.5) Describe the resource ceiling of a resource $R_k \in \mathcal{R}$ under the Stack Resource Policy (SRP) in your own words and give a formula for that resource ceiling.

      ii. (1.0) Describe the worst-case blocking B_i of a task under SRP in your own words and give a formula for B_i.

   Answer: See RTS.B4-Policies-3-RAP.

8. This question concerns guest lectures.

   (a) (0.5) During his presentation on Scalable Video Algorithms, Prof. Hentschel mentioned that the number of operations (such as additions, multiplications, and memory accesses) alone is not a convenient measure for the resource needs of an algorithm. Explain why and briefly describe an alternative.
   Answer: From his slides: "*Resources on function level are difficult to manage and to control*". As an alternative a (-n estimated) resource-budget for the processor described in terms of, e.g. a *capacity* and a *period* can be used; see RTS.C6-Resource reservation.

   (b) In his 2^{nd} lecture on real-time communication, Dr. Behnam addressed Ethernet.

      i. (0.5) Why is Ethernet not suitable for real-time?

      ii. (0.5) Describe at least two techniques to make Ethernet applicable for real-time.

   Answer: See his slides.