EINDHOVEN UNIVERSITY OF TECHNOLOGY

Department of Mathematics and Computer Science

Examination Real-time Architectures (2YN26)
on Monday, April 16^{th} 2012, 10.30h-12.00h.

First read the entire examination. There are 6 exercises in total. Grades are included between parentheses at all parts and sum up to 10 points. Good luck!

1. This question concerns general Real-Time System issues.

   (a) (0.5) Explain why the response time of a task τ (of a system) can be considerably smaller than the response time of the system to an external event, when τ handles the event.
   Answer: There can be both a *delay* and *jitter* (*i*) between the arrival of the external event and the activation of the task and (*ii*) between the completion of the task and the response of the system. Reasons include processing at I/O devices, information that needs to be transported across a bus between these devices and the CPU, and intervals of times during which interrupts are disabled at the CPU.
   The question has been taken from the exam of January 20^{th}, 2011.

   (b) (0.5) Is a task the same as a thread or process? Motivate your answer.
   Answer: See slide "Misconceptions" of set "Reference Model".

2. This question concerns fixed-priority scheduling with deferred preemption (FPDS).

   (a) (1.0) Why is the exact worst-case response time analysis for FPDS non-uniform?
   Answer: All tasks can be blocked by a lower priority task, except for the lowest priority task. Because a lower priority task τ_{L} has to start its execution of a sub job *before* the release of a higher priority task τ_{H} in order to block τ_{H}, the blocking time is a *supremum* rather than a *maximum*. The fact that blocking is a supremum is in itself not enough to cause the non-uniformity, however; see for example blocking in the context of fixed-priority pe-emptive scheduling, which results in a uniform analysis. The distinguishing characteristic for FPDS is the fact that the last sub job can't start before all pending work of previous sub jobs and all higher priority tasks has been completed. For a fixed (maximum) blocking time, this starting time of the final sub job is expressed using a floor-function. Because (non-zero) blocking gives rise to a supremum rather than a maximum, we have to consider a limit-case, resulting in a change from a floor-function (plus 1) to a ceiling-function. The latter causes the analysis to become non-uniform. See the slides for more details.
   The question has been taken from the exam of January 20^{th}, 2011.

   (b) (0.5) Give an example of a taskset that is schedulable with FPPS but not with FPDS.
   Answer: Whenever $C_1 = D_1 < T_1$, τ_1 cannot be deferred by any lower priority task.

   (c) (0.5) Which constraint must hold to allow resource sharing between tasks scheduled by FPDS without an additional resource access protocol?
   Answer: Critical sections are not allowed to cross sub job boundaries.
   The question has been taken from the exam of January 20^{th}, 2011.

3. This question concerns communication.

   (a) (1.0) Are the "ordinary" communication protocols Ethernet and Token ring appropriate for hard real-time systems? Motivate your answer.
   Answer: See slides.

   (b) (1.0) Describe the abstraction of a CAN network in your own words.
   Answer: See slides.

4. One of the lectures concerned "*A QoS approach for multimedia consumer terminals with media processing in software*". The aim of the QoS approach was *cost-effective high-quality video processing in software for multimedia consumer terminals*, motivated by the requirements for *openness* and *flexibility* of these systems, and having as boundary condition that *the existing system qualities should be preserved.*

   (a) (0.5) Explain which real-time problems were addressed.

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

   Answer: See slides. The question has been taken from the exam of April 16^{th}, 2010.

5. The following questions concern Hierarchical Scheduling Frameworks (HSFs).

   (a) (1.0) In "[Shin et al 03] I. Shin and I. Lee, *Periodic resource model for compositional real-time guarantees*, In: Proc. 24^{th} IEEE Real-Time Systems Symposium (RTSS), pp. 2-13, December 2003", a *resource supply bound function* $\texttt{sbf}_\Gamma(t)$ of a time interval of length t is defined that calculates the minimum resource supply of Γ during t units. Given a periodic resource $\Gamma(\Pi, \Theta)$, draw $\texttt{sbf}_\Gamma$ as a function of t for $0 \leq t \leq 5\Pi$, where $\Theta = \Pi/3$.
   Answer: See paper. Note that the so-called "blackout duration" (i.e. the longest time without resource supply) is equal to $2(\Pi - \Theta) = \frac{4}{3}\Pi$. Further note that a similar question was asked in the exam of April 16^{th}, 2010.

   (b) (1.0) Explain the difference between SIRAP and H-SRP.
   Answer: see papers. The question has been taken from the exam of January 23^{rd}, 2012.

6. In "[Bertogna et al 07] M. Bertogna, N. Fisher, and S. Baruah. *Static-priority scheduling and resource hold times*, In: Proc. 15^{th} International Workshop on Parallel and Distributed Real-Time Systems, pp. 1-8, March 2007.", the notion of *resource hold time* is defined for fixed-priority scheduling using the Stack Resource Policy (SRP) and algorithms are presented to (*i*) calculate resource hold times and (*ii*) to minimize resource hold times.

   (a) (0.5) Describe the notion of *resource hold time* in your own words.

   (b) (1.0) Describe *how* a resource hold time can be reduced and *what* determines its minimum.

   Answer: See paper. The question has been taken from the exam of April 11^{th}, 2011.