EINDHOVEN UNIVERSITY OF TECHNOLOGY

Department of Mathematics and Computer Science

Examination Real-time Architectures (2XN26)
on Wednesday, November 3^{rd} 2010, 14.00h-15.30h.

First read the entire examination. There are 9 exercises in total. Grades are included between parentheses at all parts and sum up to 10 points. Good luck!

1. Interrupts generated by I/O peripheral devices represent a big problem for the predictability of a real-time system. Three approaches are described in the book of Buttazzo to handle interrupts, two based on disabling interrupts and by using polling (by either application tasks or dedicated kernel routines) and a third where all external interrupts are left enabled.

   (a) (0.5) Describe the third approach. *Hint*: Distinguish between *device driver* and *device manager* and address the role of the operating system.
   Answer: See book.

   (b) (0.5) Describe at least 2 advantages of this third approach.
   Answer: Advantages (see book): (*i*) busy wait during I/O operations is eliminated; (*ii*) interrupt interference is minimized (i.e. to the *activation* of a proper task only); (*iii*) priority of the task managing the device is completely independent from other priorities and can be set according to the application requirements.

2. (0.5) Suppose you have both a *relative* and an *absolute* one-shot timer at your disposal, which type would you prefer to implement periodic tasks and why?
   Answer: A *relative* one-shot timer may give rise to jitter in the activation of jobs, whereas an *absolute* one-shot timer does not. For that reason, the latter is preferred.

3. (0.5) EDF and RMS are both optimal in some sense. Explain in which sense.
   Answer: See book.

4. (0.5) Let a set of hard periodic tasks with given characteristics be schedulable by means of both EDF and RMS. Assume an overrun happens for a particular task, e.g. that task requires more computation time than its given worst-case computation time or arrives more frequently than its minimum inter-arrival time. Which tasks can potentially miss their deadline due to such an overrun under

   (a) EDF;
   Answer: All tasks, including the task experiencing the overload.

   (b) RMS.
   Answer: The task experiencing the overload, and all tasks with a lower priority.

5. (0.5) The *absolute* finalization time of a job $\tau_{i,k}$ is given by $f_{i,k}$. Define the *relative* finalization time $F_{i,k}$ of a job $\tau_{i,k}$ of a periodic task τ_i with period T_i, activation jitter AJ_i and phasing φ_i.
   Answer: $F_{i,k} = f_{i,k} - (\varphi_i + kT_i)$; the time is relative to the start of the activation interval in which $\tau_{i,k}$ is activated; see Figure 1. Note that φ_i is the start of the activation interval of $\tau_{i,0}$.

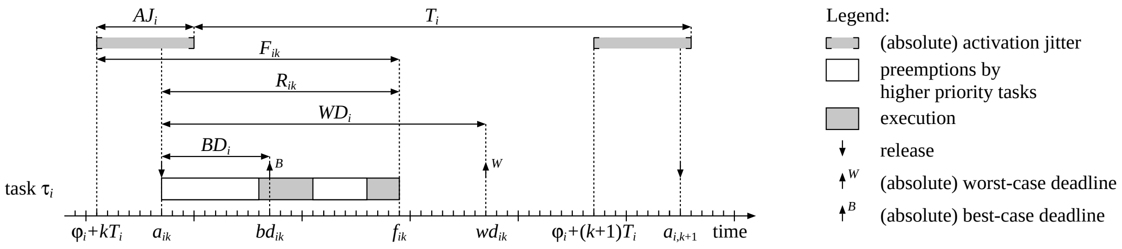

Figure 1: Relative finalization time $F_{i,k}$ of $\tau_{i,k}$.

6. Assume a set Γ of n tasks $\tau_1, \tau_2, \ldots, \tau_n$, where tasks are indexed in order of decreasing priority, i.e. τ_1 has highest and τ_n has lowest priority. When the tasks are scheduled by means of fixed-priority pre-emptive scheduling, recursive equation (1) can used to determine the worst-case response time of a task τ_i, assuming certain additional conditions.

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

(a) (0.5) Give at least 3 conditions that need to hold to use this equation.
Answer: Conditions include (see also RTS.B5-Analysis-2-FPPS): (*i*) tasks have unique priorities; (*ii*) no self-suspension of tasks; (*iii*) (worst-case) relative deadline WD_i is at most equal to the (worst-case) period WT_i minus activation jitter AJ_i, i.e. $WD_i \leq WT_i - AJ_i$; (*iv*) overhead of scheduling and context switching is ignored.
Note: arbitrary phasing is not a condition (or constraint). For a specific phasing, the equation can stil be used, but may yield a pessimistic result.

(b) (0.5) What do the left-hand side and the right-hand side of the equation represent?
Answer: LHS - amount of time *available* in $[0, x)$. RHS - *max.* amount of time *requested* by τ_i and τ_j with $1 \leq j < i$ in $[0, x)$.

(c) (0.5) Give a *sufficient* worst-case condition for schedulability of Γ using (1).
Answer: When $WD_i \leq WT_i - AJ_i$ holds for *all* $1 \leq i \leq n$, we can use

$$\underset{1 \leq i \leq n}{\forall} WD_i \geq B_i + WC_i + \sum_{1 \leq j < i} \left\lceil \frac{WD_i + AJ_j}{WT_j} \right\rceil WC_j. \tag{2}$$

7. Consider three tasks that are scheduled by means of fixed-priority pre-emptive scheduling, where τ_1 has highest and τ_3 has lowest priority, with arbitrary phasing and characteristics as given in the following table.

	T	WD	BD	AJ	WC	BC
τ_1	7	10	2	0	2	2
τ_2	9	4	2	1	3	3
τ_3	31	27	10	0	8	6

(a) Determine the best-case response time of task τ_3 by

i. (0.5) drawing a time line with an optimal instant for τ_3.
Answer: See Figure 2. Note that $BR_3 = 11$.

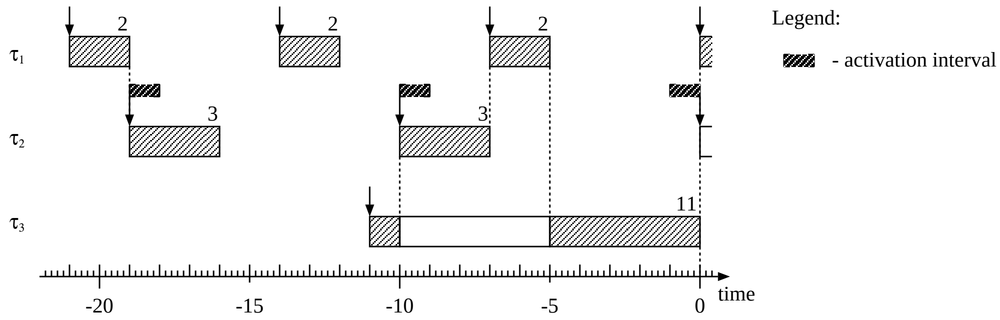

Figure 2: Timeline with an optimal instant for task τ_3.

ii. (1.0) means of 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 $WR_3 = 25$ as initial value for the iterative procedure to determine the best-case response time BR_3. Using (3), we subsequently find $BR_3 = 11$.

(b) (0.5) Is task τ_3 schedulable?
Answer: Yes, because $BD_3 = 10 \leq BR_3 = 11$ and $WR_3 = 25 \leq WD_3 = 27$.

8. (1.5) Consider a system consisting of a set of hard real-time tasks and a server, which is used to service aperiodic requests. Describe the consequence of selecting either of the following server types for the best-case and worst-case response times of hard-real-time tasks.

   (a) a *sporadic* server;
   (b) an *idling* periodic server;
   (c) a *gain-time providing* periodic server.

   Answers: The choice only potentially influences the best-case response time of hard real-time tasks with a priority lower than the server; see RTA.Exercises-5.

9. Resource access protocols:

   (a) (0.5) Semaphores were conceived by Edsger W. Dijkstra. Which specific assumption for semaphores is needed to make them applicable in the context of real-time systems?
   Answer: See slides RTS.B4-Policies-3-RAP.
   (b) (1.0) Compare PIP and SRP, using at least 5 evaluation criteria.
   Answer: See book and/or slides.
   (c) (0.5) In the paper "David Polock and Dieter Zöbel, *Conformance Testing of Priority Inheritance Protocols*, In: Proc. 7^{th} IEEE International Conference on Real-Time Computing Systems and Applications (RTCSA), pp. 404 - 408, December 2000", the authors show a defect in the original PIP protocol informally described by Sha, Rajkumar and Lehoczky.
   Describe the defect in your own words. *Hint*: it concerned transitive adjustment of priorities.
   Answer: See paper.