

PARALLEL AND DISTRIBUTED REAL-TIME SYSTEMS

EDA420

Final exam, March 17, 2005 at 08:30 – 12:30 in the V building

Examiner:

Associate professor Jan Jonsson
Phone: 031-772 5220

Content:

The written exam consists of 4 pages (including cover), containing 7 problems worth a total of 60 points.

Grading policy:

24–35 ⇒ grade 3
36–47 ⇒ grade 4
48–60 ⇒ grade 5

Restrictions:

Books, notes or calculators are NOT allowed (only writing material and dictionaries)

Solution:

No solution provided.

Results:

Posted on the department's information boards on Thursday, March 31, 2005 at 09:00.

Inspection:

Room 4128, Rännvägen 6 B, on Thursday, March 31, 2005 at 13:00–15:00.

Language:

Your solutions should be written in Swedish or English.

IMPORTANT ISSUES

1. Use separate sheets for each answered problem, and mark each sheet with the problem number.
 2. Mark the first sheet with your name and "personnummer".
 3. Justify all answers. Lack of justification can lead to loss of credit even if the answer might be correct.
 4. Explain all calculations thoroughly. If justification and method is correct then simple calculation mistakes do not necessarily lead to loss of credit.
 5. If some assumptions in a problem are missing or you consider that the made assumptions are unclear, then please state explicitly which assumptions you make in order to find a solution.
 6. Write clearly! If I cannot read your solution, it is wrong.
-

GOOD LUCK!

PROBLEM 1

State whether the following propositions are TRUE or FALSE. For each correct statement, you will be given 1 point; for each erroneous statement you will be given -1 point; if you make no statement at all, you will be given 0 points. **Quality guarantee:** The total result for this problem cannot be less than 0 points. (6 points)

- a) When there are mutual exclusion constraints in a system, it is impossible to find an optimal on-line scheduling algorithm (unless it is clairvoyant).
 - b) By the term *critical time instant* in schedulability analysis we mean a point in time where all tasks' deadlines coincide.
 - c) One of the so-called *Richard's anomalies* states that task completion times may increase as a result of increasing the number of processors.
 - d) One main differences between a *guarantee scheme* and a *robust scheme* for on-line task admission in a real-time system is that tasks may have a "second chance" in the robust scheme, even if initially rejected.
 - e) A *sufficient* feasibility test is one such that, if a given task set is not schedulable, it always reports the answer no.
 - f) In the Fault-Tolerant Average algorithm for clock synchronization, the new clock value is the mean of all collected clock readings excluding the t fastest and t slowest clocks.
-

PROBLEM 2

The following questions are related to estimation of worst-case execution times (WCET) of tasks.

- a) Explain why is it preferred that WCET estimates for tasks in a real-time system are *pessimistic* as well as *tight*. (2 points)
 - b) Explain what the term *false path* refers to in the context of WCET analysis. In conjunction with this, give an example of a program source code for which there exists a false path and also explain the how that false path could negatively affect the WCET estimate. (4 points)
 - c) In the case that preemptive scheduling is used in a system, is it appropriate to include the cost for task preemptions in the worst-case execution time of a task? Justify your answer! (2 points)
-

PROBLEM 3

The following questions are related to fault-tolerant scheduling.

- a) Describe how existing scheduling techniques are generally extended towards fault-tolerance. (4 points)
 - b) List the schedulability utilization bound derived by Pandya and Malek for fault-tolerant uniprocessor rate-monotonic scheduling. State under what assumptions this bound applies. (1 points)
 - c) Explain why it is important to assume a realistic *fault model* in fault-tolerant scheduling. Give two examples of fault models and their implications on the scheduling method used. (3 points)
-

PROBLEM 4

Consider the MULTIPROCESSOR SCHEDULING decision problem — that is, the problem of determining schedulability for a set of independent tasks with a common deadline using a set of identical processors.

- a) Describe the general procedure for proving that a decision problem, such as the MULTIPROCESSOR SCHEDULING problem, is NP-complete. (4 points)
 - b) The MULTIPROCESSOR SCHEDULING decision problem is said to be NP-complete in the strong sense for an arbitrary number of processors. It is also a so-called number problem. Describe how knowledge about whether a problem is a number problem or not can be used to decide if strong NP-completeness applies. (2 points)
 - c) Can the corresponding MULTIPROCESSOR SCHEDULING optimization problem also be proven to be NP-complete? Motivate your answer. (2 points)
-

PROBLEM 5

Describe for each of the following network communication protocols how they can be used in a real-time system. In particular, describe how upper-bounded message delivery times are guaranteed for each protocol.

- a) The Virtual-Time, Carrier-Sense Multiple Access (VTCSMA) protocol. (2 points)
 - b) The Controller Area Network (CAN) protocol. (2 points)
 - c) The Token Ring protocol (IEEE 802.5). (2 points)
 - d) The Earliest-Due-Date-Deadline (EDD-D) protocol. (2 points)
-

PROBLEM 6

Consider a real-time system with periodic tasks and a run-time system that employs preemptive scheduling using the rate-monotonic (RM) priority assignment approach.

- a) Explain the principle behind the RM priority assignment approach. (1 points)
- b) Describe shortly two different analysis methods that can be used for deciding the schedulability of tasks being scheduled with RM. (3 points)

Now assume that the given system consists of two tasks. For each task τ_i it applies that its deadline D_i is equal to the period T_i . Both tasks arrive the first time at time 0. It is known that the tasks' utilizations $U_i = C_i/T_i$ are $U_1 = 0.75$ och $U_2 = 0.25$, respectively. However, the individual execution times C_i for the two tasks are not known.

- c) If the tasks' periods T_i are not known either, is it possible to decide if the tasks are schedulable with RM? Motivate your answer. (3 points)
 - d) If we know that $T_2 = 2T_1$, is it then possible to decide whether the tasks are schedulable with RM or not? Motivate your answer. (3 points)
-

PROBLEM 7

Consider a real-time system with sporadic tasks and a run-time system that employs preemptive scheduling using the earliest-deadline-first (EDF) priority-assignment approach.

- a) Explain the concept of a sporadic task. Also give an example of a real-time application where it is suitable to model the system with sporadic tasks. (2 points)
- b) At a certain stage in the execution of the system, four sporadic tasks, τ_1, τ_2, τ_3 and τ_4 , arrive according to the following scenario:

- At time $t = t_a$, task τ_2 arrives with an execution time $C_2 = 3$ and a relative deadline $D_2 = 7$.
- At time $t = t_a + 0.5$, task τ_4 arrives with an execution time $C_4 = 1.5$ and a relative deadline $D_4 = 9$.
- At time $t = t_a + 1.5$, task τ_3 arrives with an execution time $C_3 = 1.5$ and a relative deadline $D_3 = 6.5$.
- At time $t = t_a + 2$, task τ_1 arrives with an execution time $C_1 = 3$ and a relative deadline $D_1 = 4$.

Show, by constructing a timing diagram, that the task instances given above are schedulable. Assume, for simplicity, that the system is not currently executing any other task at $t = t_a$ and that no new instances of the four tasks arrive before $t = t_a + 15$. (4 points)

- c) Assuming that the arrival scenario given in sub-problem b) still applies, now add the following event to the system:

- At time $t = t_a + 3$, task τ_0 arrives with an execution time $C_0 = 1.5$ and a relative deadline $D_0 = 2$.

What happens with the schedulability of the five tasks? Illustrate your answer with a timing diagram. (3 points)

- d) Based on the observations in sub-problems b) and c), what can be said about the competitive factor of EDF when used as an on-line scheduler? (1 points)
- e) What is the best achievable competitive factor for an on-line scheduler according to Baruah *et al.*? Under what assumptions does this result apply? (2 points)
-