Parallel and Distributed Real-Time Systems
EDA420

Final exam, March 11, 2004 at 08:45 – 12:45 in the M 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 Wednesday, March 24, 2004 at 09:00.

Inspection:
Room 4128, Rännvägen 6 B, on Wednesday, March 24, 2004 at 10:00–12:00.

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) In multiprocessor scheduling, *Dhall's effect* refers to the fact that there exists no static-priority assignment scheme that has a non-zero schedulability utilization bound.

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) An important assumption in worst-case execution-time analysis is that interferences from higher-priority tasks are accounted for.

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

For static (offline) multiprocessor scheduling, one useful algorithm is Simulated Annealing (SA).

a) Describe the major principles behind the SA algorithm when used in the general sense (not necessarily for real-time scheduling). (4 points)

b) Describe how the SA algorithm can be adapted to solve the real-time multiprocessor scheduling problem (according to Tindell, Burns and Wellings). (2 points)

c) When evaluating the scheduling performance of a technique such as the SA algorithm, what is the major evaluation methodology used (theoretical analysis, simulation of synthetic workloads or execution of real applications)? Motivate your answer! (2 points)

PROBLEM 3

The following questions are related to fault-tolerant scheduling.

a) Describe how existing *uniprocessor* scheduling techniques are generally extended towards fault-tolerance. (3 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) Describe how existing *multiprocessor* scheduling techniques are generally extended towards fault-tolerance. (3 points)

d) 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

In their research paper "On Non-Preemptive Scheduling of Periodic and Sporadic Tasks", Jeffay *et al.* address a decision problem, called Non-Preemptive Scheduling of Concrete Periodic Tasks (SCPT). This problem concerns determining whether a set of periodic tasks with specified offsets (release times) is schedulable under non-preemptive task dispatching. In that paper, the authors show that the SCPT problem is *NP-complete in the strong sense*.

a) Describe the general four-step procedure that the authors used for proving that the SCPT problem is NP-complete. (4 points)

b) The SCPT problem is NP-complete in the *strong sense*. Describe the meaning of strong NP-completeness. (2 points)

c) In the same paper, the authors also address the problem of determining whether a set of *sporadic* tasks with specified offsets (release times) is schedulable under non-preemptive task dispatching. That problem is shown to be solvable using a *pseudo-polynomial time algorithm*. Describe the difference between such an algorithm and a pure polynomial-time algorithm. (2 points)

PROBLEM 5

When a communication network is used in a real-time system it is important to know the worst-case *message delay* between the sender and receiver of an interprocessor data exchange. Two important components of the message delay are the *queuing delay* and the *transmission delay*.

a) Describe briefly the underlying causes of the queuing delay in the following networks: TDMA, Ethernet and Token Ring. (3 points)

b) Describe the mechanisms that give rise to the queuing delays on a Controller Area Network (CAN) bus system. (3 points)

c) Describe how the transmission delay is calculated for a communication link. (2 points)

d) Describe two techniques for assigning deadlines to interprocess communication. (2 points)

PROBLEM 6

One method for analyzing the performance of a system that experiences overload is to use the concept of *competitive factor*.

a) Give the formal definition of competitive factor. (4 points)

b) 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)

c) Give the competitive factor for the earliest-deadline-first (EDF) scheduling policy. Also comment on the implications of this result. (2 points)

PROBLEM 7

Consider the timing diagram below describing the execution order of two periodic tasks when rate monotonic (RM) scheduling is used. Since RM is used, it applies for each task τ_i that its deadline D_i is equal to the period T_i. The first instance of each task arrives at time 0. The tasks are schedulable, that is, they meet their deadlines. In the timing diagram we use τ_i^k to denote instance k of the periodic task τ_i. All values are given in milliseconds. The cost for switching tasks is assumed to be negligible.

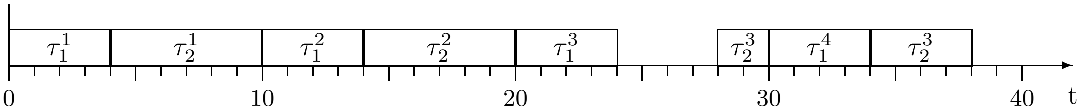

a) Decide based on the diagram above the execution time C_i and period T_i for each of the two tasks. Provide a motivation for your choice. (2 points)

b) Calculate, based on the derived values of C_i and T_i in subproblem a), the utilization U for this set of tasks. (1 points)

c) As was mentioned above, this task set meets the deadlines. However, if the execution time of any of the tasks in increased a deadline will be missed (in this case, the deadline for the first instance of τ_2). Such a task set is said to *fully utilize* the processor. Derive for the following values of C_1 the corresponding values of C_2 that make the tasks fully utilize the processor: 0, $0.5C_1'$, and $2C_1'$. Here, C_1' is the original value derived in subproblem a) (4 points)

d) Calculate, for each pair of values of C_1 and C_2 derived in subproblem c), the utilization U of the task set when the processor is fully utilized. Comment your result! (3 points)