Parallel and Distributed Real-Time Systems EDA421/DIT171

Final exam, May 23, 2012 at 14:00 – 18:00 in the V building

Examiner:
Professor Jan Jonsson
Phone: 031–772 5220

Content:
The written exam consists of 5 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 course home page on Monday, June 11, 2012 at 09:00.

Inspection:
Room 4128, Rännvägen 6 B, on Monday, June 11, 2012 at 13:00–15:00. Inspection at another occasion could be arranged by contacting the course examiner.

Language:
Your solutions should be written in English.

IMPORTANT ISSUES

1. Use separate sheets for each answered problem, and mark each sheet with the problem number.
2. Justify all answers. Lack of justification can lead to loss of credit even if the answer might be correct.
3. Explain all calculations thoroughly. If justification and method is correct then simple calculation mistakes do not necessarily lead to loss of credit.
4. 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.
5. 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) An optimization problem can never be classified as an NP-complete problem.

b) By *synthetic utilization* in server-less handling of aperiodic tasks, we mean an estimation of the workload caused by the resident periodic tasks in the system.

c) In fault-tolerant multiprocessor scheduling the term *overloading* refers to a situation where passive backups for multiple primary tasks are scheduled during the same time interval on one processor.

d) Switched Ethernet is a contention-free communication protocol.

e) If a given task set is known to be schedulable, a necessary feasibility test will always report the answer yes when applied to that task set.

f) By *high accuracy* in a clock synchronization context, we mean a small deviation between clocks.

PROBLEM 2

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

a) Explain what the term *infeasible path* refers to in the context of WCET analysis and provide an example of such a path. Also explain why it is useful to identify infeasible paths. (3 points)

b) 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! (3 points)

PROBLEM 3

The following questions are related to global multiprocessor scheduling.

a) Describe the main difference between global and partitioned multiprocessor scheduling. Also list two major advantages with global scheduling. (3 points)

b) One well-known drawback with global multiprocessor scheduling is a dilemma referred to as *Dhall's effect*. Describe the basic features of Dhall's effect and also give an example of a task set that suffers from the effect. (4 points)

c) Describe three types of system changes that can cause scheduling anomalies in global multiprocessor scheduling. (3 points)

PROBLEM 4

For static (off-line) multiprocessor scheduling, a useful method is the *branch-and-bound* (B&B) algorithm.

a) In many B&B algorithms, the search for an optimal schedule begins with an empty schedule. Describe the general approach used for extending an empty or partial schedule towards a complete schedule (that is, how child vertices are generated.) (3 points)

b) For the type of B&B algorithm mentioned in subproblem a), describe briefly four methods that can be used to avoid examining all possible schedules in search for the optimal schedule. (4 points)

c) In the B&B algorithm by Xu and Parnas, the search for an optimal schedule instead begins with an initial complete schedule. Explain briefly how new improved candidate schedules are generated in that algorithm. In your answer it is not necessary to describe all variables involved in the generation of child vertices, but rather your answer should give the intuition behind the choice of candidates. (3 points)

PROBLEM 5

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 procedure used for proving that a decision problem, such as the MULTIPROCESSOR SCHEDULING problem, is NP-complete. (4 points)

b) The MULTIPROCESSOR SCHEDULING problem is a so-called number problem. Describe the formal characteristics of a number problem. (2 points)

c) The MULTIPROCESSOR SCHEDULING problem is said to be NP-complete in the strong sense, which means that it cannot be solved by a pseudo-polynomial-time algorithm. Describe the formal characteristics of a pseudo-polynomial-time algorithm. (2 points)

PROBLEM 6

Consider the following static-priority assignment policies for preemptive scheduling of periodic independent tasks:

- Rate-Monotonic Scheduling (RM)
- Deadline-Monotonic Scheduling (DM)
- Audsley's OPA Algorithm (OPA)

a) State for each policy (RM, DM, and OPA) if it is optimal for scheduling of *synchronous* task sets on a *uniprocessor* system. (3 points)

b) State for each policy (RM, DM, and OPA) if it is optimal for scheduling of *asynchronous* task sets on a *uniprocessor* system. (3 points)

c) Which of the three policies (RM, DM, and OPA) are optimal for global scheduling on a *multiprocessor* system? (2 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) 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 using EDF. 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)

b) Show, by adding one more sporadic task arrival to the schedule, that EDF can yield a *cumulative value* of zero for the original four tasks if they have hard deadlines. Assume that a task's execution is not aborted if it misses its deadline. (3 points)

c) Give an example of another priority assignment (than EDF) that will yield a non-zero cumulative value for the original four tasks in the presence of the new sporadic task that you added in sub-problem b). Again, assume that a task's execution is not aborted if it misses its deadline. (3 points)

d) 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)