Exam for Real Time Systems

2016 Oct 17 (14-19)

Wang Yi

Important Instructions:

1. No course material or computer/calculator are allowed, only a pen and a dictionary.
2. Please mark which course you are registered for:

◯ **5hp** (1DT063) You need to solve problems 1–6 only.	◯ **10hp** (1DT004) You need to solve *all* problems.

3. For each problem, a number of choices/statements are given, that may be correct or wrong.

   You are asked to mark only the **correct** ones.

   If you mark all correct choices for each problem, it will give you all points of the problem; otherwise each wrong choice (marked) will result in "negative points" proportionally.

....

Problem 1

1. *(5p) Why is it so difficult to predict the timing behaviour of a computer system?*

   ( ) *it may run software with too many lines of code*

   (X) *it may have shared resources*

   (X) *it is impossible to decide whether a program terminates or not*

   ( ) *it may consume too much energy*

   (X) *it may have cache memory*

   ( ) *it may not support direct memory access (DMA)*

   (X) *it may have a processor with branch prediction*

2. *(5p) The characteristics of RTOS are:*

   ( ) *typical footprints: 10MB to 1GB*

   (X) *typical footprints: 0.2 to 1MB*

   (X) *deterministic and fast responsiveness*

   ( ) *fast and time-sharing*

   (X) *possible to change scheduling policies*

   ( ) *tasks are scheduled using EDF*

3. *(5p) Resource servers can be useful for*

   (X) *isolating applications of different types*

   (X) *hierarchically designing systems that may have many sub-systems*

   (X) *improving the average response times of soft real-time tasks*

   (X) *guaranteeing the deadlines of event-driven hard real-time tasks*

4. *(5p) When RMS is used, to ensure schedulability, the maximal resource utilization is bounded by:*

   (X) 100%

   ( ) 99%

   ( ) $n * (2^{1/n} - 1)$ *where* n *is the number of tasks*

   ( ) 69.3%

5. *(5p) The advantages of static scheduling include:*

   ( ) *it is flexible and easy to change*

   (X) *it can be used to implement predictable and deterministic systems*

   ( ) *it can be used to implement complex systems with large task sets*

   (X) *it may use less memory than dynamic scheduling*

   (X) *it can be used to implement safety-critical systems*

6. *(5p) We know that the bandwidth is 1Mbps for a CAN bus whose length is 50 meters. What is the bandwidth for a CAN bus of 100 meters?*

   ( ) *3-4Mbps*

   ( ) *1-2Mbps*

   (X) *0.2-0.5 Mbps*

   ( ) *0.1-0.15 Mbps*

7. *(5p) For a system with periodic tasks that are dependent on each other with input/output relations, the response times of tasks may vary. Is it possible to avoid this and thus avoid release jitters?*

   ( ) *this is not possible at all*

( ) *yes, deliver the result of a job as soon as possible*

( ) *yes, deliver the result of a job as late as possible*

(X) *yes, deliver the result of a job at the end of each period*

8. *(5p) Which statements are correct?*

( ) *EDF and DMS are non stable*

(X) *EDF is non stable; DMS is stable*

( ) *both DMS and EDF uses dynamic priorities*

( ) *EDF is optimal, DMS is not*

(X) *both EDF and DMS are optimal*

9. *(5p) Assume that a set of periodic tasks with deadlines equal to periods is schedulable on a single processor system using Shortest Job First Scheduling Algorithm. Then the task set is also schedulable using*

(X) *RMS*

( ) *Non-Preemptive EDF*

(X) *EDF*

( ) *FIFO*

(X) *DMS*

10. *(5p) When you use CAN to implement a distributed system, you should*

(X) *make sure that messages sent from different nodes have different priorities*

(X) *know that sending messages with the same identity by different nodes may lead to errors*

( ) *send messages with periodic tasks*

( ) *assign the priorities of messages according to the periods of the sending tasks*

(X) *make sure that the priority ordering of messages is given according to their id's with the lower id number, the higher priority for each message*

(X) *send messages without receiver's explicit address*

Problem 2 *(10p) For a periodic task set with utilization 1, what would be the worst case response time of a task if EDF is used for scheduling?*

(X) *the period of the task*

( ) *the worst case execution time of the task*

( ) *any number larger than the worst case execution time of the task*

( ) *it could be larger than the period of the task*

( ) *any number less than the period of the task*

(X) *any number between the worst case execution time and the period of the task*

Problem 3 *(10p) For a set of sporadic sets (C_i, D_i, T_i) where deadline $D_i \leq T_i$. Identify the correct statements in the following:*

(X) *if the task set is feasible then* $\sum C_i/T_i \leq 1$

( ) *if* $\sum C_i/T_i \leq 1$*, the task set is feasible*

( ) *if the task set is feasible then* $\sum C_i/D_i \leq 1$

(X) *if* $\sum C_i/D_i \leq 1$*, the task set is feasible*

(X) *if* $\sum C_i/D_i \leq 0.693$*, the task set is schedulable with DMS*

( ) *if the task set is schedulable then* $\sum C_i/D_i \leq 0.693$

( ) *if* $\sum C_i/T_i \leq 0.693$*, the task set is schedulable with DMS*

( ) *if the task set is schedulable then* $\sum C_i/T_i \leq 0.693$

Problem 4 *(10p) Assume a set of* 12 *non-preemptive sporadic tasks whose computation times are all* $0.27ms$ *and whose periods are* $1ms, 2ms, 2ms, 7ms, 8ms, 10ms, 18ms, 20ms, 30ms, 31ms, 43ms$ *and* $50ms$ *respectively. Assume that RMS is used for scheduling, there are no jitters and no context-switch overheads. The worst-case response time for the task with period 10ms is bounded by:*

( ) $0.27ms$, *( )* $0.61ms$, *( )* $0.82ms$, *( )* $1.22ms$, *( )* $1.81ms$, *(X)* $2.43ms$, *(X)* $2.5ms$, *(X)* $2.6ms$

(X) $2.7ms$, *(X)* $2.8ms$, *(X)* $3.1ms$, *(X)* $3.51ms$, *(X)* $4.32ms$, *(X)* $5.67ms$, *(X)* $7.02ms$

Problem 5 *(10p) What are the differences between BIP (Basic Priority Inheritance Protocol) and HLP (Immediate Priority Inheritance Protocol)?*

( ) *BIP can avoid deadlocks, HLP can not*

(X) *HLP can avoid deadlocks, BIP can not*

( ) *BIP lets a task run at a higher priority whenever it locks a semaphore*

(X) *HLP lets a task run at a higher priority whenever it locks a semaphore*

( ) *HLP is easier to implement than BIP*

( ) *Using HLP, the blocking time a higher priority task may suffer is the sum of the worst case execution times that tasks with lower priorities may run using the resources shared with the higher-priority task.*

(X) *Using HLP, the blocking time a higher priority task may suffer is the max of the worst case execution times that tasks with lower priorities may run using the resources shared with the higher priority task.*

Problem 6 *(10p) Assume a set of event-driven tasks with computing times* C_i *and deadlines* D_i *(assume* $C_i \leq D_i$*). Both may arrive at any time but the minimal distance between two arrivals of each task is* $T_i > D_i$*. Design a polling server for the tasks with period* T *and capacity* C *such that no deadline is violated when RMS is used for scheduling. Assume* $C \leq T$ *and* R *is the response time estimated using RMS analysis when the server is considered as a periodic task. Identify correct statements:*

(X) $T \leq D_i$ *is necessary*

( ) $T + T \leq D_i$ *is necessary*

(X) $T + R \leq D_i$ *is necessary and sufficient*

(X) $T + C \leq D_i$ *is necessary*

(X) $\sum C_i \leq C$ *is necessary*

(X) $\sum C_i \leq C$ *and* $T + T \leq D_i$ *is sufficient*

( ) $T + T \leq D_i$ *is sufficient*

Problem 7 *(10p) Assume 10 tasks each with execution time 2 and period 100. Assume a "large" task with execution time 100 and period 101. These 11 tasks will be scheduled on a multicore platform with 10 cores.*

( ) *the task set is schedulable with global RMS*

( ) *the task set is schedulable with global EDF*

(X) *the task set is schedulable with partitioned RMS*

(X) *the task set is schedulable with partitioned EDF*

(2) *what is the minimal number of cores needed for partitioned scheduling? (put your answer in brackets)*

Problem 8 *(10p) Assume a system with a set of DRT tasks described by digraphs whose workload are* $dbf_i(t)$ *(demand bound functions). Identify the correct ones out of the following statements:*

(X) *The system is feasible on a single-core platform if for all* $t \geq 0, \sum dbf_i(t) \leq t$

(X) *If the system is feasible on a single-core platform then for all* $t \geq 0, \sum dbf_i(t) \leq t$

( ) *The system is feasible on a two-core platform if for all* $t \geq 0, \sum dbf_i(t) \leq 2t$

(X) *If the system is feasible on a two-core platform then for all* $t \geq 0, \sum dbf_i(t) \leq 2t$

Problem 9 *(10p) Assume a set of periodic tasks with periods as deadlines and utilizations as follows:* $0.95, 0.99, 0.5, 0.1, 0.6, 0.7, 0.5, 0.2, 0.4, 0.2, 0.4, 0.1$

If you partition the tasks onto a multiprocessor platform (with job splitting if necessary) and then run the respective scheduling policy, what is the least number of processors you need?

( ) *8-9 with EDF*

(X) *6-7 with EDF*

( ) *4-5 with EDF*

( ) *1-3 with EDF*

( ) *9-10 with RMS*

(X) *7-8 with RMS*

( ) *5-6 with RMS*

( ) *1-4 with RMS*

Problem 10 *(10p)*

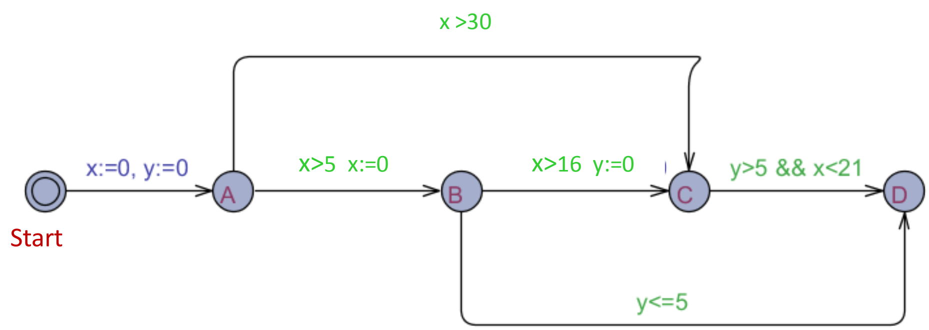

Figure 1. An Example of Timed Automaton where x and y are clocks and Start is the initial node.

Study the automaton in Figure 1. Identify the correct statements in the follwoing:

(X) *B is reachable and at B, it is always the case* $y > 5$

( ) *C is not reachable*

(X) *C is reachable and at C, it is always the case* $x > 3$

(X) *C is reachable and at C, it is possible* $x > 100$

(X) *C is reachable and at C,* y *can take any value larger than* 0

( ) *D is reachable*

( ) *D is reachable and at D, it is possible* $x < 21$ *and* $y > 5$

(X) *D is not reachable*