Tentamen vid Institutionen för Datavetenskap, Linköpings universitet

TENTAMEN TDDD07 Realtidssystem

DATUM: 20 August 2016
TID: 14-18
PLATS: TER3
ANSVARIG JOURLÄRARE: Abdeldjalil Boudjadar (0700 850 867),
Simin Nadjm-Tehrani (0702282412)

Material: English-Swedish-English dictionary
Calculator

No of assignments: 6

Total no. of points: 40

Preliminary grade limits for grades: 3, 4 and 5

3:	20 - 26 p
4:	27 - 33 p
5:	34 - 40 p

INSTRUCTIONS:

Please write your anonymous ID on each sheet of paper that you hand in. Pages should only contain answer to **one question per page** (answers to sub-questions can be on the same page). You are asked to only write on one side of each paper. Please **sort** all the sheets that you hand in, in the order of question numbers.

Make sure that **all** answers are **motivated** and supported by **clear** explanations. Figures or charts can be used to provide a clearer explanation but should be accompanied by a **textual description**. Points will not be given to answers for which the reasoning cannot be followed or that cannot be read due to bad handwriting. Wrong answers/reasoning which is embedded in partially correct ones will lead to deduction of points.

Hints: Read the question carefully to find the focus of the question. Make sure your answer is to the point and relevant for the question asked. Take the opportunity of asking questions about unclear issues during the exam session. Otherwise, whenever in doubt about the question, write down your interpretation and assumptions, and answer the question based on that interpretation. Try to dispose of your time on each question in proportion of the assignment points.

<u>Results are reported no later than 6 September 2016.</u>

Good luck!

Simin Nadjm-Tehrani

Q1: Scheduling

An electric, self-driving shuttle bus is being considered for a public transport route in the north Jutland city of Aalborg. If everything goes according to plan in testing, the completely autonomous bus could transport passengers on a 1.6 km-long route by 2018.

A self-driving bus will have thousands of parallel processes running on tens of processors connected on several real-time communication buses. Consider a dedicated processor intended for running four navigation and control processes as follows. A navigation process that decides the next chunk of the route to follow after analysing map data and any potential information on deviations from the standard route. This has a relatively slow dynamics (runs once a second) and takes 30ms to compute. A driving process that decides when to move forward and when to slow down for stopping at predetermined stops. This process runs with a frequency of 50 Hz (Hz = once a second) and has a maximum computation time of 2ms. The third process is entirely dedicated to validating obstacle notifications and issuing emergency stop commands to the driving process. This process also runs at 50 Hz and takes 3ms to compute maximally. The fourth process is a sporadic process that locks/unlocks the bus doors upon moving/stopping, and when an open/close signal overrides the current state when certain checks have passed. The minimum inter-arrival time for the process is half a second, and the maximum computation time is 50ms.

a) Consider the four processes above and determine whether the process set is schedulable with rate-monotonic scheduling, using the simplest sufficient condition that is theoretically available.

(1 point)

b) In general, how is the application developer expected to estimate the RTOS scheduler's overhead in a system so that analyzing schedulability can be applied with a margin for the overhead? Use the example to clarify the answer if you need.

(2 points)

c) Consider now that a fifth process that runs all the sensors needed for actually detecting obstacles and the data fusion for notifying that an obstacle is present. This is to be added to run on the same processor as the above four, runs twice a second, and each time has a maximum execution time of 150ms.

Discuss whether a cyclic schedule for running the set of processes consisting of the four processes under part a) above and the fifth process for sensor fusion is a good idea or not, and motivate your answer. You need not construct a full schedule here, and may ignore the scheduler overhead.

(3 points)

d) In part a) above the use of common resources by the four processes were ignored. In fact, the navigation and driving processes communicate through write/read to a common memory (variables storing the route details), and the lock/unlock process has to communicate with the driving process to know when the doors should be opened/closed. How does this knowledge affect the use of the RMS scheduling principle? What additional mechanisms are needed to guarantee the real-time execution of the four processes in part a) as well as accessing the shared resources with mutual exclusion?

(2 points)

e) A WCET estimate is a function of inputs provided to the code for which the WCET is being estimated. Describe two approaches for generating valid and relevant inputs to an application.

(2 points)

f) Describe the notion of priority inversion and explain how priority inheritance reduces its effects.

(3 points)

Q2: Dependability and predictability

a) New York Times reported on August 9th 2016: "Delta Air Lines is working to reset its operations on Tuesday morning, after a power failure grounded flights and led to cancellations and delays a day earlier. On Monday, around 1,000 of 6,000 Delta flights were canceled, the airline said. Delta said the problem was touched off by a power failure about 2:30 a.m. Eastern, shutting down computers and grounding flights for about six hours before the airline began to bring its systems back online. (As it happens, the culprit inside Delta — a failed switchgear, an item similar to a circuit-breaker box in a house — is a piece of equipment typically installed to guard against breakdowns.). Airlines were early adopters of information technology, building electronic reservation systems in the 1960s. The systems have been rebuilt over the years, but given the high volume of transactions, their data is not backed up continuously, Mr. Offutt said. He said that while airlines did have secondary systems in place — to provide power during a power failure, for example — the data was backed up only a few times each day, rather than in real time. That means that even after a malfunctioning router or power source is fixed, it can take hours to bring the systems back online.''

Use the terminology of IFIP Working Group 10.4 (from the course literature) to associate the chain of events with the fault-error-failure causal chain in this scenario, and classify the fault as permanent or intermittent.

(4 points)

b) Explain how the use of exceptions in programs fits in the classification of fault treatment approaches according to IFIP 10.4 classifications. Describe how the WCET analysis for an implemented software is affected by the existence of exceptions. Consider both exceptions that are handled by a program and those handled by the runtime system.

(3 points)

Q3: Real-time Communication

a) Describe one advantage of the TTP bus for hard real-time communication compared to using the CAN bus. Also, provide a disadvantage of the TTP bus compared to the CAN approach.

(2 points)

b) Consider a set of four messages to be transmitted on a CAN bus (where "Tx time" stands for worst case transmission time of a message on the bus). Assume that deadline is equal to period for each of the messages. Assume further that time to transmit one bit is less than 1ms. Compute the maximum response time for message m_3. Motivate the choices that you have made in the analysis, including your additional assumptions.

Message	Priority	Period (ms)	Tx time (ms)
m_1	Very High	10	1
m_2	High	5	2
m_3	Medium	30	1
m_4	Low	50	4

(3 points)

c) Consider an application in which the main fault model is the presence of Gamma radiation in the atmosphere so that arbitrary bits in the program storage areas can be flipped (from 0 to 1, or from 1 to 0). Which fault tolerance method would you consider as a means of mitigating a potential error?

(2 points)

Q4: Application design & RTOS

a) Take stand (true/false) on each of the following statements and motivate your answer!
   (1) Testing real-time systems differs from non-real-time systems as additional test code may impact the outcome of the tests.
   (2) Testing real-time systems at a later stage of development is preferable to testing at early stages of development since all the hardware and software design choices are fixed by then.
   (3) Design languages that allow non-deterministic behaviour at a higher abstraction level are better for describing real-time systems.

(3 points)

b) Real-time operating systems (RTOS) need to have "predictable latencies" for important system calls. Give three examples of system calls in the RTOS functions for which the latency needs to be demonstrably bounded.

(3 points)

Q5: Distributed systems, Quality of Service (QoS)

a) Define the notion of communication jitter in networked applications, and explain why it is relevant as a QoS indicator.

(2 points)

b) Describe the notion of skewness as defined mathematically in the context of adaptive load sharing in the paper by Xiao et al 2013.

(3 points)

c) Compare the benefits of external and internal clock synchronization algorithms to explain why you would choose one approach as opposed to the other.

(2 points)

Q6: Bonus points

a) In this question you state if you have any bonus points allocated to your attempts at bonus exercises 1, 2, and 3 during the course. Please sum up all three (if any) of your attempted exercises and write the total attained points here.

Notation for Processes

C = Worst-case execution time
B = Worst-case blocking time
D = Relative deadline
n = Number of processes
T = Period
R = Worst-case response time
J = Release jitter

Schedulability test for Rate Monotonic:

$$\sum_{i=1}^{n}\left(\frac{C_i}{T_i}\right) \le n(2^{1/n}-1)$$

Schedulability test Earliest Deadline First:

$$\sum_{i=1}^{n}\left(\frac{C_i}{T_i}\right) \le 1$$

RMS Response time analysis

$$w_i = C_i + B_i + \sum_{\forall P_j \in hp(P_i)} \left\lceil \frac{w_i + J_j}{T_j} \right\rceil C_j$$

$$R_i = w_i + J_i$$

$hp(P_i)$ is the set of processes with a higher priority than process P_i.

Timing Analysis of CSMA/CR

B = blocking time
C = transmission time of entire frame
T = period
τ_{bit} = transmission time of one bit
w = response time for the first bit of a frame to be sent
R = total response time
J = Jitter
t = Longest busy interval
$lp(m)$ = set of frames with lower priority than m.
$hp(m)$ = set of frames with higher priority than m.
$hep(m)$ = set of frames with higher or equal priority than m.
n = number of bytes in message (data field)

$$R_m = \max_{q=0..Q_m-1} (R_m(q))$$

$$R_m(q) = J_m + w_m(q) - q \cdot T_m + C_m$$

$$w_m(q) = B_m + q \cdot C_m + \sum_{\forall j \in hp(m)} \left\lceil \frac{w_m(q) + J_j + \tau_{bit}}{T_j} \right\rceil \cdot C_j$$

(with $w_m^{\,0}(q) = B_m + C_m q$)

$$Q_m = \left\lceil \frac{t_m + J_m}{T_m} \right\rceil$$

$$t_m = B_m + \sum_{j \in hep(m)} \left\lceil \frac{t_m + J_j}{T_j} \right\rceil \cdot C_j \qquad \text{(with } t_m^{\,0} = C_m \text{)}$$

$$C_m = \left(8n + 47 + \left\lfloor \frac{34 + 8n - 1}{4} \right\rfloor \right) \tau_{bit}$$

$$B_m = \max_{j \in lp(m)} (C_j)$$