Tentamen vid Institutionen för Datavetenskap, Linköpings universitet

L 11TENTAMEN TDDD07 Realtidssystem

DATUM: 25 August 2018
TID: 8-12
PLATS: TER3
ANSVARIG JOURLÄRARE: Klervie Toczé 0700895400, (Efter kl 11: Simin Nadjm-Tehrani 0702 282412)

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.

Results are reported no later than 11th September 2018.

Good luck!

Simin Nadjm-Tehrani

Q1: Scheduling

a) The Dutch company MX3D is testing the first prototype for the world's first 3D-printed bridge. It's made of a completely new type of steel, spans 40 feet, and will be installed early next year in De Wallen in Amsterdam. The team of designers and engineers that has been working on the project for years will test the bridge's integrity and add some finishing touches. The final phase will include the installation of a series of smart sensors that will monitor the bridge's health and take a slew of measurements in real-time, including air quality and weight dispersion so the bridge will actually know how many people are walking on it and how quickly they're moving. The sensor data will also be fed to a "digital twin" bridge model that will help engineers understand how the bridge is performing over time and make tweaks to future designs accordingly.

In order to rely on the outcomes of the calculations of the digital twin the sensor values must be sent over a reliable network to a server. This server is used by the engineers that monitor the operation of the bridge. Consider three processes that are running on the bridge side end of the communication network: (i) a sensor fusion process runs with a period of 120 ms, and a maximal computation time of 20ms, (ii) a command processing unit that performs a discrete control action implementing the allowed mode of operation on the bridge given the current health status as perceived in the digital twin. This is a sporadic process with a minimal interarrival time of 30 ms and a worst-case execution time of 5 ms, (iii) the communication process that manages the schedules for sending and receiving of packets from that digital twin including monitoring the status of the communication system itself. This has a period of 50 ms and a maximum execution time of 10 ms.

1. Construct a cyclic schedule for the above three processes running on the same CPU, and express any assumptions you make when constructing the schedule. What is the minor and major cycle?

   (4 points)

2. In what way would scheduling with a rate-monotonic scheduler (RMS) be better or worse in this scenario? Use the application to reason about the advantages and disadvantages of RMS.

   (3 points)

3. Assume now that the sensor fusion process is made more complex and the WCET for that process is increased to 40ms. Is the process set schedulable with RMS?

   (2 points)

b) In the part a) above the sensors involved were the minimal set of sensors and the communication between the sensors and the sensor fusion process was ignored (not considered to be significant for the overall real-time constraints). The company is now considering to add new sensors in batches of 10. These additional sensors will send their data to the sensor fusion process named under part a(i) which had 20 ms as WCET. The sensor values are to be communicated on the same communication channel that was managed by the process under a(iii) above. Assume that for each batch of 10 sensors the computation time of the sensor fusion process will increase by 5ms, and the communication process scheduling of messages will take 2 ms more. What is the maximum number of additional sensors that can be added for the system to be schedulable using Earliest Deadline First (EDF) scheduling? State at least two new assumptions for this scheduling method to be valid.

(3 points)

c) If the communication bus is treated as a resource (in addition to the CPU) what are the issues that arise if RMS scheduling were to be used? What additional piece of information would you require to be able to ensure that the process set meets its deadline?

(3 points)

Q2: Dependability and predictability

a) Model-based simulation is proposed as a means of implementing digital twins, e.g. the one that will be used in the real-time evaluation of 3D-printed bridge mentioned in Q1 above. In that question the real-time behavior of bridge-side processes were the only considered problem. If the application is considered as a distributed system (including both the digital twin simulation at the server-side and the bridge-side data analysis), which theoretical fault models should be considered when performing dependability analysis? Enumerate and describe at least 3 such generic fault models (fault modes) for distributed systems and map them on the application mentioned in question 1 (a).

(3 points)

b) Visa has said a "very rare" partial network switch failure in one of its two data centres led to the fiasco on Friday June 1st 2018 that caused millions of transactions in Europe to be declined.

The component fault within a network switch that prevented a backup switch from taking over at one of Visa's two data centres in the UK, caused interference on the two centres' mechanisms for communicating with each other, generating a backlog of work that overwhelmed the system. It took about 10 hours for normal service to resume, and affected about 1.7 million UK card holders, with around 9 per cent of transactions initiated on UK-issued cards failing to process.

The recovery took a long time. It took until 19:10 local time – the cockup having been noticed at 14:35 – to fully deactivate the system causing the transaction failures at the primary data centre, by which time the second data centre had begun processing almost all transactions normally. Normal service had resumed at both data centres by Saturday morning at 00:45.

During this time, there were two peak periods of disruption in the UK, when an average of 35 per cent of transactions failed to process: ten minutes between 15:05 and 15:15, and 50 minutes between 17:40 and 18:30.

Relate the scenario above to the attributes of dependability according to the IFIP WG10.4 terminology.
(1) Which attribute was violated?
(2) When would the breach be considered as "more frequent" or "more severe" than expected?
(3) Can you use the metrics to characterise the absence of dependability in this scenario given the information available? Motivate your answer!

(4 points)

Q3: Real-time Communication

a) The time-triggered protocol (TTP) for running a bus in a real-time communication is an alternative to the CAN protocol in applications with hard deadlines. Provide three differences between these two approaches to communication (showing where one has an advantage over the other).

(3 points)

b) The following set of messages are intended to be sent on a CAN bus.

Message #	Message priority	Max transmission time	Period
m_1	Medium	1 ms	20 ms
m_2	High	0,5 ms	10 ms
m_3	Low	3 ms	5 ms

Assume that the jitter for placing a message on the ready queue for the bus is negligible. Compute the worst-case response time for sending message m1 on this bus.

(4 points)

Q4: Application design & RTOS

a) Why is it a good idea to have standards for operating systems, such as the POSIX (IEEE 1003) standard? Expand your argument with a specific relation to real-time systems.

(2 points)

b) How is the real-time application developer expected to incorporate the knowledge about the operating system functions into the analysis of the real-time properties of an application. Provide at least three examples

(3 points)

Q5: Distributed systems, Quality of Service (QoS)

a) Take a stand (true/false) on each of the following statements and motivate your answer!

1) Adding redundancy to distributed systems as a means of fault tolerance affects the timing of the system.

2) Quality of service in a networked application is only affected by the resource allocation policies in the routing nodes of the network.

3) Vector clocks are scalable to Internet-scale distributed systems.

(3 points)

b) How can data centres enforce an energy efficiency policy when using computational resources? Is response time related to energy efficiency? Use the course literature to give two example methods to enforce energy efficiency.

(3 points)

Q6: Bonus points

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

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

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

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

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