TENTAMEN TDDD07 Realtidssystem

DATUM: 12 January 2018
TID: 8-12
PLATS: T1
ANSVARIG JOURLÄRARE: 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 31 January 2018.

Good luck!

Simin Nadjm-Tehrani

Q1: Scheduling

a) Today's agricultural machines are complex systems with many different functions that make it difficult for the driver to operate them. One way of simplifying their work is to make a number of functions automatic to enable less experienced drivers to use the machine. Research prototypes demonstrate how an autopilot can autonomously operate a harvesting machine using different sensors and regulators. Consider a simple autopilot in which three processes with the following functions run on the same processor:

- Cruise controller that regulates the number of revolutions for the wheels, where the sampling frequency is 20Hz.
- Directional regulator that decides the angle of wheels in order to follow a haystack for automatic collection of the hay, with a sampling interval of 0,2s.
- CAN controller for managing exchange of data between various subsystems.

Assume that the CAN controller is expected to manage sporadic signals with minimal inter-arrival time of 10ms. Assume further that the processes have the following maximal computation times: cruise controller 20ms, directional regulator 30ms, CAN controller 3ms, and that input jitter from sensors can be ignored.

1. How would you show whether the set of above processes is schedulable with the ”rate-monotonic” (RMS) method or not? Use the most relevant approach to analysis with RMS to answer the question and justify your choice.

   (5 points)
2. Assume that the computation taken for scheduling every arriving signal by the CAN controller increases. How does this affect the analysis above? Does the maximal response time for the above three processes become shorter or longer?

   (3 points)
3. Other than negligible utilisation for the scheduler, what other assumptions are necessary for validity of the analysis under 1 and 2 above? Give two such examples!

   (2 points)

b) Assume that a new process is to be added to the above process set under part b. The new process should manage a GPS receiver so that the harvesting machine should be able to follow a given route on the field. The process is driven by sporadic events when the driver activates a given trajectory following. Minimal inter-arrival time for this process's instances is 300ms and the worst case execution time is 60ms. Does the process set continue to be schedulable with rate-monotonic scheduling? Why/why not? Can the process set be scheduled with other scheduling algorithms? Motivate your answer!

(2 points)

c) Assume now that the GPS receiver and the cruise controller are going to share the location data in a common data structure. Assume further that the access to the common data structure is going to take a maximum of 1ms. Consider a case where the CAN controller is moved to another processor, and the cruise controller, GPS receiver, and directional regulator are to be run in a common processor. What is the maximum blocking time for each process that has to be allowed in computing the worst case response time if RMS and immediate ceiling protocol are to be combined?

(3 points)

Q2: Dependability and predictability

a) Fortune has reported an outage in the cloud services by Microsoft as follows: "Microsoft has apologized to users of its Azure cloud in Europe who could not access some services for seven hours". The episode started with an errant fire alarm. Or, as the Microsoft Azure status report put it: "During a routine periodic fire suppression system maintenance, an unexpected release of inert fire suppression agent occurred."

According to Fortune, "At that point, the data center's air handling units shut down automatically, as they are supposed to, while the conditions were assessed. Some Microsoft Azure cloud services were difficult or impossible to access between 1:27 p.m. and 8:15 p.m. local time on September 29, 2017".

Mishaps and accidents in physical systems impacting delivery of cyber services is likely to be escalating in coming years. Using the knowledge you gained from this course, how should the providers of cloud services deal with such scenarios?

(2 points)

b) IEEE Spectrum in Dec 2017 reports: "fires that killed more than 40 people in California in recent months have also jolted the state's biggest utilities, Pacific Gas & Electric (PG&E) and Southern California Edison (SCE). The utilities have had to work around the clock to keep power flowing to fire-afflicted communities, even as their equipment and policies face scrutiny as potential contributors to the deadly fires."

The devices under investigation are specifically the automatic reclosers. These "are pole-mounted circuit breakers that can quickly restore power after outages, but they can also multiply the fire risk from damaged lines" by creating sparks on each reconnection.

The phenomenon under scrutiny is further explained as follows. "…most network faults are transient. In such cases the recloser detects a power surge, momentarily interrupts electricity flow, and then automatically re-closes its contacts to restart flow down the affected line. Reclosers usually try restarting a line 2-3 times before giving up and "locking out" a line. Sometimes multiple attempts are needed to do the job… such as when high-temperature electrical arcing at the site of the fault burns away hung trees or tree limbs. Under the wrong conditions, however, such arcing and ignition can obviously spark a fire." Another utility company at San Diego uses 172 "so-called pulse reclosers that probe lines after a fault rather than simply restarting power flows. The intelligent breaker recloses its contacts for just 1-2 milliseconds and then evaluates the power that flows back. If the flow looks normal it restarts the line. And if the power signal matches the signature of a permanent fault, it locks the line out."

A suite of lawsuits filed by residents affected by the October wine country fires allege that PG&E's reclosers are to blame. And one state lawmaker has called for PG&E to be broken up if an ongoing investigation by Sacramento-based state agency CalFire finds it caused the fires.

Analyse the above scenario using the IFIP WG 10.4 fault-error-failure causal chain. Classify the use of a pulse recloser as a means of managing faults (in terms of fault removal, fault prevention, fault tolerance, or fault forecasting).

(4 points)

Q3: Real-time Communication

a) The time-triggered protocol (TTP) for running a bus in a real-time communication setting uses two hardware processors in addition to the host hardware on which the application processes run. Which are these two processors and what is the function(s) running on each?

(3 points)

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

Message #	Message ID
m_1	01110011001
m_2	11010011110
m_3	01100011100

Assume that all three messages arrive simultaneously at the CAN bus interface. Using these messages, describe how the CAN arbitration method determines which message is to be sent first. (3 points)

Q4: Application design & RTOS

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

(1) Automatic software updates while a car is running on the highway is a breach of robustness.

(2) When criticality with respect to safety are evaluated at design stage (e.g. using ASIL in ISO 26262 standard), higher emphasis is placed on exposure and controllability than on severity.

(3) The platform-independent approach to design does not bring any benefit in an area where the underlying platform is not changed frequently.

(3 points)

b) What is meant by "memory locking" in a real-time operating system?

(2 points)

c) POSIX has ten services that can be used in compliance. Name 4 of these ten services

(2 points)

d) In November 2017 it was disclosed that Intel chipsets used by billions of devices use an embedded *management engine* with a (proprietary version of) MINIX operating system. This is independent of the operating system otherwise used on a device (Windows, Linus, iOS, Android,…). The code has been shown to have security vulnerabilities and any fixes to these devices worldwide is going to be costly and slow if not impossible. Does this revelation support the arguments by the "bare machine" proponents in system software? Motivate your answer! (1 point)

Q5: Distributed systems, Quality of Service (QoS)

a) Give two methods for applying static and dynamic redundancy in hardware (one method for each), and compare (for similar machines and similar programs) the response times of programs run in each configuration during normal (fault-free) operation. You can use qualitative terms like "shorter, faster, longer, higher, …" and draw diagrams to support your arguments.

(3 points)

b) In the scheduling approach described by Xiao et al. 2013 there is a notion of skewness that defines the unevenness of utilization levels for each resource at each server. Consider a server that has three common resources for all its virtual machines: CPU, memory, and network bandwidth. Assume that at some point in time there is a 50% utilisation for the CPU, 60% utilization for the memory, and 10% utlisation for the bandwidth. What is the skewness level for this server at this point in time?

(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:

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

Schedulability test Earliest Deadline First:

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

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

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

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

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

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