

TENTAMEN TDDD07 Realtidssystem

DATUM: 12 January 2015

TID: 8-12

PLATS: T1-T2

ANSVARIG JOURLÄRARE: Simin Nadjm-Tehrani (0702 282412)

Material: English-Swedish-English dictionary
Calculator

No of assignments: 5

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. You may answer the questions in English (the course language) or Swedish.

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 the interpretation. Try to dispose of your time on each question in proportion of the assignment points.

Results are reported no later than 28th January 2015.

Good luck!

Simin Nadjm-Tehrani

Q1: Scheduling

A autonomous underwater vehicle (AUV) is a complex system with navigation and control functions that work autonomously using information gathered from a range of sensors, and include several processors. Consider an AUV processor on which the following three functions (computational processes) will be implemented on the same CPU: (1) A motion control process which will be in charge of controlling the motion in the x-y-z dimensions using motors, rudders, and propeller(s), (2) a sensor fusion component that evaluates and updates the information based on collected values from compasses, depth sensors, sidescans, sonars, and other sensors, (3) a planning and navigation process that constructs a longer range trajectory based on a given mission.

Assume that the motion control process has a period of 10ms and has a maximum computation time of 2ms, the sensor fusion process has a period of 5ms and has a WCET of 3ms, and the planning and navigation process has a period of 20 and a WCET time of 3ms.

- a) Construct a cyclic schedule for the above set of processes, justify your choices, and present your minor and major cycle.
(2 points)
- b) Consider now that new sensors are to be added to improve the location awareness of the AUV and your task is to decide how much additional CPU usage is available for managing the additional sensor fusion tasks that need to be done. Given that the cyclic scheduling policy is to be preserved, how would you answer that question? Motivate your answer by including the change to your cyclic schedule in part a) above as part of the reasoning.
(2 points)
- c) Assume now that a new CPU is being considered for the implementation which will lead to the WCET for all the processes in part a) to be reduced by half. Assume further that the additional capacity is to be used to include two new processes. One for updating the display unit with the estimated position of the nearest obstacles (with a period of 5ms), and another process for updating the current position of the AUV on the same screen (every 10ms). Assume that the five processes will be scheduled to run using a rate-monotonic scheduler (RMS). Assume further that each of the two new processes need a 1ms interval to update the screen with mutual exclusion. What is the maximum blocking time that you would use in the response time analysis for each process in the set when using RMS in combination with the immediate ceiling protocol (ICP)?
(3 points)
- d) Present a proof sketch that any process instance being run in the RMS-ICP set up will be only blocked once during its runtime.
(3 points)
- e) Discuss the problem of deadlocks in the context of cyclic scheduling.
(1 point)

Q2: Dependability and predictability

- a) The US Federal Communications Commission (FCC) has issued an investigative report on last April's 911 meltdown that affected 81 public safety answering points (PSAP) in seven states and blocked over 5,600 calls for help, saying the event was entirely

preventable. It explained that the outage occurred because of a “software coding error” at Intrado Inc.’s Colorado network center, which provides routing services for several states.

“At 11:54 p.m. on April 9, 2014, the PSAP trunk member’s counter at Intrado’s emergency call management centre exceeded its threshold and could send no more 911 calls to PSAPs using CAMA¹ trunks. Under normal operations, the PSAP trunk member assigns a unique identifier for each call that terminates using CAMA trunks. This is how Intrado has implemented the protocol commonly used to complete 911 calls over CAMA trunks, which (unlike SS7) require additional features to carry the signaling.

In this case, the trunk assignment counter reached a pre-set capacity limit to assign trunks, which meant that no additional database entries to reserve a PSAP CAMA trunk could be created, no trunk assignments for call delivery could be made for PSAPs with CAMA trunks and, therefore, no 911 calls could be completed to these PSAPs or any backup PSAP through the Englewood emergency call management centre.

When the software stopped making trunk assignments, it prevented calls being routed through the Englewood hub from reaching these PSAPs. Further, inadequate alarm management resulted in significant delays in determining the software fault and restoring 911 service to full functionality. Intrado operated a redundant hub in Miami, Florida to which 911 traffic could have been immediately rerouted, but because the malfunction was not detected promptly and mitigation actions were not efficiently developed, Intrado did not execute either an automatic or manual switchover of traffic to the Miami hub until six hours had elapsed. This switchover almost immediately restored the service.”

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 the notion of fault forecasting (according to the IFIP 10.4 terminology) and describe why it is useful for justifying dependability attributes.

(2 points)

- c) The Scania Truck company uses a 3-colour scheme (green-yellow-red) for scheduling messages sent on three CAN bus segments connected via a gateway, where each message is generated by one of 1000 functions operating in the trucks. Describe one advantage and one disadvantage of this architecture with respect to dependability.

(1 point)

Q3: Real-time Communication

- a) The time-triggered architecture TTA supports redundancy in nodes and the transmission infrastructure as a means of fault tolerance. What is meant by “replica determinism” in the context of this architecture?

(3 points)

- b) The following table presents 5 messages to be scheduled on a CAN bus (where “Tx time” stands for worst case transmission time of a message frame on the bus). Assume that deadline is equal to period for each of the messages. Assume further that time to transmit

¹ CAMA: Centralised Automatic Message Accounting

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	5	1
m_2	High	50	2
m_3	Medium	20	2
m_4	Low	10	5
m_5	Very Low	20	2

(3 points)

Q4: Application design & RTOS

a) Consider an air traffic control system where the goal of distributed traffic controller nodes is to collectively keep flying aircrafts within safety distances from each other. Decide which of the following properties is a functional property and which is an extra-functional property. Motivate your answer!

1. The controller node should accept flight plans from each running aircraft and monitor the current position of the aircraft on a common screen.
2. The amount of flight data regularly transmitted by each aircraft within a controller's coverage range should be limited to X kb.
3. The accuracy of the recorded position of an aircraft on each controller screen should be within Z km of the actual position.
4. If a flight data for one aircraft is lost the controller node should be able to retrieve the information on a second redundant poll or through contacting the other aircrafts in vicinity.

(4 points)

b) Software engineering methods and processes are applied to both real-time and non-real-time systems for better quality assurance.

1. Why is the discovery of design faults in real-time systems at an early stage harder than discovering them at a late development stage?
2. When can misconceptions in requirements for a system be discovered? In any development phase?

(2 points)

c) Why is the organisation of memory as a pool of fixed sized units a preferred form of memory allocation within real-time operating systems?

(2 points)

Q5: Distributed systems, Quality of Service (QoS)

- a) Take a stand on the following statements, motivating your answer by referring to the relevant terminology discussed in the course literature:
- 1) Live video broadcast and video streaming applications have different packet loss requirements.
 - 2) A system that uses Lamport-MellierSmith algorithm for clock synchronisation does not provide adequate synchronisation in presence of two-face clocks.
 - 3) Buffer management in Internet middle nodes is a means of prioritizing packets that have short deadlines.
- (3 points)
- b) Present the rules that govern the operation of vector clocks in distributed systems.
- (3 points)
- c) Zhen et al. have suggested a method for managing overloads in physical machines at data centres by migrating virtual machines from overloaded ones to less overloaded ones. They use notions of hot, cold, skewness, etc with respect to the current load of a machine. How do they ensure that the spreading of the load on too many physical machines does not lead to excessive energy consumption?
- (2 points)



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)

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

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

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

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

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

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

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

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

