

Theory Exercises for TDDD07 Real-time Systems

Massimiliano Raciti
Jordi Cucurull
Simin Nadjm-Tehrani

Real-Time Systems Laboratory
Department of Computer and Information Science
Linköping University, Sweden

November 2019

Copyright © 2019 Simin Nadjm-Tehrani

1. Scheduling and Resource Sharing

Q1.1:

Construct a cyclic schedule for the following process set and present the minor and major cycles respectively. Argue the result and analyse the case assuming that no output jitter is tolerated.

<i>Task</i>	<i>Period (ms)</i>	<i>WCET (ms)</i>
A	50	5
B	30	10
C	20	5

Q1.2:

Assume that the following task set is to be analysed for its real-time properties in a new generation of the Mars Rover: (1) the *trajectory following task* for finding the way to a well-defined physical coordinate, (2) the *sensor & measurement task* that configures and reads various sensor's values, (3) a *disk storage task* that is responsible for periodically writing the buffered measured values on a log on persistent memory, (4) a *communication task* for contacting the ground station and sending cached data. The following table summarises the task set's periods and worst-case execution times (WCET) in milliseconds.

<i>Task</i>	<i>Period</i>	<i>WCET</i>
Trajectory follower	1500	100
Sensor & measurement	500	100
Disk storage	3000	1000
Ground communication	500	200

a) Assume that the tasks can be scheduled as independent tasks and the task set is to be scheduled using a cyclic executive. Provide such a cyclic schedule, and in particular, the minor and major cycles for the schedule. Comment on the advantages and disadvantages of a cyclic schedule.

b) Assume that to lower the processor utilization the communication overhead is reduced by using a dedicated channel that is used every 500 ms with 100 ms instead of 200ms as WCET. Furthermore, the sensor configuration frequency can be reduced to once every second. Assume that the trajectory following process is the only process with strict jitter requirements (no output jitter is acceptable). Explain whether the task set is schedulable with rate-monotonic scheduling (RMS). Include all the assumptions in your analysis. Comment on the satisfaction of the jitter requirements.

c) Assume now that the task that records sensor values in the buffer and the Disk storage task (that reads from that buffer) are in fact not independent from each other. The analysis of the operation of the system shows that under RMS scheduling, when the arrival of sensor values are occasionally delayed, the disk storage task locks the buffer memory. Thereby the sensor

recording task is at times waiting for the slow disk storage task, as well as the trajectory follower task that can preempt the disk storage task. What is this phenomenon called? How can the impact of sharing the buffer be accounted for and minimized, and the overall task set real-time response determined?

Q1.3:

- 1) A space stationed telescope that has been sent to take pictures from far away stars has three different functions that are implemented on the same CPU: A position stabilizer that uses a gyro to stabilize the movement around own axes; A star-follower process that tries to fix the orientation so that sharp pictures can be taken; An energy management system that deals with charging of batteries and adaptation to sun hours and angles. Assume that the position stabilizer is run every 50 ms and takes maximally 5ms to run, that the energy manager is run every 30 ms with a maximal computation time of 10ms, and the star-follower process takes a picture from a given star every 20 ms and takes maximally 5ms to fix orientation with the help of the picture.
Construct a cyclic schedule for the above process set and present the minor and major cycles respectively.
- 2) Assume now that the process set is scheduled using rate monotonic scheduling (RMS) and we want to increase the functions to be implemented on the same CPU with 2 new ones. A picture storage process that shares common (data) memory with the star-follower process, and a communication process that shares position data with the position stabilizer (i.e. two pairs of processes sharing one memory resource each). The length of execution in critical sections for all processes is provided in the following table. Assume that we use immediate ceiling protocol for managing access to common resources, and that picture storage process has a periodicity of 50ms and communication process has a periodicity of 100ms. Compute the blocking time (one component in the response time) for the processes that share common resources.

Process	WCET for running while holding a resource
Position stabilizer	1 ms
Star follower	2 ms
Picture storage	2 ms
Communication	1 ms

- 3) Consider the following periodic processes and their parameters (in *ms*):

Process	WCET	Period	Deadline
P1	3	10	10
P2	4	16	8
P3	2	12	12

Given the above parameters, can utilisation based tests be used for guaranteeing schedulability with RMS? Motivate your answer!

Q1.4:

A system for quality assurance and packing of chocolate bars has been designed using three spectrometers one for each type of chocolate (dark, milk and white chocolate). A chocolate block arrives on a conveyor belt and when reaching the end point it will be subject to sensing of its colour and structure. The block then moves on an adjacent conveyor belt and depending on the classification of a quality test will be discarded or packed into the right packaging. There are five computational processes involved, which run on the same CPU. Two servo motor controllers (one for each conveyor belt) running with the period of 4ms, one sensing process that also analyses the outcome of the spectrometer measurements (same process for all three spectrometers and their analysis), one process for both making the decision on the packaging component and actuating it, and one logging component that collects statistics of packed blocks and sends the data away for stock management. Assume that the sensor analysis process is run every 10ms and takes a maximum of 2 ms to run. Assume further that the actuator process is run at the same rate as the sensor analysis process and takes a maximum of 1ms to complete, and the logging process has a worst case execution time of 2ms. The servo steering processes take a maximum of 0,5ms each to compute.

- 1) How often can the logging process be run if the process set is to be scheduled with the earliest deadline first scheduling technique? Give one additional assumption necessary for your analysis to hold.
- 2) Assume now that the logging process will be run with a periodicity of 8ms. Construct a cyclic schedule for the above process set and present the minor and major cycles. You are expected not to increase the processor utilisation, but jitter can be accepted.
- 3) Now assume that the process set is scheduled using rate monotonic scheduling (RMS) and we intend to account for the access to shared resources in the system using the immediate ceiling protocol. The sensor analysis process outcome is output to a shared memory which the Logging process reads from when reporting on scrapping/packaging outcome. The logging process sends the collected statistics via a local area network. Consider the same parameter set as in part a) above and assume that processes with the same period are run as one process. Include the period for the logging process as 8ms.

Process	Worst case locking to read/write the common data structure	Increase in the WCET for the logging process due to access to the communication network
Sensor analysis & actuator	0,5 ms	
Logging	1 ms	
Logging		1ms

Compute the maximum response time for each process in the rate-monotonic setting to show whether the process set is schedulable.

- 4) What is meant by a sporadic task in a real-time system? Give two examples of such tasks.

Q1.5:

In modern train systems many safety systems are implemented in software. For example, functions for maintaining a given speed and a system to avoid passing a red light signal. Also in order to stop the wheels from skidding there are regulators for each pair of wheels so that their speed is kept under control separately. Consider the set of three processes that implement the named functions:

- The cruise controller to keep a given a speed, running with a period of 10 ms and a maximal execution time of 3ms.
- The group of wheel-pair regulators all having a period of 5 ms and collectively take a maximum of 1ms to run.
- The monitoring of stop light signal and automatic braking will have a maximum running time of 1 ms.

- 1) How often should the monitoring process be scheduled if the processor utilisation is desired to be at most 70%?
- 2) Assume that a new process for controlling the power usage by the wheel motors is added. This process will need to run with a period of 30 ms and takes a maximum of 2ms to run. Assume further that the requirement of 70% utilisation is dropped, but the jitter for the cruise control process is to be minimised. Let the minimum inter-arrival time for activating the monitoring/stopping process be 5 ms. Construct a cyclic schedule for the four processes, including its major/minor cycles.
- 3) The following three processes share two resources R and S with the given worst case locking times:

Process	Shared Resource	Worst case locking for shared resource	Priority (highest is 4)
P1	R	0.5 ms	2
P2	R	1.0 ms	3
P2	S	2.0ms	3
P3	S	1.0ms	1
P4	-	-	4

Assume that the processes are to be analysed for their response time according to rate monotonic scheduling and using the immediate ceiling protocol. Find the blocking term for each process.

- 4) Take a stand on the following statement (true/false) and motivate your answer!
“If a task set includes a task that is neither periodic nor sporadic, then the task set cannot be analysed for schedulability in the rate-monotonic context .”

Q1.6:

Consider a set of three processes, P1, P2, and P3, with the following instances of arrival time and following execution patterns. The processes share two common resources **a** and **b**. The patterns describe a process executing without any resource during one time unit using the notation **E**. When the resource **a** (respectively **b**) is used for one time unit it is listed as one instance of **a** (respectively **b**). Assume that fixed priority scheduling is used combined with immediate ceiling protocol, and highest priority is denoted by 3. Show the actual execution of the three process instances by drawing charts over time lines. Use the notation “P” for a preempted process, “B” for a blocked process in your solution chart. You can use the printed chart at the end of the exam for this purpose.

Process	Pattern	Priority	Release time
P1	EaEbE	2	2
P2	EaabbE	1	0
P3	EbbaaE	3	7

Q1.7:

Consider the following processes where the Immediate Priority Ceiling protocol is used for resource management and rate-monotonic is used for scheduling policy. Assume that the time denoted by “locking time” denotes the maximum time that a process uses a given resource (the WCET of the critical section with that resource in use).

Process	Period	WCET
A	100	10
B	40	12
C	50	6
D	1000	50

Process	Resource	Locking time
A	S1	1
B	S1	2
C	S2	1
D	S2	4

- Compute the ceiling for each resource
- Compute the blocking time B for every process.

Q1.8:

Consider the following processes where the Immediate Ceiling Priority protocol is used for resource management and RM is used for scheduling the processes. Locking time denotes how long the process locks the respective semaphore. Compute the maximum blocking time and the response time for each process.

Process	Period	WCET
A	16	4
B	10	3
C	8	2


Process	Resource	Locking time
A	S1	1
B	S1	1
B	S2	1
C	S1	1
C	S2	2
C	S3	1

2. Distributed systems

Q2.1:


For the given events in the following scenario:

- Identify the Logical and Vector clock timestamps;
- Can you identify at least a couple of “concurrent” events? Motivate your answer according to the definition given for vector clock.



Q2.2:

Identify two events x and y in the following scenario such that $LC(x) < LC(y)$ but it is not the case that $x \rightarrow y$ (x happens before y).




Q2.3:

Give an example of a distributed real-time system with hard deadlines, and another example with soft deadlines.

Q2.4:

- Consider 4 nodes in a distributed system in which the local clocks are to be synchronised using the Lamport/Melliar-Smith algorithm (i.e. upon exchange of local clock values, a value with a difference higher than δ is replaced by own clock and then the reported values are averaged to produce the value that replaces the own clock value). Consider a case where none of the clocks is a two-face clock (i.e. acts Byzantinely). Explain if the synchronisation requirement is met at node A if node A reads its own clock value t , and receives 3 values from nodes B, C, and D where two out of the readings differ from t by exactly δ .
- Consider the following three nodes in a client-server architecture which we model using the asynchronous model of distributed systems. Assume that a request by client 2 reaches the server before a request by client 1, followed by the server responding to the queries in FIFO order. Derive the vector clocks of five of the events in the system including at least a send and a receive event, and mark two concurrent events included among the five events.



Q2.5:

Consider the following two terms used in quality of service (QoS) requirements descriptions, and identify for each term whether it is an application level description or an enforcement level indication: loss ratio, video quality.

Q2.6:

Consider the following set of messages that are intended to be sent over a CAN network. Message 1 has an ID = 000 and sends its frames periodically every 5 ms. Message 2 has the ID = 001 to be transmitted every 7ms. Message 3 has ID = 010 sent with a periodicity of every 10ms. Assume that every frame is maximally 135 bits long and the transmission time per bit (t_{bit}) is 1×10^{-6} s.

- Describe how the nodes are going to agree on the order that the frames are going to be sent when all nodes send simultaneously. Exemplify with the nodes n1, n2 and n3.
- Compute the response time for the lowest priority message assuming that none of the nodes sending these messages create a jitter.
- How is the response time of the lowest priority message affected if the three messages are expecting to suffer jitter? Assume maximum jitter values of 5, 5 and 8 ms for messages 1..3 respectively. Recompute the response time for the lowest priority message.

Q2.7:

The following messages take one ms each for being sent over a CAN-bus, and have the following periods and maximal jitter at release time:

Message	period (ms)	Jitter
m1 (high priority)	20	1
m2 (middle priority)	10	2
m3 (low priority)	5	0

Compute the maximum response time for all the messages.

Q2.8:

Consider a set of messages to be scheduled on a CAN bus with the following priorities and parameters. Assume the deadline for a message is equal to its period, and that all messages have a maximum transmission time of 1 ms. Assume further that transmission time for one bit is less than 1 ms.

Message	period (ms)	Jitter
m1 (high priority)	30	5
m2 (middle priority)	15	0
m3 (low priority)	5	0

Compute the worst case response time for the message with the lowest priority.

Q2.9:

A distributed real-time system's behaviour is strongly affected by both scheduling of CPU as a resource in the nodes, and allocation of bandwidth in the communication channel as a resource. Take a stand on the following statements and declare whether the statement is true or false. Motivate your answer!

- a) The priorities for messages in a CAN network must be the same fixed priorities used in the priority-based scheduling of the processes in the nodes that are connected to the bus.
- b) If one node that is connected to a TTP bus crashes, this can be detected easier by the other nodes than if the system would use a CAN bus.
- c) If a process exceeds its assumed worst case execution time (WCET) at some point in time, it is stopped from sending its final output on a TTP bus.
- d) On a CAN bus a high priority message can only be delayed (blocked) once by messages of lower priority.

3. Dependability

Q3.1:

In the following scenarios identify the causal chain from fault to error and failure:

- The transmission of a world cup football game in all countries stopped for 35 minutes due to a software misconfiguration that crashed the central live streaming server in the head quarters of the media company which is in charge of relaying the game.
- Broadband services were experienced as being slow several hours after the electricity outage was created due to lightning.
- In Java 1.6.0_21, the company field was changed from 'Sun Microsystems, Inc' to 'Oracle.' Some applications depend on that field to identify the virtual machine. For example, all Eclipse versions since 3.3 including the recent Helios release (2010) have been reported to crash with an OutOfMemoryError due to this change.

Q3.2:

- a) The following messages take one ms each for being sent over a CAN-bus, and have the following periods and maximal jitter at release time:

Message	period (ms)	Jitter
m1 (high priority)	20	1
m2 (middle priority)	10	2
m3 (low priority)	5	0

Assume that the time taken to transmit one bit is less than 1 ms. Compute the maximum response time for the middle priority message.

- b) Where is the MEDL (message descriptor list) of a TTA bus stored and what is its role?

Q3.3:

- a) Give an example technique that helps to discover early design faults in embedded real-time systems.
- b) What is meant by platform independence, and why is it a good property in modelling languages for real-time systems?

Q3.4:

- a) Take a stand on the following propositions (true or false), and motivate your answer:
 - 1) Application program timing faults can never be detected by the run-time environment.
 - 2) Simulation of the design of a program can be used to study run-time fault tolerance properties.
 - 3) Voting systems (e.g. triple modular redundancy) cannot be used to tolerate the same software design fault appearing in every replica.
- b) Describe four drawbacks with checking real-time systems for correctness and timeliness only on the platform for which they are intended, as opposed to being tested in a platform-independent design phase first.

Q3.5:

- a) Consider a collision avoidance component in a flight control system. Decide which of the following properties is a functional property and which is an extra-functional property (also sometime called a non-functional property):
 - 1) When another aircraft is within X meters of own aircraft on the same altitude, the marking on both pilot display screens should change colour within Y milliseconds.
 - 2) For two aircrafts on the same altitude if own aircraft is instructed to rise, the other aircraft has agreed to descend (go down).
 - 3) If the altitude measurement delivered by the altitude-metering system is inaccurate then the collision avoidance system should switch to an alternative source for altitude value.
- b) Name two types of languages that can be used to describe a system at a higher level of abstraction than the platform-dependent programming language. For each language describe one benefit that the language brings to the system development process.

Q3.6:

- a) Explain whether production defects in microchips are an example of faults, errors or failures.
- b) Describe the relation between “degraded mode” and “system failure”. You may use an example to explain whether these terms are synonyms or have differences.

Q3.7:

- a) Explain the notion of graceful degradation and give one example of it in a real application setting.
- b) Take a stand on the following propositions (true or false), and motivate your answer:
 - 1) Redundancy in hardware through triple modular redundancy does not increase the response time of an application compared to a non-replicated solution.
 - 2) TCP employs redundancy in data when a message is retransmitted to achieve reliable communication.
 - 3) TCP employs redundancy in time when a message is retransmitted to achieve reliable communication.

Q3.8:

Identify the causal chain of fault-error-failure in the following scenario, including whether the fault was permanent, transient or intermittent.

On 7th October a newspaper in Colorado reported that real-time alerts had ceased for a number of persons being electronically monitored instead of being in jail. BI Inc., a company that provides electronic monitoring for several nationwide agencies, experienced a problem with one of its offender monitoring servers at 7:29 a.m., temporarily disabling the server's notification system and delaying violation notifications to customers. The technical glitch happened when one of BI's servers exceeded its threshold of 2.1 billion records. "The offenders and suspects on the monitoring system did not know their devices were down at the time, the company officials said, and there were no major problems reported as a result of the technical failure."

Q3.9:

Identify the causal chain of fault-error-failure in the following scenario.

The University College London Hospitals Trust (a consortium of hospitals in London) was on 22nd February 2011 forced to halt a number of services, including the cancellation of 50 per cent of its operations, due to a faulty network switch. The faulty switch left computers across the connected London hospitals unable to access various systems such as the trust's patient administration system and its patient records software CareCast.

4. Mixed questions

Q4.1:

- a) Explain the notion of reliability, and provide a metric (= means of measuring) that can be used for measuring this attribute in a system.

- b) Take a stand on the following propositions (true or false), and motivate your answer:
 - 1) A real-time operating system may tolerate program design faults by redundancy in time, i.e. a scheduling algorithm that runs a process again if the first run of the process fails.
 - 2) Simulation of the design of a program can be used to eliminate all requirements faults.
 - 3) Voting systems (e.g. triple modular redundancy) can be used to tolerate transient faults but not permanent faults.
 - 4) A method for fault forecasting is to build in adaptive load control in the design of the system.

- c) Describe four functions in a real-time operating system that need to be implemented in a different manner from an ordinary operating system.