

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich Institut für Technische Informatik und Kommunikationsnetze Computer Engineering and Networks Laboratory



Exam Spring 2015 Embedded Systems

Prof. L. Thiele

Note:

The given solution is only a proposal. For correctness, completeness, or understandability, no responsability is taken.

Spring 2015	Embedded Systems	Page 1

Task 1 : Real-Time Scheduling

1.1: Periodic Task Scheduling

A harmonic periodic task set is a periodic task set where all periods of the tasks are pairwise multiples or divisors of each other. Table 1 shows an example of a harmonic periodic task set.

	Computation Time	Period = Deadline
Task 1	1	2
Task 2	1	4
Task 3	2	8

Table 1: Harmonic periodic task set.

(a) (5 points) Schedule all tasks in Table 1 using Rate Monotonic (RM) scheduling. Draw the RM schedule in Figure 1. Do all tasks meet their deadlines?



Figure 1: RM Schedule.

(maximal 40 points)

(maximal 20 points)

Spring 2015	Embedded	Systems	Page 2

(b) (5 points) Schedule all tasks in Table 1 using Earliest Deadline First (EDF). In case of equal deadlines, a task with lower *index* has higher priority. Draw the EDF schedule in Figure 2. Is the EDF schedule the same as the RM schedule?



Figure 2: EDF Schedule.

swiss reactar institute of rea		inputer Engineering	
Spring 2015	Embedded	Systems	Page 3

(c) (10 points) Let C_i and T_i denote the computation time and the period of the periodic task *i*, respectively. Assume that the relative deadline of each task is equal to its period. Prove the following statement:

If a periodic task set with n tasks is *harmonic* (i.e., for any two tasks i and j, if $T_i \ge T_j$ then $\frac{T_i}{T_j}$ is a positive integer) and it holds that $\sum_{1 \le i \le n} \frac{C_i}{T_i} \le 1$, then all tasks of the task set meet their deadline using the RM scheduling policy.

(Hint: Start by using the necessary and sufficient schedulability test of the RM scheduling algorithm for the lowest priority task)



1.2: Aperiodic Task Scheduling

(maximal 7 points)

An application consists of periodic tasks and three aperiodic tasks. The total utilization of periodic tasks is 0.6. Parameters of aperiodic tasks are given in Table 2.

	Arrival Time	Computation Time
J_1	57	10.5
J_2	60	37.2
J_3	180	65.4

Table 2: Aperiodic Tasks.

Assume that the application is scheduled using EDF and all aperiodic tasks are scheduled using a Total Bandwidth Server (TBS).

(a) (2 points) What is the maximum utilization of the TBS?

(b) (5 points) Determine the *worst-case* finish time of aperiodic tasks J_1 , J_2 , J_3 .

(maximal 13 points)

1.3: Resource Sharing

Three tasks share a single resource protected by a critical section. The execution patterns for the tasks are shown in Figure 3 and their release times are given in Table 3. The critical sections are shaded in grey. P_1 , P_2 , and P_3 represent the priorities of Task 1, Task 2, and Task 3, respectively. Lower index tasks have higher priority; i.e., $P_1 > P_2 > P_3$.



Figure 3:	Execution	Patterns.

	Release Time
Task 1	5
Task 2	3
Task 3	0

Table 3: Release Times.

(a) (5 points) Use the Priority Inheritance Protocol to schedule all tasks. Draw your schedule in Figure 4.



(b) (3 points) Draw the priority level of Task 3 with respect to time in Figure 5.





Institute of Technology Zurich Computer Engineering and Networks Laboratory Embedded Systems Page 6

(c) (5 points) Explain with an example how a deadlock can occur when the Priority Inheritance Protocol is used. (Note: This question is independent of questions 1.3a and 1.3b.)

Spring 2015	Embedded Systems	Page 7

Task 2 : Components and Communication

2.1: Components

- (a) (1 point) Digital Signal Processors are especially suited for ... (check one box). Control Dominated Systems Data Dominated Systems
- (b) (1 point) Briefly outline (in about two sentences) the main characteristics of Control Dominated Systems and of Data Dominated Systems.
- (c) (1 point) State in two sentences, what are the main characteristics of FPGAs and ASICs? What are the respective advantages and disadvantages for implementing an embedded application?

2.2: CSMA/CR

(maximal 6 points)

Consider a wired communication system with Carrier Sense Multiple Access Collision Resolution (CSMA/CR). The system is *dominant high*, meaning that a device sending a high (1)has higher priority than a device sending a low (0). The communication packets consist of [Device ID | Data] (most significant bit sent first); there are four devices A, B, C, and D, with device IDs A = 12, B = 7, C = 13, and D = 14.

(a) (4 points) Sort the devices A, B, C, and D in descending order of priority. (Hint: start with the device that has the highest priority in accessing the communication system.)

T

(maximal 45 points)

(maximal 3 points)

	0,		0	0	,
Spring 2015	Emb	edded Systems			Page 8

(b) (2 points) In the system described above, if the highest priority device is continuously sending packets, the other devices are not be able to send any packets. This effect is called *starvation*. Can starvation also happen when using CSMA/CD? Provide a short motivation (about two sentences) for your answer.

2.3: Token Ring

(maximal 18 points)

Consider the token ring with four clients of Figure 6. The maximum transmission rate is $16 M b p s^{\dagger}$, communication proceeds in rounds, each round can last **at most** 5 m s, and there is one token. Within each round, the client that has the token performs the following steps:

- 1. If there is data ready, send a packet consisting of a $4 k b^{\ddagger}$ header, followed by data.
- 2. If a packet was sent, wait until an ACK of $4\,kb^{\ddagger}$ is received.
- 3. Pass on the $8 k b^{\ddagger}$ token to the next client.

Step 3 always takes place once per round; steps 1 and 2 only happen, once per round, if the client has data to send. If data is too big to be sent in one round, it has to be partitioned



Figure 6: Token ring status at time 0; the four clients need to send the specified data amounts^{\dagger}.

into multiple rounds. All clients receive messages without any delay and the processing time for tokens, packets, and ACKs is negligible. At time 0, client A has the token; at the end of

 $^{^{\}dagger}1\,Mbps = 10^{6}\,bps$ (bits per second), $1\,kbps = 10^{3}\,bps$

 $^{{}^{\}ddagger}1 Mb = 10^6 b$ (bits), $1 kb = 10^3 b$

Swiss rederar institute of feel	mology zurich	computer Engineering a	na Networks Laboratory
Spring 2015	Embe	edded Systems	Page 9

each round, the token is passed to the next client on the network, in a clockwise direction. At time 0, every client has data ready to send; Figure 6 shows the data size.

(a) (12 points) **Calculate** the total time needed until all clients have sent all their packets and the token is in possession of Client A again.

(b) (6 points) **Calculate** the aggregated data throughput of the network (in Mbps) for this scenario.

2.4: Bluetooth

(maximal 18 points)

A Bluetooth connection is defined with the following characteristics:

- Slaves can only send packets directly after they have received a packet from the master.
- One slot has a duration of $625 \, \mu s$.
- Each packet contains a 126 *bit* header, a 24 *bit* CRC, and the **maximum** integer number of bytes[‡] of payload (user data) to fill up the rest of the packet.
- It is possible to send packets with the length of 1 slot, 3 slots or 5 slots.
- Only a **maximum** of $366 \, \mu s$ of the first slot can be used for packet transmission.
- The data rate is $1 \ Mbps^{\dagger}$
- (a) (10 points) Determine the channel throughput (only transmitted user data, excluding header, CRC, etc.) for the packet lengths given in Table 4 and the Bluetooth connection defined above. Please show all your calculations.

Packet Length in Slots		Throughput in kbps	(1 kbps = 1000 bps)
To Slave	To Master	To Slave	To Master
1	1		
1	3		
1	5		

Table 4: Packet lengths and channel throughput calculation.

 $^{^{\}ddagger}1 \, byte = 8 \, bits$

[†]1 $Mbps = 10^6$ bits per second

wiss Federal Institute of Technology Zurich	Computer Engineering and Networks Lab	oratory
Spring 2015	Embedded Systems P	'age 11

(b) (2 points) Why is the first $625\,\mu s$ slot not entirely used for packet transmission? Answer in one short sentence.

(c) (6 points) Which two connection types are defined in the Bluetooth Standard, and what are those types used for? Please state in two short sentences.



	0,		-	0	•	
Spring 2015		Embedded Sys	stems			Page 12

Task 3 : Low Power Design

3.1: Energy Minimization

Consider a heterogeneous dual-core platform consisting of processors π_1 and π_2 . A task of 10^7 clock cycles is going to be executed on this platform. The workload can be freely partitioned between the two processors to run in parallel. In addition, we make the following assumptions:

- The platform has two operation modes: active and sleep.
- In active mode, the frequency, dynamic and leakage power consumptions of π_i are denoted as f_i , P_{di} and P_{li} , respectively. We assume they are constant: $f_1 = 1$ MHz, $P_{d1} = 3 \text{ mW}$, $P_{l1} = 0 \text{ mW}$, $f_2 = 4 \text{ MHz}$, $P_{d2} = 10 \text{ mW}$ and $P_{l2} = 6 \text{ mW}$. We assume dynamic power is only dissipated when a processor is processing tasks, while leakage power is always dissipated in active mode.
- In sleep mode, the processor frequency is zero and no power is dissipated, neither dynamic nor leakage. The switching overheads in terms of time and energy can be neglected.
- (a) (10 Points) Assume that a processor can switch to sleep mode whenever it is idle, i.e., when it is not processing tasks. What is the optimal (i.e., minimal) energy consumption of the system?

(b) (8 Points) Assume that the platform can only switch to sleep mode if *both* processors are idle. What is the optimal (i.e., minimal) energy consumption of the system now? (Hint: assume that x and y number of clock cycles are assigned to π_1 and π_2 , respectively. You can formulate and solve the problem with respect to x and y, or find directly the relation between x and y.)



(maximal 35 points)

(maximal 18 points)



Spring 2015	Embedded System	ns Page 13

3.2: Energy Harversting

(maximal 17 points)

Consider a processor with negligible leakage power dissipation and dynamic power dissipation specified as $P_{dynamic} = \left(\frac{f}{MHz}\right)^3 mW$, where f is the frequency in Hz. The processor is put into a zero power state whenever it is idle. The set of hard real-time tasks of Table 5 needs to be executed on the processor:

Tasks	T_1	T_2	T_3
Period (ms)	6	4	12
Relative Deadline (ms)	6	4	12
Cycles (x 10^3)	2	1	2

Table 5: Characteristics of the hard real-time tasks to be executed.

All tasks initially arrive at time zero. The system has a battery with initial energy C microjoule (μ J) and it is replenished by a constant power source of A microjoule per millisecond (μ J/ms).

(a) (9 Points) Assume $C = 6 \ \mu$ J, $A = 0.5 \ \frac{\mu J}{ms}$ and $f = 1 \ MHz$. Apply EDF scheduling and draw in Figure 7 the battery energy during the time interval $[0 \ ms, 12 \ ms]$.



Figure 7: Battery energy diagram.

	87		0	0	·····,
Spring 2015	Embeda	led Systems			Page 14

(b) (8 Points) Assume the battery does not run out of charge. Prove or disprove the following statement:

To have the maximum possible battery energy after each hyper-period (12 ms), all tasks should run at the same frequency.

(Hint: providing main arguments or formal proof are both accepted.)

(maximal 12 points)

Spring 2015	Embedded Systems	Page 15
Task 4 : Architecture	e Synthesis	(maximal 60 points)

4.1: Architecture Synthesis Fundamentals

Mark the following statements as *true* or *false* and provide a brief explanation (1 sentence).

• (1 point) The different paths in a dependence graph represent branches in the control flow of a program.

		\Box False
Explanation: _		
(1 point) In a fire) if there is	marked graph, a node w a token on at least one □ True	vith more than one input edge is activated (it car e of its input edges. □ False
Explanation: _		
(1 point) As-S the latency of	oon-As-Possible (ASAP an operation set, under □ True	 P) is an <i>exact</i> scheduling algorithm for minimizing r no resource constraints. □ False
Explanation: _		
(1 point) Und anteed minima	er resource constraints, al latency.	the LIST algorithm returns a schedule with guar-
Explanation: _	□ True	□ False
(2		
with limited re	e LIST scheduling algori esources.	thm can be applied to pipelined implementations
- :	□ True	□ False
Explanation: _		

idgenössische Technisch	Institut ne Hochschule Zürich und Kor f Technology Zurich Compu	für Technische Informatik mmunikationsnetze ter Engineering and Networks Laboratory
Spring 2015	Embedded System	ns Page 16
 (2 points) When such that their s 	loop folding is enabled in a functio starting and finishing times are alw] True	nal pipeline, operations are schedulec ways in the <i>same</i> physical iteration.
Explanation:		
 (2 points) Consider a unit of each read allocation that n 	der the sequence graph of Fig.8(a) source type (ADD or MUL) costs (minimizes <i>both</i> the schedule latence True	and the resource graph of Fig.8(b). In D.5 SFr, there is exactly <i>one</i> resource by and implementation cost.
Explanation:		
Explanation:	True False	
Explanation:	True False	v ₀ nop
Explanation: $v_0 \xrightarrow{nop}$ $v_1 \xrightarrow{v_2}$ $(+) \xrightarrow{v_2}$ $(+) \xrightarrow{v_4}$	True V_1 V_1 V_1 V_3 V_1 V_3 V_5 V_5 V_5	$\begin{array}{c} & & & & \\ & & & & \\ & & & & \\ & & & & $
Explanation: v_0 v_1 v_2 v_1 v_2 v_3 v_4 v_4 v_5 v_6 v_6	True \square False V_1 V_1 V_1 V_3 V_1 V_3 V_1 V_3 V_1 V_3 V_1 V_3 V_1 V_3 V_1 V_1 V_3 V_1 V_1 V_3 V_1 V_1 V_3 V_1 V_3 V_1 V_2 V_1 V_2 V_1 V_2 V_1 V_2 V_1 V_2 V_1 V_2 V_3 V_1 V_2 V_3 V_1 V_2 V_3 V_1 V_2 V_3 V_1 V_3 V_2 V_3 V_1 V_2 V_3 V_1 V_2 V_3 V_2 V_3 V_2 V_3 V_4 V_1 V_2 V_3 V_4 V_2 V_4	$\begin{array}{c} & \overbrace{v_0}^{nop} \\ & \overbrace{v_2}^{-2} \\ & \overbrace{v_3}^{-1} \\ & \overbrace{v_1}^{-1} \\ & \overbrace{v_1}^{-1} \\ & \overbrace{v_1}^{-1} \\ & \overbrace{v_1}^{-1} \\ & \overbrace{v_5}^{-1} \\ & v_$
Explanation: v_0 v_0 v_1 v_2 v_3 v_4 \downarrow v_4 \downarrow v_4 \downarrow v_5 v_6 (nop) a) Sequence	True \square False V_1 $+$ V_1 $+$ V_3 $ 1$ $+$ ADD V_5 $+$ $+$ V_2 $+$ 1 $+$ WUL (b) Resource Graph.	(c) Weighted Constraint Graph.

Figure 8: Architecture Synthesis Fundamentals

und Kommunikationsnetze Computer Engineering and Networks Laboratory

Institut für Technische Informatik

	-		-	
Spring 2015		Embedded Systems		Page 17

4.2: Scheduling with Resource Constraints

(maximal 25 points)

Determine a resource allocation and a schedule for implementing the following functions:

```
int func (int a, int b, int c, int d, int e, int f, int g, int h){
    int x = (a *1 b) +3 (c *2 d) +4 2;
    int y1 = (e +5 f);
    int y2 = (g +6 h);
    int y = mod (y1, y2);
    int z = (x *8 y);
    return z;
}
int mod (int arg1, int arg2){
    int res = arg1 %7 arg2;
    return res;
}
```

Notation $*_1$ indicates that the index of this multiplication is 1 (node v_1 in a sequence graph).

Additions (+) need 1 cycle to be executed on an adder, ADD. Multiplications (*) and modulo operations (%) need 2 cycles on a multiplier, MUL. Calling a function and returning take 0 time.

(a) (6 points) Provide the hierarchical sequence graph $G_S = (V_S, E_S)$ for function **func**. Denote each node as v_i , where *i* is the index of the corresponding operation in the code.

Spring 2015	Embeddeo	d Systems	Page 18

(b) (6 points) Apply the ASAP and ALAP algorithms to compute the earliest (l_i) and latest (h_i) starting times of all operations $v_i \in V_S$, $i \in \{1, \dots, 8\}$. For ALAP, assume the maximum latency $\overline{L} = 7$. Fill in the starting times in Table 6.

	$l_i(ASAP)$	$h_i(ALAP)$
v_1		
v_2		
v_3		
v_4		
v_5		
v_6		
v_7		
v_8		

Table 6: Starting time bounds given no resource constraints

(c) (13 points) An Integer Linear Program (ILP) needs to be formulated for the problem of *area minimization* of the implementation, under latency constraints. On a given chip, an adder (ADD) requires s(ADD) = 0.15mm² and a multiplier s(MUL) = 0.35mm², respectively. We seek a resource allocation and a feasible schedule with latency not greater than $L_{max} = 8$.

	v_1	v_2	v_3	v_4	v_5	v_6	v_7	v_8
l_i	0	0	2	4	0	0	2	6
h_i	2	2	4	5	2	2	4	6

Table 7: Starting time bounds for operations v_1, \ldots, v_8

For the ILP formulation, the binary variables $x_{i,t} \in \{0,1\}$ are defined $\forall v_i \in V_S$ and $\forall t : l_i \leq t \leq h_i$. Please consider the l_i , h_i values in Table 7 for this question (**not** the values from (b)). $x_{i,t} = 1$ if operation v_i starts executing at time t in a schedule or $x_{i,t} = 0$, otherwise. Function $\tau : V_S \to \mathbb{N}$ can be used to denote the starting time of an operation $v_i \in V_S$ and function $\alpha : V_R \to \mathbb{N}^+$ to denote the allocation of resource instances, where $V_R = \{\text{ADD}, \text{MUL}\}$. Given the above notations:

- (4 points) Express the objective function of the ILP.
- (2 points) Express the latency constraint.
- (2 points) Define $\tau(v_1)$ as a function of $x_{1,t}$, where $l_1 \leq t \leq h_1$.
- (5 points) Express all resource constraints at time t = 2 (without \sum formulation).



4.3: Iterative Algorithms

(maximal 23 points)

Consider the marked graph G_M in Figure 9. The nodes labelled with +, *, **x2** represent addition, multiplication of two input values, and multiplication of an input value by 2, respectively. The edge $f_3 \rightarrow f_2$ has m initial tokens, and the edge $f_3 \rightarrow f_1$ has n initial tokens, where m > n. The input u generates a sequence of numbers, with u(k) being the k-th number.



Figure 9: Marked graph G_M

(a) (6 points) Determine the output value v(k) as a function of the input values for k > m.

(b) (6 points) Assume m = 3 and n = 1. Illustrate the data dependencies among the operations with an equivalent extended sequence graph $G_S = (V_S, E_S, d)$, where $V_S = \{f_0, \dots, f_4\}$ and d_{ij} denotes the index displacement for each $(f_i, f_j) \in E_S$.

(c) (6 points) An addition needs 1 time unit and a multiplication needs 2 time units to execute. Suppose that we implement the algorithm using functional pipelining. Express all data dependency constraints in G_S in the form:

$$\tau(f_j) - \tau(f_i) \ge w(f_i) - d_{ij} \cdot P, \quad \forall (f_i, f_j) \in E_S,$$
(1)

where P is the iteration interval of the pipeline and $w(f_i)$ the execution time of f_i .

(d) (5 points) Assuming unlimited resources, what is the highest throughput $\frac{1}{P}$ that can be achieved with functional pipelining?