



Eidgenössische Technische Hochschule Zürich  
Swiss Federal Institute of Technology Zurich



Institut für Technische Informatik  
und Kommunikationsnetze  
Computer Engineering and Networks Laboratory

Exam Autumn 2017

# Hardware/Software Codesign

Prof. L. Thiele

**Task 1 : Models of Computation**

(maximal 38 points)

**1.1: State Charts**

(maximal 12 points)

Observe the state chart in Figure 1. The following events can trigger a state change: {'red', 'blue', 'powerOn', 'powerOff', 'plus', 'minus', 'sO', 'sA', 'sB'}. 'vA' and 'vB' are two internal variables. For this state chart answer the questions below.

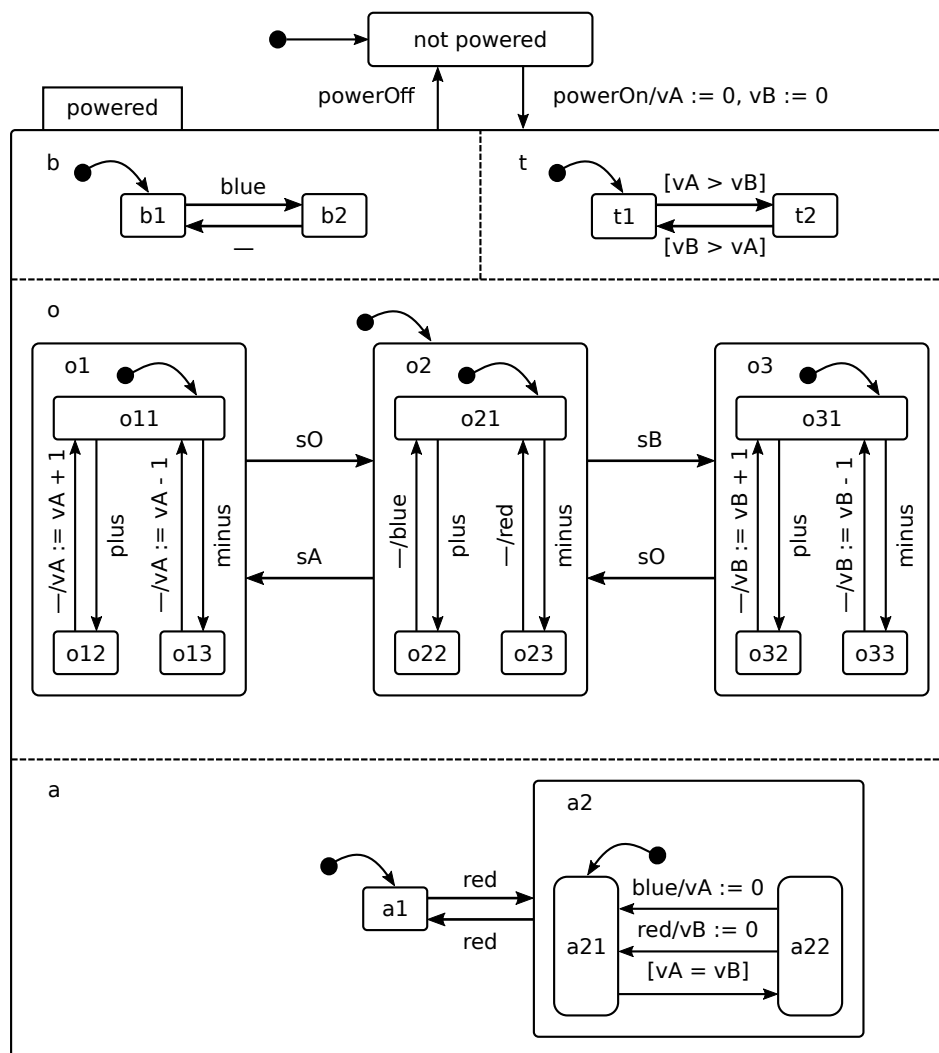


Figure 1: A state chart.

(a) (3 points) In how many different states can the system be? For this question, exclude the variables 'vA' and 'vB'.

(b) (9 points) Determine the sequence of states, starting from the state 'not powered', if the following sequence of external events occurs: 'powerOn', 'red', 'sA', 'plus', 'sO' and 'plus'. Fill out Table 1 by specifying only the **stable** state reached after each external event.

event	b	o	a	t
'powerOn'				
'red'				
'sA'				
'plus'				
'sO'				
'plus'				

Table 1: System state after external events.

**1.2: Dataflow Languages**

(maximal 18 points)

A system modeled by an SDF graph has three processes named 'A', 'B' and 'C', and has the following topology matrix:

$$\begin{pmatrix} 3 & -4 & 0 \\ 0 & \beta & -\alpha \\ -2 & 0 & 4 \end{pmatrix} \quad (1)$$

Suppose that  $\alpha$  and  $\beta$  are positive integers, and initial tokens exist only on the edge between nodes 'A' and 'C'.

- (a) (3 points) Draw the SDF graph and attach labels to the heads and tails of edges that represent the number of tokens consumed and produced at each firing of the corresponding process.

(b) (9 points) Given that the system has sufficient number of initial tokens, determine minimal integer values for the parameters  $\alpha$  and  $\beta$  such that the system has a periodic schedule. Determine the relative execution rates of the processes.

(c) (6 points) Determine a periodic schedule that requires the minimal number of initial tokens on the edge between nodes 'A' and 'C'.

**1.3: KPNs and SDFs**

(maximal 8 points)

- (a) (2 points) A process is described below. Based on the description, argue shortly whether this is a Kahn process or not.

The process 'MERGE' has two input channels and one output channel. When a token appears on any of the input ports, and no token appears at the same time on the other port, it is forwarded to the output channel immediately. If tokens appear on the two input ports simultaneously, they are forwarded to the output immediately, but one predetermined input port's token is always forwarded first. This way, the two input channels are merged into one output.

---

---

- (b) (2 points) A process is described below. Based on the description, argue shortly whether this is a Kahn process or not.

The process 'SPLIT' has one input channel and two output channels. Tokens appearing on the input port are forwarded immediately to one of the output channels in an alternating fashion (i.e. first to one, then to the other, and so on). This way, the input channel is split into two output channels, with each of the output channels receiving half of the events.

---

---

- (c) (2 points) A function is given below. It uses the function `send(int, OutChannel)` to place one token on the specified communication channel. Values `a` and `b` are integers that were read from input channels, while `ch` is an output channel. Based on the code given below, argue shortly whether this could be the fire function of an SDF actor.

```
function fire() {  
    if (a < b) {  
        send(a + b, ch);  
    }  
    else {  
        send(a - b, ch);  
    }  
}
```

---

---

- (d) (2 points) A function is given below. It uses the function `send(int, OutChannel)` to place one token on the specified communication channel. Values `a` and `b` are integers that were read from input channels, while `ch` is an output channel. Based on the code given below, argue shortly whether this could be the fire function of an SDF actor.

```
function fire() {  
    if (a < b) {  
        printf("Warning: The result is not positive");  
    }  
    else {  
        send(a - b, ch);  
    }  
}
```

---

---

**Task 2 : Performance Analysis**

(maximal 36 points)

Determine the Best Case Execution Time (BCET) of the following program, which is written in pseudo code:

```
j = 0;
a = k + 2;

if (i < 10) {
  if (b)
    j = k + a;
  else {
    b = TRUE;
    c = 2 * a + c;
  }
  i = i + 1;
}
k = i + j;
```

The underlying machine executing the program is specified as follows:

- the processor has no pipeline.
- the processor has no registers, i.e. all variables are stored in memory.
- in an assignment, predicate evaluation, or arithmetic operation, variables are accessed in order from right to left e.g.

$a = b + c$

variables are accessed in the order c, b, a.

- consider only the data cache.
- reads and writes from/to a variable are treated the same way. If the required data is in the cache, there is a cache HIT and it takes **1 cycle** to read/write the data. If the data is not in the cache, there is a cache MISS and it takes **100 cycles** to read/write the data.
- execution model is very simplified. No temporary results are considered. Execution times of all processor operations are 1 cycle. In addition, the time for loading and storing a variable needs to be taken into account.



- the data cache is 2-way set-associative. The cache can store a total of 4 data blocks in 2 sets. Each cache block can contain at most one variable. The variables are mapped to sets in the following way:

Variable	Set
a,b,c	1
i,j,k	2

- the cache uses LRU replacement policy within each set once it is full.

Initial conditions:  $a, c, i, j, k \in (-100, 100)$ ,  $b \in \{\text{TRUE}, \text{FALSE}\}$  and cache state is  $((\underbrace{a, -}_{\text{set 1}}, \underbrace{-, -}_{\text{set 2}})$ ,

i.e. set 1 has variable  $a$  in its most recently used block, set 2 is empty.

- (a) (6 points) Determine the basic blocks of the program and draw the corresponding control flow graph.

- (b) (20 points) Use cache MAY analysis to determine the best-case execution time for each basic block.

- (c) (10 points) Write down the structural constraints and the ILP to determine the BCET of one program execution.

**Task 3 : Optimization, Mapping, Partitioning, and Design Space Exploration**  
(maximal 46 points)**3.1: Multi-Criteria Optimization**

(maximal 11 points)

Table 2 shows a design space with 8 points in 3 dimensions. All objectives, the cost  $C$ , the power dissipation  $P$ , and the execution time  $T$  are to be minimized.

Point	Cost [CHF]	Power [mW]	Time [ms]
P1	20	450	800
P2	70	300	500
P3	80	150	300
P4	40	100	400
P5	90	50	200
P6	50	300	700
P7	40	350	900
P8	70	200	500

Table 2: Three dimensional design points.

- (a) (5 points) Identify the points from Table 2 which belong to the Pareto front.

- (b) (6 points) Consider only the Cost and Power dimensions. Draw the points in Fig. 2. Given the reference point  $R(100[\text{CHF}], 500[\text{mW}])$ , find the subset with two points that has the *largest* hypervolume and determine the value of this hypervolume.

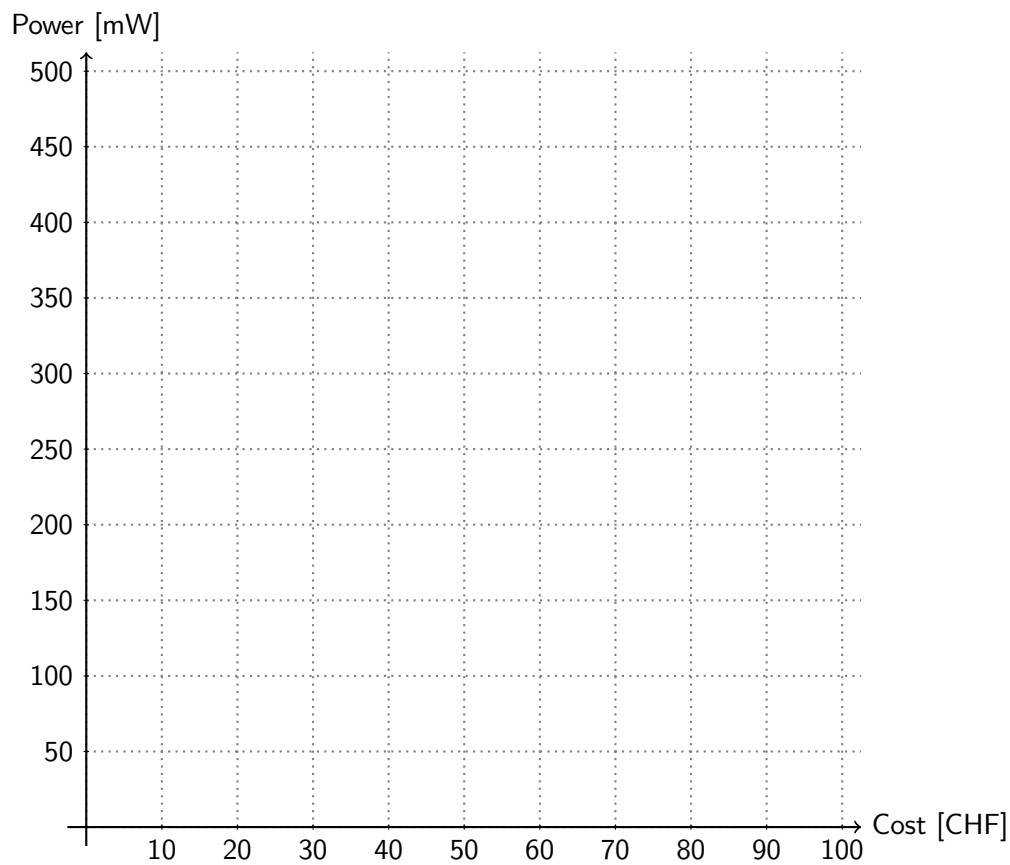


Figure 2: Two-dimensional design space.

**3.2: Design Space Exploration**

(maximal 19 points)

Figure 3 shows a specification at the system level consisting of a synchronous data flow (SDF) graph and an architecture template. Table 3 presents the execution time of the different actors on the available processors. We have one instance of each processor available. Each token has a size of 100 bit, and the bus has a bandwidth of 5000 bit/s. Communication between tasks bound to the same processor does not require any time. The SDF graph is executed 10 times per second.




Figure 3: System-level specification.

Table 3: Execution times for tasks.

processor	execution time [ms]		
	T1	T2	T3
MIPS	5	30	-
ARM	50	10	20

- (a) (4 points) Construct the specification graph. Your graph should include the problem graph with communication nodes, the architecture graph, and all valid bindings.

- (b) (3 points) Determine the minimal number of activations by each actor that are needed to complete one SDF graph execution.

- (c) (8 points) List all design points (allocation, binding, schedule) and evaluate their maximum CPU utilization ( $f_1$ ) and average bus utilization ( $f_2$ ), based on the following definitions:

$$f_1 = \max(\text{busy}(\text{ARM}), \text{busy}(\text{MIPS}))$$

$\text{busy}(P)$  denotes the long-term average utilization of processor  $P$ .

$$f_2 = \sum_{u \text{ mapped to bus}} \frac{b(u)}{t(\text{bus})}$$

$b(u)$  denotes the long-term average communication request rate in [bit/s] and  $t(\text{bus})$  denotes the bus bandwidth in [bit/s].

Plot the points in Figure 4 and determine the Pareto-optimal solutions.



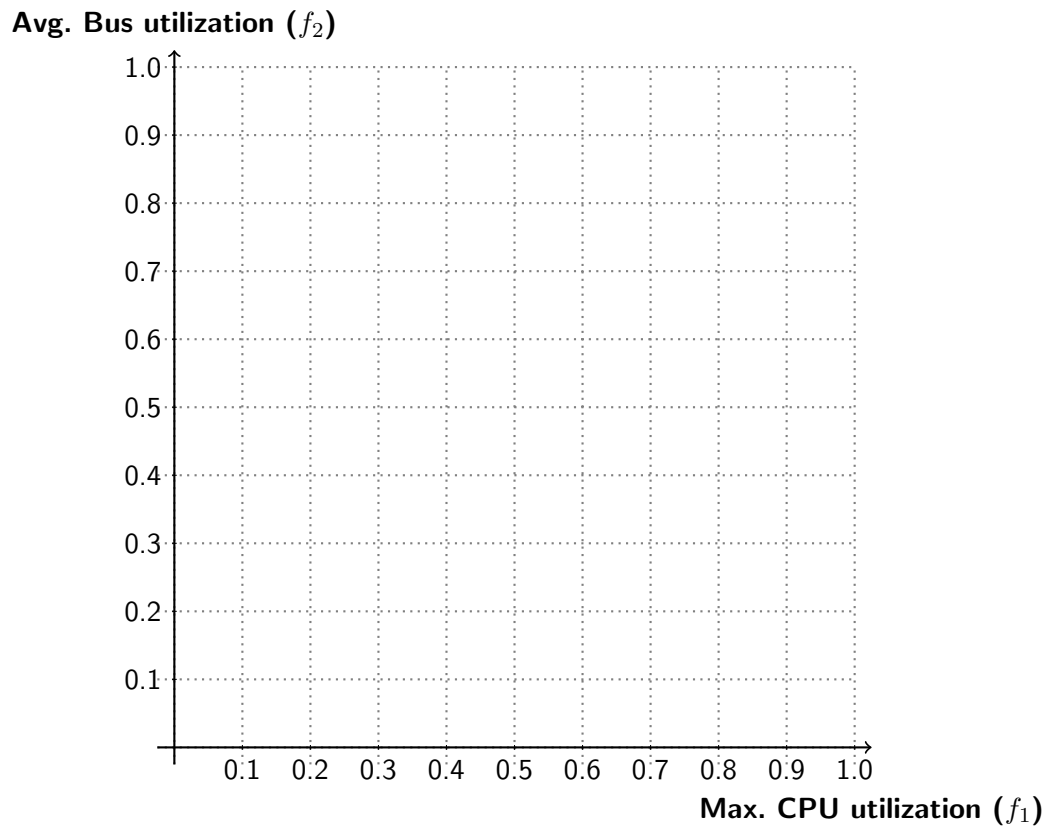


Figure 4: CPU and bus utilization plot.

- (d) (4 points) Consider a new Pareto front with two points  $P_a$  and  $P_b$  whose  $(f_1, f_2)$  values are  $(0.3, 1.0)$  and  $(0.5, 0)$ , respectively. The multi-objective problem is now transformed to a single objective problem using the following function:

$$\text{cost}(P) = K_1 \times f_1 + K_2 \times f_2$$

This function, which needs to be minimized, assigns a single objective value  $\text{cost}(P)$  to an objective vector  $P$ . Determine the range of  $\frac{K_2}{K_1}$  such that  $P_a$  is the optimal solution.

**3.3: Modular Performance Analysis**

(maximal 16 points)

An input stream S1 is processed by CPU, producing an output stream S2. This stream is in turn sent via BUS. Stream S1 has events arriving with a period of 2ms, jitter of 3ms and zero inter-arrival time. CPU handles its incoming events with a constant execution rate of 1 event/ms. BUS is a TDMA resource with a bandwidth of 1 event/ms, where a slot of 1ms is available in every 2ms cycle.

- (a) (6 points) Draw the upper and lower arrival curves for stream S1, and the input service curve for CPU, for time intervals up to 8ms. Determine the maximum delay in [ms].

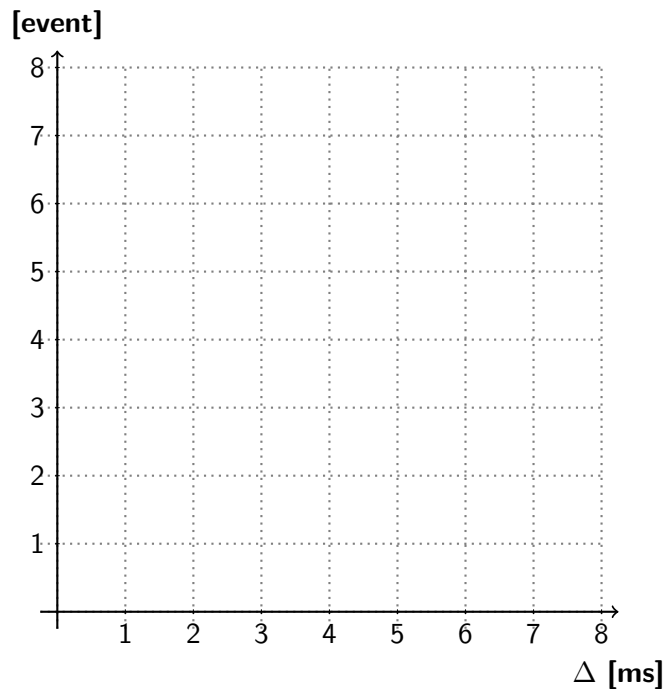


Figure 5: Arrival curves for S1, and input service curve for CPU.

- (b) (10 points) Draw the upper arrival curve for stream S2, and the upper and lower input service curves for BUS, for time intervals up to 8ms. Determine the maximum buffer capacity needed for S2 in [event], and the maximum delay in [ms].

*Hint:* Assume that the output upper arrival curve  $\alpha^{u'}$  of a component can be computed as:

$$\alpha^{u'} = \alpha^u \otimes \beta^l$$

where  $\alpha^u$  is the upper input arrival curve,  $\beta^l$  is the lower input service curve, and  $f \otimes g$  is the min-plus de-convolution:

$$(f \otimes g)(t) = \sup_{u \geq 0} \{f(t + u) - g(u)\}$$




Figure 6: Upper arrival curve for S2, and input service curves for BUS.