



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Institut für Technische Informatik
und Kommunikationsnetze
Computer Engineering and Networks Laboratory



Exam Autumn 2016

Hardware/Software Codesign

Prof. L. Thiele

Task 1 : Models of Computation

(maximal 43 points)

1.1: StateCharts

(maximal 13 points)

Consider the StateChart as shown in Figure 1.

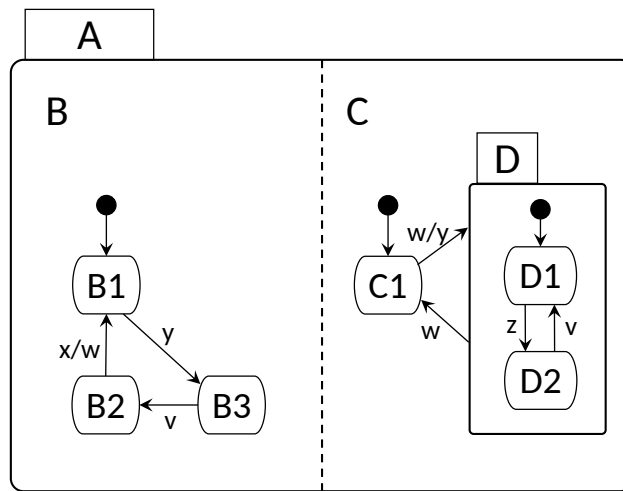


Figure 1: StateChart

- (a) (5 points) Draw the state space of the StateChart as a tree, which shows the hierarchy of states and denote the state types (basic states, sequential states, and parallel states).

- (b) (8 points) Consider the following sequence of external events: x, w, y, y, z, v, x. Determine the sequence of states, the automaton in Figure 1 passes through, starting from the initial state.

event	state
initial	
x	
w	
y	
y	
z	
v	
x	

1.2: Dataflow Languages

(maximal 21 points)

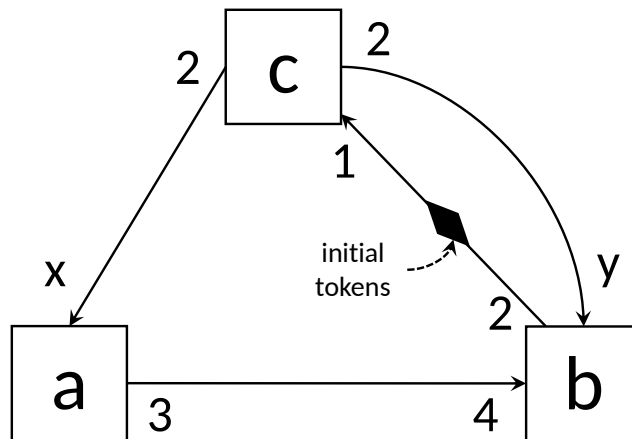


Figure 2: SDF graph

Figure 2 shows a system modeled by a SDF graph with processes $\{a, b, c\}$. The numbers attached to the heads and tails of edges represent the number of tokens consumed and produced respectively, at each firing of the corresponding process.

- (a) (3 points) Write down the topology matrix of the system.

- (b) (8 points) Determine the parameters x and y such that the system has a periodic schedule and note the relative execution rates of the processes.

- (c) (10 points) Determine a periodic schedule that requires a minimum number of initial tokens. **Note:** Only edge bc has initial tokens.

Task 2 : Performance Analysis

(maximal 39 points)

2.1: WCET Analysis

(maximal 21 points)

Determine the WCET of the following program which is written in pseudo code.

```
s=0;

WHILE (k < 5) do {

    if (f)
        j = j + 1;
    else {
        j = 0;
        f = TRUE;
    }

    k = k + 1;
    s = s + k;
}
```

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 **40 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 fully associative. It can store 4 blocks of data. Each cache block can contain exactly one variable.
- the cache uses LRU replacement policy.

Initial conditions: $s, j \in (-\infty, \infty)$, $f \in [0, 1]$, and $k \in [0, +\infty)$ and cache state is $(-, -, -, -)$, i.e. all blocks are empty.

- (a) (5 points) Determine the basic blocks of the program.

- (b) (10 points) Use cache MUST analysis to determine the worst-case execution time for each basic block. **Do not** unroll the loop.

- (c) (6 points) Write down the structural constraints and the ILP to determine the WCET of the program. Given the initial value intervals of $s, j \in (-\infty, +\infty)$, $f \in [0, 1]$, and $k \in [0, +\infty)$, can you specify any additional constraints?

2.2: Modular Performance Analysis

(maximal 18 points)

A bus has to transmit data of two independent packet streams. The arbitration policy used is fixed priority, i.e. the packets of the first stream have the high priority and the bus communicates data from the second stream only if there is no waiting data from the first stream. Data arrives in packets but the transmission of data can be interrupted and resumed any time (preemptive arbitration).

The input streams and the bus are characterized as follows:

- *Stream 1*: Periodic stream with jitter, packet size 1KByte, period 2ms, jitter 1.5ms.
- *Stream 2*: Periodic stream, packet size 1KByte, period 2ms.
- *Bus*: Bandwidth 1MByte/s.

- (a) (4 points) Draw the upper arrival curve for stream 1 and the service curve for the bus using Figure 4.




Figure 4: Arrival curve for stream 1 and bus service

- (b) (2 points) Determine the maximum data delay (in ms) and maximum capacity of the buffer (in KBytes) for stream 1. Justify your answer.

- (c) (10 points) Draw the upper arrival curve for stream 2 and the lower service curve that is left over for stream 2 after communicating data of stream 1. Use Figure 5.




Figure 5: Arrival curve for stream 2 and the bus service remaining from stream 1

- (d) (2 points) Determine the maximum data delay (in ms) and maximum capacity of the buffer (in KBytes) for stream 2. Justify your answer.

Task 3 : Optimization, Mapping and Partitioning, and Design Space Exploration

(maximal 38 points)

3.1: Multi-Criteria Optimization

(maximal 14 points)

Figure 6 shows a 2-dimensional design space with nine design points. Both objectives, the execution time T , and the power dissipation P , are to be minimized.




Figure 6: Design space

(a) (2 points) Indicate the Pareto points in Figure 6.

(b) (3 points) Given the reference point $R(500\text{mW}, 900\text{ms})$, compute the hypervolume of the given nine design points.

(c) (3 points) Given the above mentioned reference point, find the subset of three points with the *largest* hypervolume (environmental selection).

(d) (6 points) A single-objective optimization method is used which minimizes a weighted cost function: $cost(d) = k_1 \cdot T(d) + k_2 \cdot P(d)$, where d denotes a design point and $k_1 + k_2 = 1$.

(a) (5 points) Which design point is returned for $(k_1, k_2) = (0.8, 0.2)$ and what is its cost?

(b) (1 point) Which of the Pareto points in Figure 6 will not be found by this method, irrespective of the weights chosen?

3.2: Design Space Exploration

(maximal 24 points)

Figure 7 shows a specification at the system level consisting of a data flow graph and an architecture template. Table 1 presents the available components with cost and execution times, as well as the time taken for bus communications. We have one instance of each of the components available. Communication between tasks bound to the same component does not require any communication time.




Figure 7: System-level specification

component	cost [CHF]	execution time [ms]		
		T1	T2	T3
ARM	50,-	10	5	8
FPGA	100,-	8	2	4

component	cost [CHF]	communication time [ms]		
		T1→T2	T1→T3	T2→T3
Bus	10,-	5	3	3

Table 1: Components with cost and execution times for tasks and bus communications

- (a) (3 points) Provide tight upper bounds, without considering the task precedence constraints, for the number of distinct bindings and schedules.

- (b) (3 points) Construct the specification graph. Your graph should include the problem graph with communication nodes, the architecture graph, and all possible bindings.

- (c) (14 points) List all design points (allocation, binding, schedule) together with their total execution time and cost. Insert the design points into the time-cost diagram in Figure 8. Identify and mark the Pareto points.




Figure 8: Time-cost diagram

(d) (2 points) The design space exploration is done with an evolutionary algorithm. Present a possible encoding for the allocation and binding.

(e) (2 points) Evaluate the redundancy of the encoding chosen in the task before by comparing the sizes of the search and solution space.