

Family Name:

First Name:

Student ID:

Final Exam

Design of Digital Circuits (252-0028-00L)

ETH Zürich, Spring 2019

Prof. Onur Mutlu

Problem 1 (12 Points):	Boolean Algebra	
Problem 2 (20 Points):	Verilog	
Problem 3 (30 Points):	Finite State Machines (FSM)	
Problem 4 (20 Points):	ISA vs. Microarchitecture	
Problem 5 (20 Points):	Performance Evaluation	
Problem 6 (40 Points):	Pipeline (Reverse Engineering)	
Problem 7 (36 Points):	Tomasulo's Algorithm	
Problem 8 (30 Points):	Systolic Arrays	
Problem 9 (35 Points):	GPUs and SIMD	
Problem 10 (40 Points):	Reverse Engineering Caches	
Problem 11 (30 Points):	Dataflow	
Problem 12 (BONUS: 30 Points):	Branch Prediction	
<hr/>		
Total (343 (313 + 30 bonus) Points):		

Examination Rules:

1. Written exam, 180 minutes in total.
2. **No books, no calculators, no computers or communication devices.** 3 double-sided A4 sheets of handwritten notes are allowed.
3. Write all your answers on this document; space is reserved for your answers after each question.
4. You are provided with scratchpad sheets. Do not answer questions on them. **We will not collect them.**
5. Clearly indicate your final answer for each problem. Answers will only be evaluated if they are readable.
6. Put your Student ID card visible on the desk during the exam.
7. If you feel disturbed, immediately call an assistant.
8. Write with a black or blue pen (no pencil, no green or red color).
9. Show all your work. For some questions, you may get partial credit even if the end result is wrong due to a calculation mistake. If you make assumptions, state your assumptions clearly and precisely.
10. Please write your initials at the top of every page.

Tips:

- **Be cognizant of time. Do not spend too much time on one question.**
- **Be concise. You may be penalized for verbosity.**
- **Show work when needed. You will receive partial credit at the instructors' discretion.**
- **Write legibly. Show your final answer.**

Initials: _____

Design of Digital Circuits

23rd of August, 2019

This page intentionally left blank

1 Boolean Algebra [12 Points]

- (a) [6 Points] Find the simplest sum-of-products representation of the following Boolean equation. Show your work step-by-step.

$$F = (\bar{A} + B + C).(A + B + \bar{C}).C + A$$

- (b) [6 Points] Convert the following Boolean equation so that it contains only NAND operations. Show your work step-by-step.

$$F = \bar{A} + \overline{(B.C + \bar{A}.C)}$$

2 Verilog [20 Points]

Please answer the following three questions about Verilog.

- (a) [5 Points] Does the following code result in a single D Flip-Flop with a synchronous active-low reset? Please explain your answer.

```
1  module mem (input clk, input reset, input [1:0] d, output reg [1:0] q);
2  always @ (posedge clk or negedge reset)
3      begin
4          if (!reset) q <= 0;
5          else q <= d;
6      end
7  endmodule
```

- (b) [5 Points] Does the following code result in a sequential circuit or a combinational circuit? Please explain your answer.

```
1  module Mask (input [1:0] data_in, input mask, output reg [1:0] data_out);
2  always @ (*)
3      begin
4          data_out[1] = data_in[1];
5          if (mask)
6              data_out[0] = 0;
7      end
8  endmodule
```

(c) [10 Points] Is the following code syntactically correct? If not, please explain the mistake(s) and how to fix it/them.


```
1  module fulladd(input a, b, c, output reg s, c_out);
2      assign s = a^b;
3      assign c_out = (a & b) | (b & c) & (c & a);
4  endmodule
5
6  module top ( input wire [5:0] instr, input wire op, output z);
7
8      reg[1:0] r1, r2;
9      wire [3:0] w1, w2;
10
11     fulladd FA1 (.a(instr[0]), .b(instr[1]), .c(instr[2]),
12                 .c_out(r1[1]), .z(r1[0]));
13     fulladd FA2 (.a(instr[3]), .b(instr[4]), .c(instr[5]),
14                 .z(r2[0]), .c_out(r2[1]));
15
16     assign z = r1 | op;
17     assign w1 = r1 + 1;
18     assign w2 = r2 << 1;
19     assign op = r1 ^ r2;
20
21 endmodule
```

3 Finite State Machines (FSM) [30 Points]

You are given two *one-bit* input signals (T_A and T_B) and one *one-bit* output signal (O) for the following modular equation: $2N(T_A) + N(T_B) \equiv 2 \pmod{4}$. In this modular equation, $N(T_A)$ and $N(T_B)$ represent the **total number of times** the inputs T_A and T_B are high (i.e., logic 1) at each positive clock edge, respectively. The one-bit output signal, O , is set to 1 when the modular equation is satisfied (i.e., $2N(T_A) + N(T_B) \equiv 2 \pmod{4}$), and 0 otherwise. An example that sets $O = 1$ at the end of the fourth cycle would be:

- (1st cycle) $T_A = 0$ ($N(T_A) = 0$), $T_B = 0$ ($N(T_B) = 0$), $2N(T_A) + N(T_B) \equiv 0 \pmod{4} \Rightarrow O = 0$
- (2nd cycle) $T_A = 1$ ($N(T_A) = 1$), $T_B = 1$ ($N(T_B) = 1$), $2N(T_A) + N(T_B) \equiv 3 \pmod{4} \Rightarrow O = 0$
- (3rd cycle) $T_A = 1$ ($N(T_A) = 2$), $T_B = 0$ ($N(T_B) = 1$), $2N(T_A) + N(T_B) \equiv 1 \pmod{4} \Rightarrow O = 0$
- (4th cycle) $T_A = 0$ ($N(T_A) = 2$), $T_B = 1$ ($N(T_B) = 2$), $2N(T_A) + N(T_B) \equiv 2 \pmod{4} \Rightarrow O = 1$

- (a) [10 Points] You are given a partial **Moore** machine state transition diagram that corresponds to the modular equation described above. However, the input labels of most of the transitions are still missing in this diagram. Please label the transitions with the correct inputs so that the FSM correctly implements the above specification.



- (b) [10 Points] Describe the FSM with Boolean equations assuming that the states are encoded with **one-hot encoding**. Assign state encodings while using the **minimum** possible number of bits to represent the states. Please indicate the values you assign to each state.

- (c) [10 Points] Describe the FSM with Boolean equations assuming that the states are encoded with **binary encoding** (i.e., fully encoding). Assign state encodings while using the **minimum** possible number of bits to represent the states. Please indicate the values you assign to each state.

4 ISA vs. Microarchitecture [20 Points]

A new CPU has two comprehensive user manuals available for purchase as shown in Table 1.

Manual Title	Cost	Description
the_isa.pdf	CHF 1 million	describes the ISA in detail
the_microarchitecture.pdf	CHF 10 million	describes the microarchitecture in detail

Table 1: Manual Costs

Unfortunately, the manuals are extremely expensive, and you can only afford one of the two. If both manuals might be useful, you would prefer the cheaper one.

For each of the following questions that you would like to answer, decide which manual is more likely to help. *Note: we will subtract 1 point for each **incorrect** answer. For an unanswered question, you will get +0 points.*

1. [2 Points] The latency of a branch predictor misprediction.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
2. [2 Points] The size of a physical memory page.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
3. [2 Points] The memory-mapped locations of exception vectors.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
4. [2 Points] The function of each bit in a programmable branch-predictor configuration register.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
5. [2 Points] The bit-width of the interface between the CPU and the L1 cache.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
6. [2 Points] The number of pipeline stages in the CPU.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
7. [2 Points] The order in which loads and stores are executed by a multi-core CPU.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
8. [2 Points] The memory addressing modes available for arithmetic operations.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
9. [2 Points] The program counter width.
 1. the_isa.pdf
 2. the_microarchitecture.pdf
10. [2 Points] The number of cache sets at each level of the cache hierarchy.
 1. the_isa.pdf
 2. the_microarchitecture.pdf

5 Performance Evaluation [20 Points]

You are the leading engineer of a new processor. Both the design of the processor and the compiler for it are already done. Now, you need to decide if you will send the processor to manufacturing at its current stage or if you will delay the production to introduce last-minute improvements to the design. To make the decision, you meet with your team to brainstorm about how to improve the design. Together, after profiling the target applications for the processor, you come up with two options:

- **Keep the current project.** For version A of the processor, the clock frequency is 600 MHz, and the following measurements are obtained:

Instruction Class	CPI	Frequency of Occurrence
A	2	40%
B	3	25%
C	3	25%
D	7	10%

- **Include optimizations to the design.** For version B of the processor, the clock frequency is 700 MHz. The ISA for processor B includes three new types of instructions. Those three new types of instructions increase the total number of executed instructions for processor B by 50%, in comparison to processor A. The following measurements are obtained:

Instruction Class	CPI	Frequency of Occurrence
A	2	15%
B	2	15%
C	4	10%
D	6	10%
E	1	10%
F	2	20%
G	2	20%

- (a) [7 Points] What is the CPI of each version? Show your work.

CPI_A :

CPI_B :

- (b) [6 Points] What are the MIPS (Million Instructions Per Second) of each version? Show your work.

$MIPS_A$:

$MIPS_B$:

- (c) [7 Points] Considering your team is aiming to release to the market the processor that gives better performance when executing the target application, which processor version will you choose as the final design? Show your work.

6 Pipeline (Reverse Engineering) [40 Points]

The following piece of code runs on a pipelined microprocessor as shown in the table (F: Fetch, D: Decode, E: Execute, M: Memory, W: Write back). Instructions are in the form “Instruction Destination, Source1, Source2.” For example, “ADD A, B, C” means $A \leftarrow B + C$.

Cycles		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
0	MUL R5, R6, R7	F	D	E1	E2	E3	E4	M	W										
1	ADD R4, R6, R7		F	D	E1	E2	E3	-	M	W									
2	ADD R5, R5, R6			F	D	-	-	E1	E2	E3	M	W							
3	MUL R4, R7, R7				F	-	-	D	E1	E2	E3	E4	M	W					
4	ADD R6, R7, R5							F	D	-	E1	E2	E3	M	W				
5	ADD R3, R0, R6								F	-	D	-	-	E1	E2	E3	M	W	
6	ADD R7, R1, R4										F	-	-	D	E1	E2	E3	M	W

Use this information to reverse engineer the architecture of this microprocessor to answer the following questions. Answer the questions as precise as possible with the provided information. If the provided information is not sufficient to answer a question, answer “Unknown” and explain your reasoning clearly.

- (a) [5 Points] How many cycles does it take for an adder and for a multiplier to calculate a result?

- (b) [5 Points] What is the minimum number of register file read/write ports that this architecture implements? Explain.

- (c) [5 Points] Can we reduce the execution time of this code by enabling more read/write ports in the register file? Explain.

- (d) [5 Points] Does this architecture implement any data forwarding? If so, how is data forwarding done between pipeline stages? Explain.

- (e) [5 Points] Is it possible to run this code faster by adding more data forwarding paths? If it is, how? Explain.

- (f) [5 Points] Is there internal forwarding in the register file? If there is not, how would the execution time of the same program change by enabling internal forwarding in the register file? Explain.

7 Tomasulo’s Algorithm [36 Points]

In this problem, we consider an in-order fetch, out-of-order dispatch, and out-of-order retirement execution engine that employs Tomasulo’s algorithm. This engine behaves as follows:

- The engine has four main pipeline stages: Fetch (F), Decode (D), Execute (E), and Write-back (W).
- The engine can fetch **FW** instructions per cycle, decode **DW** instructions per cycle, and write back the result of **RW** instructions per cycle.
- The engine has two execution units: 1) an *integer ALU* for executing integer instructions (i.e., addition and multiplication) and 2) a *memory unit* for executing load/store instructions.
- Each execution unit has an **R**-entry reservation station.
- An instruction always allocates the first available entry of the reservation station (in top-to-bottom order) of the corresponding execution unit.

The reservation stations are all initially empty. The processor fetches and executes *six* instructions. Table 2 shows the six instructions and their execution diagram.

Using the information provided above and in Table 2 (see the next page), fill in the blanks below with the configuration of the out-of-order microarchitecture. Write “Unknown” if the corresponding configuration cannot be determined using the information provided in the question.

The latency of the ALU and memory unit instructions:	
In which pipeline stage is an instruction dispatched?	
Number of entries of each reservation station (R):	
Fetch width (FW):	
Decode width (DW):	
Retire width (RW):	
Is the integer ALU pipelined?	
Is the memory unit pipelined?	
If applicable, between which stages is data forwarding implemented?	

Instruction/Cycle:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33		
1: ADD R1 ← R0, R1	F	D	E1	E2	W																														
2: LD R2 ← [R1]	F	D	-	-	E1	E2	E3	E4	E5	E6	E7	E8	E9	E10	W																				
3: ADDI R1 ← R1, #4	F	D	D	-	-	E1	E2	W																											
4: LD R3 ← [R1]	F	D	D	-	-	-	-	-	E1	E2	E3	E4	E5	E6	E7	E8	E9	E10	W																
5: MUL R4 ← R2, R3			F	-	-	D	-	-	-	-	-	-	-	-	-	-	-	-	E1	E2	W														
6: ST [R0] ← R4			F	-	-	-	-	-	-	-	-	-	-	-	-	-	D	-	-	-	-	E1	E2	E3	E4	E5	E6	E7	E8	E9	E10	W			

Table 2: Execution diagram of the six instructions.

8 Systolic Arrays [30 Points]

A systolic array consists of 3x4 Processing Elements (PEs), interconnected as shown in Figure 1. The inputs of the systolic array are labeled as H0, H1, H2 and V0,V1,V2,V3. Figure 2 shows the PE logic, which performs a multiply and accumulate operation (MAC), and it saves the result in an internal register (reg). Figure 2 also shows how each PE propagates its inputs. We make the following assumptions:

- The latency of each MAC is one cycle.
- The propagation of the values from i_0 to o_0 , and from i_1 to o_1 , takes one cycle.
- The initial value of all registers is zero.
- You can input a value more than once in the systolic array.

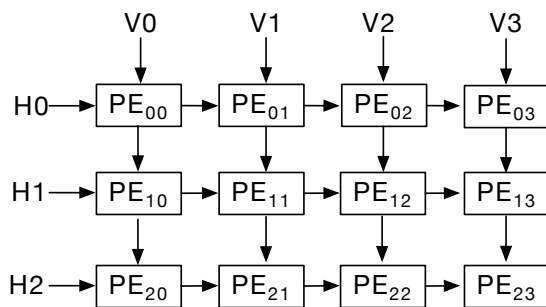


Figure 1: PE array




Figure 2: Processing Element (PE)

Your goal is to use this systolic array to perform the convolution of a 3x3 image (matrix I) with three 2x2 filters (matrices F, G, and H), to obtain three outputs (matrices O, U, and E):

$$\begin{matrix} I_{00} & I_{01} & I_{02} \\ I_{10} & I_{11} & I_{12} \\ I_{20} & I_{21} & I_{22} \end{matrix} \otimes \begin{matrix} F_{00} & F_{01} \\ F_{10} & F_{11} \end{matrix} = \begin{matrix} O_{00} & O_{01} \\ O_{10} & O_{11} \end{matrix}$$

$$\begin{matrix} I_{00} & I_{01} & I_{02} \\ I_{10} & I_{11} & I_{12} \\ I_{20} & I_{21} & I_{22} \end{matrix} \otimes \begin{matrix} G_{00} & G_{01} \\ G_{10} & G_{11} \end{matrix} = \begin{matrix} U_{00} & U_{01} \\ U_{10} & U_{11} \end{matrix}$$

$$\begin{matrix} I_{00} & I_{01} & I_{02} \\ I_{10} & I_{11} & I_{12} \\ I_{20} & I_{21} & I_{22} \end{matrix} \otimes \begin{matrix} H_{00} & H_{01} \\ H_{10} & H_{11} \end{matrix} = \begin{matrix} E_{00} & E_{01} \\ E_{10} & E_{11} \end{matrix}$$

As an example, the convolution of the matrix I with the filter F is computed as follows:

- $O_{00} = I_{00} * F_{00} + I_{01} * F_{01} + I_{10} * F_{10} + I_{11} * F_{11}$
- $O_{01} = I_{01} * F_{00} + I_{02} * F_{01} + I_{11} * F_{10} + I_{12} * F_{11}$
- $O_{10} = I_{10} * F_{00} + I_{11} * F_{01} + I_{20} * F_{10} + I_{21} * F_{11}$
- $O_{11} = I_{11} * F_{00} + I_{12} * F_{01} + I_{21} * F_{10} + I_{22} * F_{11}$

You should compute the three convolutions in the minimum possible amount of cycles. Fill the following table with:

1. The input values (matrices I, F, G, and H) in the correct input ports of the systolic array (the values can be repeated).
2. The output values and the corresponding PE where the outputs (matrices O, U, and E) are generated.

Fill the gaps only with relevant information.

cycle	H0	H1	H2	V0	V1	V2	V3	PE ₀₀	PE ₀₁	PE ₀₂	PE ₀₃	PE ₁₀	PE ₁₁	PE ₁₂	PE ₁₃	PE ₂₀	PE ₂₁	PE ₂₂	PE ₂₃	
0																				
1																				
2																				
3																				
4																				
5																				
6																				
7																				
8																				
9																				
10																				
11																				
12																				
13																				
14																				
15																				

9 GPUs and SIMD [35 Points]

We define the *SIMD utilization* of a program that runs on a GPU as the fraction of SIMD lanes that are kept busy with *active threads* during the run of a program. As we saw in lecture and practice exercises, the SIMD utilization of a program is computed across the *complete run* of the program.

The following code segment is run on a GPU. Each thread executes a **single iteration** of the shown loop. Assume that the data values of the arrays A and B are already in vector registers so there are no loads and stores in this program. (Hint: Notice that there are 4 instructions in each iteration.) A warp in the GPU consists of 32 threads, and there are 32 SIMD lanes in the GPU.

```
for (i = 0; i < 1026; i++) {  
    if (A[i] < 33) {           // Instruction 1  
        B[i] = A[i] << 1;    // Instruction 2  
    }  
    if (A[i] > 33) {           // Instruction 3  
        B[i] = A[i] >> 1;    // Instruction 4  
    }  
}
```

Please answer the following five questions.

- (a) [2 Points] How many warps does it take to execute this program?

- (b) [10 Points] What is the maximum possible SIMD utilization of this program? Show your work. (Hint: The warp scheduler does not issue instructions where no threads are active).

- (c) [5 Points] Please describe what needs to be true about array A to reach the maximum possible SIMD utilization asked in part (b). (Please cover all cases in your answer.)

- (d) [13 Points] What is the minimum possible SIMD utilization of this program? Show your work.

- (e) [5 Points] Please describe what needs to be true about array A to reach the minimum possible SIMD utilization asked in part (d). (Please cover all cases in your answer.)

10 Reverse Engineering Caches [40 Points]

You are trying to reverse-engineer the characteristics of a cache in a system, so that you can design a more efficient, machine-specific implementation of an algorithm you are working on. To do so, you have come up with three sequences of memory accesses to various *bytes* in the system in an attempt to determine the following four cache characteristics:

- Cache block size (8, 16, 32, 64, or 128 B).
- Cache associativity (1-, 2-, 4-, or 8-way).
- Cache size (4 or 8 KB).
- Cache replacement policy (LRU or FIFO).

The only statistic that you can collect on this system is *cache hit rate* after performing each sequence of memory accesses. Here is what you observe:

Sequence	Addresses Accessed (Oldest → Youngest)								Hit Rate
1.	31	8192	63	16384	4096	8192	64	16384	3/8
2.	32768	0	129	1024	3072	8192			0
3.	0	4	8	4096	64	128			1

Assume that the cache is initially empty at the beginning of the first sequence, but *not* at the beginning of the second and third sequences. The sequences are executed back-to-back, i.e., no other accesses take place in between the three sequences. Thus, **at the beginning of the second (third) sequence, the contents are the same as at the end of the first (second) sequence.**

Based on what you observe, what are the following characteristics of the cache? Explain to get points. If a characteristic cannot be known, then write "Unknown" and explain.

(a) [10 Points] Cache block size (8, 16, 32, 64, or 128 B)?

(b) [10 Points] Cache associativity (1-, 2-, 4-, or 8-way)?

(c) [10 Points] Cache size (4 or 8 KB)?

(d) [10 Points] Cache replacement policy (LRU or FIFO)?

11 Dataflow [30 Points]

- We define the *switch node* in Figure 3 to have 2 inputs (**I**, **Ctrl**) and 1 output (**O**). The *Ctrl* input always enters perpendicularly to the switch node. If the *Ctrl* input has a *True* token (i.e., a token with a value of 1), the **O** wire propagates the value on the **I** wire. Else, the 2 input tokens (**I**, **Ctrl**) are consumed, and no token is generated at the output (**O**).
- We define the *inverter node* in Figure 4 to have 1 input (**I**) and 1 output (**O**). The node negates the input token (i.e., $O = !I$).
- We define the *TF node* in Figure 5 to have 3 inputs (I_F , I_T , **Ctrl**) and 1 output (**O**). When **Ctrl** is set to True, **O** takes I_T . When **Ctrl** is set to False, **O** takes I_F .
- The \geq node outputs True only when the left input is greater than or equal to the right input.
- The +1 node outputs the input plus one.
- The + node outputs the sum of the two inputs.
- A node generates an output token when tokens exist at *every* input, and *all* input tokens are consumed.
- Where a single wire splits into multiple wires, the token travelling on the wire is replicated to all wires.




Figure 3: Switch Node




Figure 4: Inverter Node





Figure 5: TF Node

Consider the dataflow graph on the following page. Numbers in dashed boxes represent tokens (with the value indicated by the number) in the initial state. The **X** and **Y** inputs automatically produce tokens as soon as the previous token on the wire is consumed. The order of these tokens follows the pattern (*note, the following are all single digit values spaced appropriately for the reader to easily notice the pattern*):

X: 0 01 011 0111 01111

Y: 1 22 333 4444 55555

Consider the dataflow graph on the following page. Please clearly describe the sequence of tokens generated at the output (OUT).



12 BONUS: Branch Prediction [30 Points]

Assume a machine with a two-bit global history register (GHR) shared by all branches, which starts with Not Taken, Not Taken (2'b00). Each pattern history table entry (PHTE) contains a 2-bit saturating counter. The saturating counter values are as follows:

2'b00 - Strongly Not Taken
 2'b01 - Weakly Not Taken
 2'b10 - Weakly Taken
 2'b11 - Strongly Taken

Assume the following piece of code runs on this machine. The code has two branches (labeled B1 and B2). When we say that a branch is taken, we mean that the code inside the curly brackets is executed. For the following questions, assume that this is the only block of code that will ever be run, and the loop-condition branch (B1) is resolved first in the iteration before the if-condition branch (B2).

```
for (int i = 0; i < 1000000; i++) { /* B1 */
    /* TAKEN PATH for B1 */
    if (i % 3 == 0) { /* B2 */
        j[i] = k[i] - 1; /* TAKEN PATH for B2 */
    }
}
```

- (a) [20 Points] Is it possible to observe that the branch predictor mispredicts 100% of the times in the first 5 iterations of the loop? If yes, fill in the table below with all possible initial values each entry can take. We represent Not Taken with N, and Taken with T.

Table 3: PHT

PHT Entry	Value
TT	
TN	
NT	
NN	

Show your work here.

- (b) [10 Points] At steady-state, we observe the following pattern which repeats over time: TTTNTN, with T representing Taken, and N representing Not Taken. When GHR pattern equals to NT or TT, the predictor will observe that the branch outcome will be either T or N. Therefore, no matter what the initial values for these two entries are in the pattern history table (PHT), only one of the branches can be predicted correctly. Thus prediction accuracy will never reach 100%. Explain how using local history registers instead of the global history register will help bring the prediction accuracy up to 100% during the steady state, by showing what each PHTE will saturate to.