## Final Exam

# Digital Design and Computer Architecture (252-0028-00L) ETH Zürich, Spring 2021 

| Prof. Onur Mutlu |  |  |
| :---: | :---: | :---: |
| Problem 1 (20 Points): | Boolean Logic Circuits |  |
| Problem 2 (60 Points): | Verilog |  |
| Problem 3 (45 Points): | Finite State Machines |  |
| Problem 4 (30 Points): | ISA vs. Microarchitecture |  |
| Problem 5 (45 Points): | Performance Evaluation |  |
| Problem 6 (65 Points): | Pipelining |  |
| Problem 7 (60 Points): | Tomasulo's Algorithm |  |
| Problem 8 (75 Points): | GPUs and SIMD |  |
| Problem 9 (45 Points): | Branch Prediction |  |
| Problem 10 (70 Points): | Caches |  |
| Problem 11 (BONUS: 25 Points): | Prefetching |  |
| Problem 12 (BONUS: 35 Points): | Systolic Arrays |  |
| Total (575 (515 + 60 bonus) Points): |  |  |

## Examination Rules:

1. Written exam, 180 minutes in total.
2. No books, no calculators, no computers or communication devices. 3 double-sided (or 6 onesided) 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, red or any other 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.

This page intentionally left blank

## 1 Boolean Logic Circuits [20 points]

(a) [10 points] Using Boolean algebra, find the simplest Boolean algebra equation for the following min-terms:
$\sum(1111,1110,1000,1001,1011,1010,0000)$. Show your work step-by-step.
$\square$
(b) [10 points] Convert the following Boolean equation so that it only contains NOR operations. Show your work step-by-step.

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

## 2 Verilog [60 points]

### 2.1 Complete the Verilog code [30 points]

For each numbered blank (1)-5 in the following Verilog code, mark the choice below (i.e., one of options A, B, C, D) that makes the Verilog module operate as described in the comments. The resulting code must have correct syntax.

```
module my_module (input clk, input rst,
    input[15:0] idata, input[1:0] op, (1)[31:0] odata);
    (2) nval = 32'd0; // defining a 32-bit signal with an initial value of 0
    always@* begin
        case (op)
            2'b00:
                nval = odata + idata; // when 'op' is decimal 0, add 'idata' to
                // 'odata' and assign the result to 'nval'
                2'b01:
                    nval = odata - idata; // when 'op' is decimal 1, subtract 'idata'
                        // from 'odata' and assign the result to 'nval'
                2'b10:
                nval = idata; // when 'op' is decimal 2, assign 'idata' to 'nval'
                (3):
                    nval = 0; // when 'op' is decimal 3, assign 0 to 'nval'
        endcase
    end
    // executing the following always block on the rising edge of 'clk'
    always@ (posedge clk) begin
        if (rst)
            (4) // resetting 'odata' to 0 for the next cycle
        else
            (5) /l assigning 'nval' to 'odata' for the next cycle
    end
endmodule
```

Provide your choice for each blank (1)-5 below:
(1):
A. output B. output reg
C. output wire
D. input reg
(2):
A. $\operatorname{reg}[31: 0]$
B. input[31:0]
C. wire[31:0]
D. int[31:0]
(3):
A. $2^{\prime} \mathrm{b} 3$
B. $3^{\prime} \mathrm{b} 3$
C. 2'h11
D. default
(4):
A. assign odata $<=0$;
B. assign odata $=0$;
C. odata $==0$;
D. odata $<=0$;
(5):
A. assign odata $<=$ nval;
B. assign odata $=$ nval;
C. odata $==$ nval;
D. odata $<=$ nval;

### 2.2 What Does This Code Do? [30 points]

You are given a Verilog code that you are asked to analyze and find out what it does.

```
module my_module2 (input clk, output[1:0] out);
    reg state = 1'b0;
    reg[1:0] my_reg = 0;
    always@(posedge clk) begin
        state <= &out ? ~state : state;
    end
    always@(posedge clk) begin
        case(state)
            1'b0: begin
            my_reg <= my_reg + 1;
            end
            1'b1: begin
            my_reg <= my_reg - 1;
            end
        endcase
    end
    assign out = my_reg;
endmodule
```

Show the values (as unsigned decimal numbers) that the out signal takes, starting from the initial state of the module, for 16 consecutive clock (i.e., clk) cycles. Explain your answer briefly.
$\square$

## 3 Finite State Machines [45 points]

### 3.1 Simplifying an FSM [20 points]

You are given the Mealy state machine of a one input / one output digital circuit design. Answer the following questions for the given state diagram.

(a) [10 points] Is it possible to simplify this state diagram and reduce the number of states? If so, simplify it to the minimum number of states. Explain each step of your simplification. Draw the simplified state diagram. If not, explain why it is not possible to simplify the state diagram.
$\square$
(b) [10 points] Assume this state machine is used to process binary numbers from the least significant bit to the most significant bit. You are given an input bit stream: "10110100". Please show the output bit stream produced by this FSM.
$\square$

### 3.2 Designing an FSM [25 points]

Design a Moore finite state machine (FSM) with one input and one output. The input provides an unsigned binary number in a bit-serial fashion from the most-significant bit to the least-significant bit. The output should be logic-1 in a clock cycle if the provided input so far is divisible by 8 (i.e., [the input number $] \bmod 8=0$ ). (Hint: Recall that the output depends only on the current state in a Moore FSM.) Below are some example bit-streams that should output a logic-1 value.

- 1000
- 10000
- 11000
- 111000
- 101000

To start an input bit stream, the user should reset the FSM. Draw the state diagram and explain why it works. Your state machine should use as few states as possible and each state should have a precise definition and output.

## 4 ISA vs. Microarchitecture [30 points]

A new CPU has two comprehensive user manuals available for purchase which describe the ISA and the microarchitecture of the CPU, respectively.
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 ISA manual since it is much cheaper than the microarchitecture manual.
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 and award 0 points for unanswered questions (the minimum number of total points you can get for this question is 0).

1. [2 points] Number of uniquely identifiable memory locations.
2. ISA
3. Microarchitecture
4. [2 points] Number of instructions fetched per clock cycle.
5. ISA
6. Microarchitecture
7. [2 points] Support for branch prediction hints conveyed by the compiler.
8. ISA
9. Microarchitecture
10. [2 points] Number of general-purpose registers.
11. ISA
12. Microarchitecture
13. [2 points] Number of non-programmable registers.
14. ISA
15. Microarchitecture
16. [2 points] SIMD processing support.
17. ISA
18. Microarchitecture
19. [2 points] Number of integer arithmetic and logic units (ALUs).
20. ISA
21. Microarchitecture
22. [2 points] Number of read ports in the physical register file.
23. ISA
24. Microarchitecture
25. [2 points] Endianness (big endian vs. small endian).
26. ISA
27. Microarchitecture
28. [2 points] Size of a virtual memory page.
29. ISA
30. Microarchitecture
31. [2 points] Cache coherence protocol.
32. ISA
33. Microarchitecture
34. [2 points] Number of cache blocks in the L3 cache.
35. ISA
36. Microarchitecture
37. [2 points] Ability to flush (i.e., invalidate) a cache line using the operating system code.
38. ISA
39. Microarchitecture
40. [2 points] Number of pipeline stages.

> 1. ISA 2. Microarchitecture
15. [2 points] How many prefetches the hardware prefetcher generates in a clock cycle.

1. ISA
2. Microarchitecture

## 5 Performance Evaluation [45 points]

A multi-cycle processor $P 1$ executes load instructions in $\mathbf{6}$ cycles, store instructions in $\mathbf{6}$ cycles, arithmetic instructions in 2 cycles, and branch instructions in 2 cycles. Consider an application $A$ where $40 \%$ of all instructions are load instructions, $20 \%$ of all instructions are store instructions, $30 \%$ of all instructions are arithmetic instructions, and $10 \%$ of all instructions are branch instructions.
(a) [10 points] What is the CPI of application $A$ when executing on processor $P 1$ ? Show your work.
$\square$
(b) [10 points] A new design of the processor doubles the clock frequency of $P 1$. However, the latencies of all instructions increase by 4 cycles. We call this new processor $P 2$. The compiler used to generate instructions for $P 2$ is the same as for $P 1$. Thus, it produces the same number of instructions for program $A$. What is the CPI of application $A$ when executing on processor P2? Show your work.
$\square$
(c) [5 points] Which processor is faster $(P 1$ or $P 2)$ ? By how much (i.e., what is the speedup)? Show your work.
$\square$
$\qquad$
(d) [20 points] You want to improve the original P1 design by including one new optimization without changing the clock frequency. You can choose only one of the following options:
(1) ALU: An optimized $A L U$, which halves the latency of both arithmetic and branch instructions.
(2) LSU: An asymmetric load-store unit, which halves the latency of load operations but doubles the latency of store operations.

Which optimization do you add to $P 1$ for application $A$ ? Show your work and justify your choice.

## 6 Pipelining [65 points]

Consider two pipelined machines implementing the MIPS ISA, Machine A and Machine B. Both machines have one $A L U$ and the following five pipeline stages, very similar to the basic 5 -stage pipelined MIPS processor we discussed in lectures:

1. Fetch (one clock cycle)
2. Decode (one clock cycle)
3. Execute (one clock cycle)
4. Memory (one clock cycle)
5. Write-back (one clock cycle).

Machines A and B have the following specifications:

|  | Machine A | Machine B |
| :--- | :--- | :--- |
| Data <br> Forward- <br> ing/Inter- <br> locking | Does NOT implement interlocking in hard- <br> ware. Relies on the compiler to order in- <br> structions or insert nop instructions such <br> that dependent instructions are correctly <br> executed. | Implements data dependence detection and <br> data forwarding in hardware. On detection <br> of instruction dependence, it forwards an <br> operand from the memory stage or from the <br> write-back stage to the execute stage. The <br> result of a load instruction (lw) can only be <br> forwarded from the write-back stage. |
| Internal <br> register file <br> forwarding | Implemented (i.e., an instruction writes into <br> a register in the first half of a cycle and <br> another instruction can correctly access the <br> same register in the second half of the cycle). | Same as Machine A <br> Branch |
| Predicts all branches as always-taken, and <br> the next program counter is available after <br> the decode stage. | Same as Machine A |  |

Consider the following code segment:

```
Loop: lw $1, 0($4)
    lw $2, 400($4)
    add $3, $1, $2
    sw $3, 0($4)
    sub $4, $4, #4
    bnez $4, Loop
Initially, $1 = 0, $2 = 0, $3 = 0, and $4 = 400.
```

$\qquad$
(a) [15 points] Re-write the code segment above with minimal changes so that it gets correctly executed in Machine A with minimal latency. You can either insert nop instructions or reorder instructions as needed.
$\square$
(b) [15 points] Fill the table below with the timeline of the first loop iteration of the code segment in Machine A.

| Instruction | Clock cycle number |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 1 | 2 | 3 | 4 | 5 | 6 |  | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |

$\qquad$
(c) [10 points] Calculate the number of cycles it takes to execute the code segment on Machine A. Show your work in the box.
$\square$
(d) [15 points] Fill the table below with the timeline of the first loop iteration of the code segment in Machine B.

| Instruction | Clock cycle number |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 |
| 1w \$1, 0(\$4) |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| lw \$2, 400 (\$4) |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| add \$3, \$1, \$2 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| sw \$3, 0 (\$4) |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| sub \$4, \$4, \#4 |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| bnez \$4, Loop |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |

(e) [10 points] Calculate the number of cycles it takes to execute the code segment on Machine B. Show your work in the box.
$\square$

## 7 Tomasulo's Algorithm [60 points]

Consider an in-order fetch, out-of-order dispatch, and in-order retirement execution engine that employs Tomasulo's algorithm. This engine has the following characteristics:

- The engine has four main pipeline stages: Fetch (F), Decode (D), Execute (E), and Write-back (W).
- The engine can fetch one instruction per cycle, decode one instruction per cycle, and write back the result of one instruction per cycle.
- The engine has two execution units: 1) an adder to execute ADD instructions and 2) a multiplier to execute MUL instructions.
- The execution units are fully pipelined. The adder has two stages (E1-E2), and the multiplier has four stages (E1-E2-E3-E4). Execution of each stage takes one cycle.
- The adder has a two-entry reservation station, and the multiplier has a three-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.
- Full data forwarding is available, i.e., during the last cycle of the E stage, the tags and data are broadcast to the reservation station and the Register Alias Table (RAT). For example, an ADD instruction updates the reservation station entries of the dependent instructions in the E2 stage. So, the updated value can be read from the reservation station entry in the next cycle. Therefore, a dependent instruction can potentially begin its execution in the next cycle (after E2).
- The multiplier and adder have separate output data buses, which allow both the adder and the multiplier to update the reservation station and the RAT in the same cycle.
- An instruction continues to occupy a reservation station slot until it finishes the Write-back (W) stage. The reservation station entry is deallocated after the Write-back (W) stage.


### 7.1 Problem Definition

The processor is about to fetch and execute five instructions. Assume the reservation stations ( $R S$ ) are all initially empty, and the initial state of the register alias table ( $R A T$ ) is given below in Figure (a). Instructions are fetched, decoded, and executed as discussed in class. At some point during the execution of the five instructions, a snapshot of the state of the RS and the RAT is taken. Figures (b) and (c) show the state of the RS and the RAT at the snapshot time. A dash (-) indicates that a value has been cleared. A question mark (?) indicates that a value is unknown to you.

| Reg | Valid | Tag | Value |
| :---: | :---: | :---: | :---: |
| R0 | 1 | - | 1900 |
| R1 | 1 | - | 82 |
| R2 | 1 | - | 1 |
| R3 | 1 | - | 3 |
| R4 | 1 | - | 10 |
| R5 | 1 | - | 5 |
| R6 | 1 | - | 23 |
| R7 | 1 | - | 35 |
| R8 | 1 | - | 61 |
| R9 | 1 | - | 4 |

(a) Initial state of the RAT

| Reg | Valid | Tag | Value |
| :---: | :---: | :---: | :---: |
| R0 | 1 | $?$ | 1900 |
| R1 | 1 | $?$ | 82 |
| R2 | 1 | $?$ | 1 |
| R3 | 1 | $?$ | 45 |
| R4 | 0 | A | $?$ |
| R5 | 0 | F | $?$ |
| R6 | 1 | $?$ | 23 |
| R7 | 1 | $?$ | 35 |
| R8 | 0 | L | $?$ |
| R9 | 0 | B | $?$ |

(b) State of the RAT at the snapshot time

(c) State of the RS at the snapshot time

### 7.2 Questions

### 7.2.1 Dataflow Graph [40 points]

Based on the information provided above, identify the instructions and provide the dataflow graph below for the instructions that have been fetched. Please appropriately connect the nodes using edges and specify the direction of each edge. Label each edge with the destination architectural register and the corresponding Tag.
$\square$

### 7.2.2 Program Instructions [20 points]

Fill in the blanks below with the five-instruction sequence in program order. There can be more than one correct ordering. Please provide only one correct ordering. When referring to registers, please use their architectural names (R0 through R9). Place the register with the smaller architectural name on the left source register box.
For example, ADD R8 $\Leftarrow$ R1, R5.

$\square$

$\Leftarrow$ $\qquad$ $\square$
$\square$

$\square$
$\square$
$\square$
$\square$ $\Leftarrow$ $\qquad$ ,$\square$
$\square$
$\square$
$\square$ ,$\square$
$\qquad$

## 8 GPUs and SIMD [75 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.

The following code segments are run on a GPU. We assume that (1) A resides in memory and is shared by all threads, $(2)$ s resides in a register and is private to each thread, and (3) the code segments are correct (i.e., do not think about any correctness issues when answering this question).

A warp in the GPU consists of 32 threads, and there are 32 SIMD lanes in the GPU. Each thread executes a single iteration of the outermost loop (with index i). Assume that the data values of the array A are already in vector registers so there are no memory loads and stores in this program. (Hint: Notice that there are 4 instructions in each iteration of the outermost loop of both code segments.)

```
s = 1;
for (i = 0; i < 1024; i++) {
    for (j = 0; j < 10; j++) { // Inst. 1
        if (i % (2 * s) == 0) // Inst. 2
            A[i] += A[i + 1]; // Inst. 3
        s = s << 1; // Inst. 4
    }
}
```


## Code Segment 1

$s=512 ;$
for (i = 0; i < 1024; i++) \{
for ( $j=0 ; j<10 ; j++$ ) $\{1 /$ Inst. 1
if (i<s) // Inst. 2
A[i] +=A[i + s]; // Inst. 3
$\mathrm{s}=\mathrm{s} \gg 1 ; \quad / /$ Inst. 4
\}
\}

## Code Segment 2

Please answer the following questions.
(a) [5 points] How many warps does it take to execute these code segments?
$\qquad$
(b) [10 points] What is the SIMD utilization of the first iteration of the inner loop $(j=0)$ for Code Segment 1? Show your work. (Hint: The warp scheduler does not issue instructions when no thread is active).
$\square$
(c) [10 points] What is the SIMD utilization of the first iteration of the inner loop $(j=0)$ for Code Segment 2? Show your work. (Hint: The warp scheduler does not issue instructions when no thread is active).
$\square$
$\qquad$
(d) [15 points] What is the SIMD utilization of any iteration of the inner loop ( $0<=j<10$ ) for Code Segment 1? Show your work. (Hint: Derive an analytical expression, which may be piecewise).
(e) [15 points] What is the SIMD utilization of any iteration of the inner loop $(0<=j<10)$ for Code Segment 2? Show your work. (Hint: Derive an analytical expression, which may be piecewise).
$\qquad$
(f) [10 points] Is there any iteration $(0<=j<10)$ where both code segments have the same utilization? Explain your reasoning.
(g) [10 points] Which code is expected to run faster on a GPU? Explain your reasoning.

## 9 Branch Prediction [45 points]

You are given the following piece of code that iterates through two large arrays, j and k , each populated with completely (i.e., truly) random positive integers. The code has five branches (labeled B1, B2, B3, B4, and B5). When we say that a branch is taken, we mean that the code inside the curly brackets is executed. Assume that the code is run to completion without any errors or interruptions (i.e., there are no exceptions). For the following questions, assume that this is the only block of code that will ever be run on the machines, and that the loop condition branch is resolved first in the iteration (i.e., the if statements execute only after resolving the loop condition branch).

```
for(int i = 0; i < 1000; i++) { //B1
    //TAKEN PATH for B1
    if (i % 2=0) { //B2
        j[i] = k[i] * i; //TAKEN PATH for B2
    }
    if (i < 250) { //B3
        j[i] = k[i] - i; //TAKEN PATH for B3
    }
    if (i<500) { //B4
        j[i] = k[i] + i; //TAKEN PATH for B4
    }
    if (i >= 500) { //B5
        j[i] = k[i] / i; //TAKEN PATH for B5
    }
}
```

Listing 1: Application to evaluate.

You are given three machines whose components are identical in every way, except for their branch predictors.

- Machine A uses an always-taken branch predictor.
- Machine B uses one single-level global two-bit saturating counter branch predictor shared by all branches, which starts at Weakly Taken (2'b10).
- Machine C uses a per-branch two-bit saturating counter as its branch predictor. All counters start at Weakly Not Taken ( $2^{\prime}$ b01).

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
$\qquad$

Answer the following questions:

1. [15 points] What is the branch misprediction rate when the above piece of code runs on Machine A? Show your work.
$\square$
2. [15 points] What is the branch misprediction rate when the above piece of code runs on Machine B? Show your work.
$\square$
3. [15 points] What is the branch misprediction rate when the above piece of code runs on Machine C? Show your work.
$\qquad$

## 10 Caches [70 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 (2-, 4-, or 8 -way).
- Cache replacement policy (LRU or FIFO).
- Cache size ( 4 or 8 KiB ).

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 $\rightarrow$ Youngest) |  |  |  |  |  |  |  | Hit Rate |
| :---: | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :---: |
| 1. | 0 | 16 | 24 | 25 | 1024 | 255 | 1100 | 305 | $2 / 8$ |
| 2. | 31 | 65536 | 65537 | 131072 | 262144 | 8 | 305 | 1060 | $3 / 8$ |
| 3. | 262145 | 65536 | 4 |  |  |  |  |  | $2 / 3$ |

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

Based on what you observe, what are the following characteristics of the cache? Explain to get points.
(a) [20 points] Cache block size $(8,16,32,64$, or 128 B)?
$\square$
(b) [20 points] Cache associativity (2-, 4-, or 8-way)?
$\square$
$\qquad$
(c) [20 points] Cache replacement policy (LRU or FIFO)?
(d) [10 points] To identify the cache size ( 4 or 8 KiB ), you can access two addresses right after sequence 3 (i.e., the contents are the same as at the end of the third sequence) and measure the cache hit rate. Which two addresses would you choose? Explain your answer (there may be several correct answers).
$\square$

## 11 BONUS: Prefetching [25 points]

A runahead execution processor is designed with an unintended hardware bug: every other instruction in runahead mode is dropped by the processor after the fetch stage. Recall that the runahead mode is the speculative processing mode where the processor executes instructions solely to generate prefetch requests. All other behavior of the runahead mode is exactly as we described in lectures. When a program is executed, which of the following scenarios could happen compared to a runahead processor without the hardware bug and why? Circle YES if there is a possibility to observe the described behavior and explain in the box (either if you answer YES or NO). Assume that the program has no bug in it and executes correctly on the processor without the hardware bug.
(a) [8 points] The buggy runahead processor finishes the program correctly and faster than the nonbuggy runahead processor.

$$
Y E S \quad N O
$$

Why?
$\square$
(b) [8 points] The buggy runahead processor finishes the program correctly and slower than the nonbuggy runahead processor.

$$
Y E S \quad N O
$$

Why?
$\square$
(c) [9 points] The buggy runahead processor executes the program incorrectly.

$$
Y E S \quad N O
$$

Why?
$\qquad$

## 12 BONUS: Systolic Arrays [35 points]

A systolic array consists of $4 \times 4$ Processing Elements (PEs), interconnected as shown in Figure 1. The inputs of the systolic array are labeled as $H_{0}, H_{1}, H_{2}, H_{3}$ and $V_{0}, V_{1}, V_{2}, V_{3}$. Figure 2 shows the PE logic, which performs a multiply and accumulate MAC operation and saves the result to an internal register (reg). Figure 2 also shows how each PE propagates its inputs. We make the following assumptions:

- The latency of each MAC operation 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 values of all internal registers is zero.


Figure 1: PE array

Processing Element (PE)


Figure 2: Processing Element (PE)

Your goal is to use the example systolic array shown in Figure 1 to perform the convolution $(\circledast)$ of a 3 x 3 image (matrix $I_{3 x 3}$ ) with four 2 x 2 filters (matrices $A_{2 x 2}, B_{2 x 2}, C_{2 x 2}$, and $D_{2 x 2}$ ), to obtain four 2 x 2 outputs (matrices $W_{2 x 2}, X_{2 x 2}, Y_{2 x 2}$, and $Z_{2 x 2}$ ):

| $I_{00}$ <br> $I_{10}$ <br> $I_{20}$ | $\begin{aligned} & I_{01} \\ & I_{11} \\ & I_{21} \end{aligned}$ | $\begin{aligned} & I_{02} \\ & I_{12} \\ & I_{22} \end{aligned}$ | $(*)$ | $\begin{aligned} & A_{00} \\ & A_{10} \end{aligned}$ | $\begin{aligned} & A_{01} \\ & A_{11} \end{aligned}$ | $=$ | $\begin{aligned} & W_{00} \\ & W_{10} \end{aligned}$ | $\begin{aligned} & W_{01} \\ & W_{11} \end{aligned}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $I_{00}$ <br> $I_{10}$ <br> $I_{20}$ | $\begin{aligned} & I_{01} \\ & I_{11} \\ & I_{21} \end{aligned}$ | $\begin{aligned} & I_{02} \\ & I_{12} \\ & I_{22} \end{aligned}$ | $(\circledast)$ | $\begin{aligned} & B_{00} \\ & B_{10} \end{aligned}$ | $\begin{aligned} & B_{01} \\ & B_{11} \end{aligned}$ | $=$ | $\begin{aligned} & X_{00} \\ & X_{10} \end{aligned}$ | $\begin{aligned} & X_{01} \\ & X_{11} \end{aligned}$ |
| $\begin{aligned} & I_{00} \\ & I_{10} \\ & I_{20} \end{aligned}$ | $\begin{aligned} & I_{01} \\ & I_{11} \\ & I_{21} \end{aligned}$ | $\begin{aligned} & I_{02} \\ & I_{12} \\ & I_{22} \end{aligned}$ | $(\circledast)$ | $C_{00}$ $C_{10}$ | $C_{01}$ $C_{11}$ | $=$ | $Y_{00}$ $Y_{10}$ | $Y_{01}$ $Y_{11}$ |
| $\begin{aligned} & I_{00} \\ & I_{10} \\ & I_{20} \end{aligned}$ | $\begin{aligned} & I_{01} \\ & I_{11} \\ & I_{21} \end{aligned}$ | $\begin{aligned} & I_{02} \\ & I_{12} \\ & I_{22} \end{aligned}$ | $(*)$ | $\begin{aligned} & D_{00} \\ & D_{10} \end{aligned}$ | $\begin{aligned} & D_{01} \\ & D_{11} \end{aligned}$ | $=$ | $Z_{00}$ $Z_{10}$ | $Z_{01}$ $Z_{11}$ |

As an example, the convolution of the matrix $I_{3 x 3}$ with the filter $A_{2 x 2}$ is computed as follows:

- $W_{00}=I_{00} * A_{00}+I_{01} * A_{01}+I_{10} * A_{10}+I_{11} * A_{11}$
- $W_{01}=I_{01} * A_{00}+I_{02} * A_{01}+I_{11} * A_{10}+I_{12} * A_{11}$
- $W_{10}=I_{10} * A_{00}+I_{11} * A_{01}+I_{20} * A_{10}+I_{21} * A_{11}$
- $W_{11}=I_{11} * A_{00}+I_{12} * A_{01}+I_{21} * A_{10}+I_{22} * A_{11}$
$\qquad$

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

1. The input elements (from matrices $I_{3 x 3}, A_{2 x 2}, B_{2 x 2}, C_{2 x 2}$, and $D_{2 x 2}$ ) in the correct input ports of the systolic array $\left(H_{0}, H_{1}, H_{2}, H_{3}\right.$ and $\left.V_{0}, V_{1}, V_{2}, V_{3}\right)$. (Hint: If necessary, an input element can be concurrently streamed into several input ports of the array.)
2. The output values and the corresponding PE where the output elements (of matrices $W_{2 x 2}, X_{2 x 2}$, $Y_{2 x 2}$, and $Z_{2 x 2}$ ) are generated.
Fill the blanks only with relevant information.


