

DESIGN OF DIGITAL CIRCUITS (252-0028-00L), SPRING 2019
OPTIONAL HW 4: DATA FLOW, PIPELINING, TOMASULO'S ALGORITHM,
AND OUT OF ORDER EXECUTION

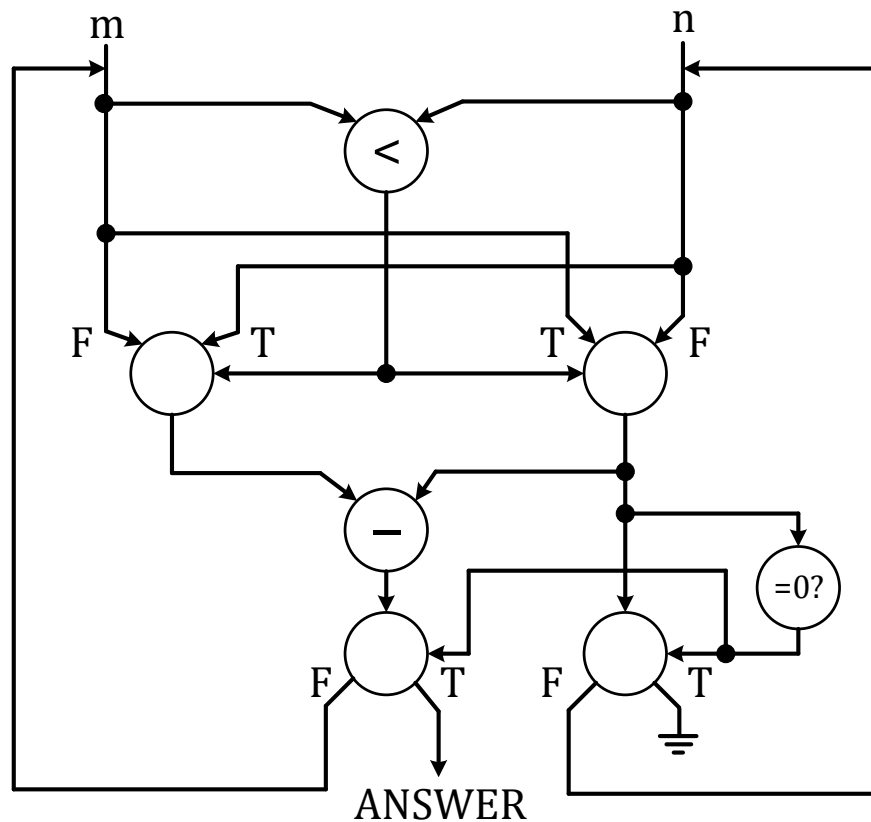
Instructor: Prof. Onur Mutlu

TAs: Mohammed Alser, Can Firtina, Hasan Hassan, Juan Gomez Luna, Lois Orosa, Giray Yaglikci



Released: Friday, April 12, 2019

1 Data Flow (Reverse Engineering)


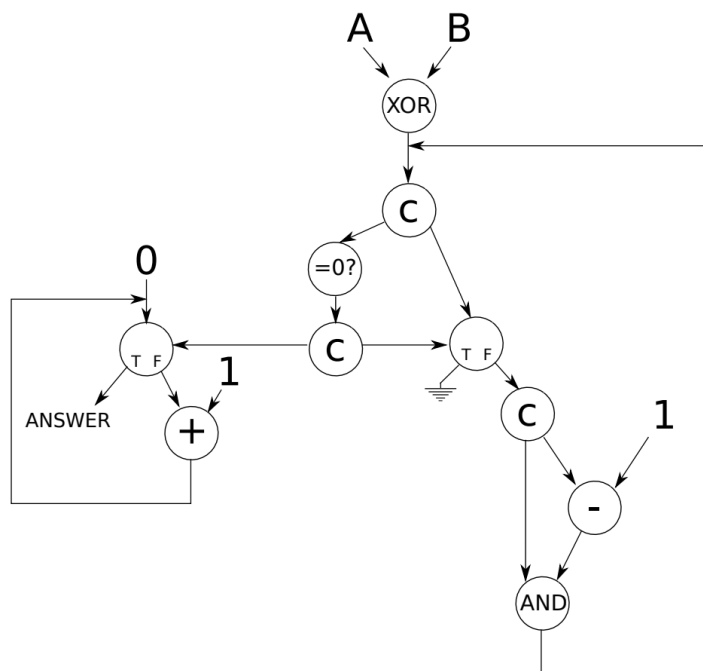
- (a) What does the following dataflow program do? Specify clearly in less than 10 words (one could specify this function in three words).



- (b) What does the following dataflow graph do (10 words or less)? (Hint: Identify what Module 1 and Module 2 perform.)



- (c) What does the following dataflow graph do (15 words or less)?
 (Note that the inputs, A and B , are non-negative integers.)



2 Pipelining I

Given the following code:

```
1 MUL R3, R1, R2
2 ADD R5, R4, R3
3 ADD R6, R4, R1
4 MUL R7, R8, R9
5 ADD R4, R3, R7
6 MUL R10, R5, R6
```

Calculate the number of cycles it takes to execute the given code on the following models:

Note 1: Each instruction is specified with the destination register first.

Note 2: Do not forget to list any assumptions you make about the pipeline structure (e.g., how is data forwarding done between pipeline stages)


Note 3: For all machine models, use the basic instruction cycle as follows:

- Fetch (one clock cycle)
- Decode (one clock cycle)
- Execute (MUL takes 6, ADD takes 4 clock cycles). The multiplier and the adder are not pipelined.
- Write-back (one clock cycle)


(a) A non-pipelined machine

(b) A pipelined machine with scoreboard and five adders and five multipliers without data forwarding


- (c) A pipelined machine with scoreboarding and five adders and five multipliers with data forwarding.



- (d) A pipelined machine with scoreboarding and one adder and one multiplier without data forwarding



- (e) A pipelined machine with scoreboarding and one adder and one multiplier with data forwarding



3 Pipelining II

Consider two pipelined machines implementing MIPS ISA, Machine I and Machine II:

Both machines have the following *five pipeline stages*, very similarly to the basic 5-stage pipelined MIPS processor we discussed in lectures, and *one ALU*:

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).

Machine I does *not* implement interlocking in hardware. It assumes all instructions are independent and relies on the compiler to order instructions such that there is sufficient distance between dependent instructions. The compiler either moves other independent instructions between two dependent instructions, if it can find such instructions, or otherwise, inserts **nops**. Assume internal register file forwarding (an instruction writes into a register in the first half of a cycle and another instruction can correctly access the same register in the next half of the cycle). Assume that the processor predicts all branches as always-taken.

Machine II implements data forwarding in hardware. On detection of a flow dependence, it forwards an operand from the memory stage or from the write-back stage to the execute stage. The load instruction (**lw**) can *only* be forwarded from the write-back stage because data becomes available in the memory stage but not in the execute stage like for the other instructions. Assume internal register file forwarding (an instruction writes into a register in the first half of a cycle and another instruction can access the same register in the next half of the cycle). The compiler does *not* reorder instructions. Assume that the processor predicts all branches as always-taken.

Consider the following code segment:

```
Copy: lw    $2, 100($5)
      sw    $2, 200($6)
      addi  $1, $1, 1
      bne   $1, $25, Copy
```

Initially, \$5 = 0, \$6 = 0, \$1 = 0, and \$25 = 25.

- (a) When the given code segment is executed on Machine I, the compiler has to reorder instructions and insert **nops** if needed. Write the resulting code that has minimal modifications from the original.

- (b) When the given code segment is executed on Machine II, dependencies between instructions are resolved in hardware. Explain when data is forwarded and which instructions are stalled and when they are stalled.

- (c) Calculate the *machine code size* of the code segments executed on Machine I (part (a)) and Machine II (part (b)).

- (d) Calculate the number of cycles it takes to execute the code segment on Machine I and Machine II.

(e) Which machine is faster for this code segment? Explain.

4 Tomasulo's Algorithm

Remember that Tomasulo's algorithm requires tag broadcast and comparison to enable wake-up of dependent instructions. In this question, we will calculate the number of tag comparators and size of tag storage required to implement Tomasulo's algorithm in a machine that has the following properties:

- 8 functional units where each functional unit has a dedicated separate tag and data broadcast bus
- 32 64-bit architectural registers
- 16 reservation station entries per functional unit
- Each reservation station entry can have two source registers

Answer the following questions. Show your work for credit.

(a) What is the number of tag comparators per reservation station entry?

(b) What is the total number of tag comparators in the entire machine?

(c) What is the (minimum possible) size of the tag?

(d) What is the (minimum possible) size of the register alias table (or, frontend register file) in bits?

(e) What is the total (minimum possible) size of the tag storage in the entire machine in bits?

5 Tomasulo's Algorithm (Reverse Engineering)

In this problem, we will give you the state of the Register Alias Table (RAT) and Reservation Stations (RS) for an out-of-order execution engine that employs Tomasulo's algorithm, as we discussed in lectures. Your job is to determine the original sequence of **four instructions** in program order.

The out-of-order machine in this problem behaves as follows:

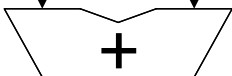
- The frontend of the machine has a one-cycle fetch stage and a one-cycle decode stage. The machine can fetch one instruction per cycle, and can decode one instruction per cycle.
- The machine executes *only* register-type instructions, e.g., $OP\ R_{dest} \leftarrow R_{src1},\ R_{src2}$.
- The machine dispatches one instruction per cycle into the reservation stations, in program order. Dispatch occurs during the decode stage.
- An instruction always allocates the first reservation station that is available (in top-to-bottom order) at the required functional unit.
- When an instruction in a reservation station finishes executing, the reservation station is cleared.
- The adder and multiplier **are not** pipelined. An add operation takes 2 cycles. A multiply operation takes 3 cycles.
- The result of an addition and multiplication is broadcast to the reservation station entries and the RAT in the writeback stage. A dependent instruction can begin execution in the next cycle after the writeback if it has all of its operands available in the reservation station entry.
- When multiple instructions are ready to execute at a functional unit at the same cycle, the oldest ready instruction is chosen to be executed first.

Initially, the machine is empty. Four instructions then are fetched, decoded, and dispatched into reservation stations. Pictured below is the state of the machine when the final instruction has been dispatched into a reservation station:


RAT

Reg	V	Tag	Value
R0	—	—	—
R1	0	A	5
R2	1	—	8
R3	0	E	—
R4	0	B	—
R5	—	—	—

ID	V	Tag	Value	V	Tag	Value
A	0	D	—	1	—	8
B	0	A	—	0	A	—
C	—	—	—	—	—	—



ID	V	Tag	Value	V	Tag	Value
D	1	—	5	1	—	5
E	0	A	—	0	B	—
F	—	—	—	—	—	—



- (a) Give the four instructions that have been dispatched into the machine, in program order. The source registers for the first instruction can be specified in either order. Give instructions in the following format: “opcode destination \leftarrow source1, source2.”

		\leftarrow		,	
		\leftarrow		,	
		\leftarrow		,	
		\leftarrow		,	

- (b) Now assume that the machine flushes all instructions out of the pipeline and restarts fetch from the first instruction in the sequence above. Show the full pipeline timing diagram below for the sequence of four instructions that you determined above, from the fetch of the first instruction to the writeback of the last instruction. Assume that the machine stops fetching instructions after the fourth instruction.

As we saw in lectures, use “F” for fetch, “D” for decode, “En” to signify the nth cycle of execution for an instruction, and “W” to signify writeback. You may or may not need all columns shown.


Cycle:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Inst.:																		
Inst.:																		
Inst.:																		
Inst.:																		
Inst.:																		

- (c) Finally, show the state of the RAT and reservation stations at the end of the **12th cycle** of execution in the figure below. Complete all blank parts.


RAT

Reg	V	Tag	Value
R0			
R1			
R2			
R3			
R4			
R5			

ID	V	Tag	Value	V	Tag	Value
A						
B						
C						



ID	V	Tag	Value	V	Tag	Value
D						
E						
F						



6 Out-of-Order Execution

In this problem, we consider an in-order fetch, out-of-order dispatch, and in-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 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 for executing ADD instructions and 2) a multiplier for executing 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 four-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 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.

6.1 Problem Definition

The processor is about to fetch and execute *six* instructions. Assume the *reservation stations (RS)* are all initially empty and the initial state of the *register alias table (RAT)* is given below in Figure (a). Instructions are fetched, decoded and executed as discussed in class. At some point during the execution of the six 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.


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

Initial state of the RAT

Reg	Valid	Tag	Value
R0	1	?	1900
R1	0	Z	?
R2	1	?	12
R3	1	?	3
R4	1	?	10
R5	0	B	?
R6	1	?	23
R7	0	H	?
R8	1	?	350
R9	0	A	?

Snapshot state of the RAT

ID	V	Tag	Value	V	Tag	Value
A	1	?	350	1	?	12
B	0	A	?	0	Z	?



ID	V	Tag	Value	V	Tag	Value
–	–	–	–	–	–	–
T	1	?	10	1	?	35
H	1	?	35	0	A	?
Z	1	?	82	0	H	?



6.2 (a) Data Flow Graph

Based on the information provided above, identify the instructions and complete the dataflow graph below for the six 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. *Note that you may **not** need to use all registers and/or nodes provided below.*

Register IDs:



Output

6.3 (b) Program Instructions

Fill in the blanks below with the six-instruction sequence in program order. 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`.

<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	,	<input type="text"/>
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	,	<input type="text"/>
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	,	<input type="text"/>
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	,	<input type="text"/>
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	,	<input type="text"/>
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	,	<input type="text"/>


7 Out-of-Order Execution - Reverse Engineering I

In this problem, we will give you the state of the Register Alias Table (RAT) and Reservation Stations (RS) for an out-of-order execution engine that employs Tomasulo's algorithm. Your job is to determine the original sequence of **five instructions** in program order.

The out-of-order machine in this problem behaves as follows:

- The frontend of the machine has a one-cycle fetch stage and a one-cycle decode stage. The machine can fetch one instruction per cycle, and can decode one instruction per cycle.
- The machine dispatches one instruction per cycle into the reservation stations, in program order. Dispatch occurs during the decode stage.
- An instruction always allocates the first reservation station that is available (in top-to-bottom order) at the required functional unit.
- When a value is captured (at a reservation station) or written back (to a register) in this machine, the old tag that was previously at that location is *not cleared*; only the valid bit is set.
- When an instruction in a reservation station finishes executing, the reservation station is cleared.
- Both the adder and multiplier are fully pipelined. An add instruction takes 2 cycles. A multiply instruction takes 4 cycles.
- When an instruction completes execution, it broadcasts its result. A dependent instructions can begin execution in the next cycle if it has all its operands available.
- When multiple instructions are ready to execute at a functional unit, the oldest ready instruction is chosen.

Initially, the machine is empty. Five instructions then are fetched, decoded, and dispatched into reservation stations. When the final instruction has been fetched and decoded, one instruction has already been written back. Pictured below is the state of the machine at this point, after the fifth instruction has been fetched and decoded:



- (a) Give the five instructions that have been dispatched into the machine, in program order. The source registers for the first instruction can be specified in either order. Give instructions in the following format: “opcode destination \leftarrow source1, source2.”

		\leftarrow		,	
		\leftarrow		,	
		\leftarrow		,	
		\leftarrow		,	
		\leftarrow		,	

- (b) Now assume that the machine flushes all instructions out of the pipeline and restarts fetch from the first instruction in the sequence above. Show the full pipeline timing diagram below for the sequence of five instructions that you determined above, from the fetch of the first instruction to the writeback of the last instruction. Assume that the machine stops fetching instructions after the fifth instruction.


As we saw in class, use “F” for fetch, “D” for decode, “En” to signify the nth cycle of execution for an instruction, and “W” to signify writeback. You may or may not need all columns shown.

Cycle:	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Instruction:														
Instruction:														
Instruction:														
Instruction:														
Instruction:														

Finally, show the state of the RAT and reservation stations after **10 cycles** in the blank figures below.

RAT

Reg	V	Tag	Value
R0			
R1			
R2			
R3			
R4			
R5			
R6			
R7			



Src 1

Src2

	Tag	V	Value	Tag	V	Value
A						
B						
C						



Src 1

Src2

	Tag	V	Value	Tag	V	Value
x						
y						
z						

8 Out-of-Order Execution - Reverse Engineering II

A five instruction sequence executes according to Tomasulo's algorithm. Each instruction is of the form ADD DR,SR1,SR2 or MUL DR,SR1,SR2. ADDs are pipelined and take 9 cycles (F-D-E1-E2-E3-E4-E5-E6-WB). MULs are also pipelined and take 11 cycles (two extra execute stages). An instruction must wait until a result is in a register before it sources it (reads it as a source operand). For instance, if instruction 2 has a read-after-write dependence on instruction 1, instruction 2 can start executing in the next cycle after instruction 1 writes back (shown below).

```
instruction 1      |F|D|E1|E2|E3|.....|WB|
instruction 2      |F|D|-|-|.....|-|E1|
```

The machine can fetch one instruction per cycle, and can decode one instruction per cycle.

The register file before and after the sequence are shown below.

	Valid	Tag	Value
R0	1		4
R1	1		5
R2	1		6
R3	1		7
R4	1		8
R5	1		9
R6	1		10
R7	1		11

	Valid	Tag	Value
R0	1		310
R1	1		5
R2	1		410
R3	1		31
R4	1		8
R5	1		9
R6	1		10
R7	1		21

- (a) Complete the five instruction sequence in program order in the space below. Note that we have helped you by giving you the opcode and two source operand addresses for the fourth instruction. (The program sequence is unique.)


Give instructions in the following format: "opcode destination \leftarrow source1, source2."

<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	, <input type="text"/>
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	, <input type="text"/>
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	, <input type="text"/>
MUL	<input type="text"/>	\leftarrow	R6	, R6
<input type="text"/>	<input type="text"/>	\leftarrow	<input type="text"/>	, <input type="text"/>

- (b) In each cycle, a single instruction is fetched and a single instruction is decoded.

Assume the reservation stations are all initially empty. Put each instruction into the next available reservation station. For example, the first ADD goes into "a". The first MUL goes into "x". Instructions remain in the reservation stations until they are completed. Show the state of the reservation stations at the end of cycle 8.

Note: to make it easier for the grader, when allocating source registers to reservation stations, please always have the higher numbered register be assigned to source2.



(c) Show the state of the Register Alias Table (Valid, Tag, Value) at the end of cycle 8.

	Valid	Tag	Value
R0			
R1			
R2			
R3			
R4			
R5			
R6			
R7			