

COMP3710, Computer Microarchitecture

Homework 1 (Weight: 20%)

Due date: October 5, 2022 (11:59 pm)

Total Points: 100

Important Instructions: (1) Write down the names and UIDs of each student in a group (if applicable) on the first page of your submission. (2) Submit the solution as a single pdf file.

Instructions for Q12: You can fill out the PowerPoint slide deck and convert it to a pdf document. You can then combine the pdf document with a second pdf file with responses to all other questions. Alternatively, you can copy and paste the main structures from the PowerPoint slide deck to your word document. And then submit a single pdf file. You can also print out the slide, fill the contents by hand, scan the document, and convert it to pdf.

Submission: Please submit a single Pdf file to the email below:
comp3710arch2022@gmail.com

(2 Points)

Q1. Compilers impact the performance of applications in different ways. For a program, compiler **X** results in a dynamic instruction count of 1 billion instructions, and an execution time of one second. A second compiler **Y** results in an execution time of 1.5 seconds, and a dynamic instruction count of 1.2 billion instructions. For a processor with a clock cycle time of one nano seconds, find the average CPI for each of the two programs.

(8 Points)

Q2. We are interested in adding a register-memory arithmetic instruction to the MIPS architecture. The new instruction exploits the I-format for load, store, and branch instructions. The new instruction has the label *ACCM* and employs an unused opcode in the ISA (the exact opcode is irrelevant). The semantics of the ACCM instruction is shown below:

Instruction: ACCM Rt, Const(Rs)

Interpretation: $\text{Reg}[\text{Rt}] = \text{Reg}[\text{Rt}] + \text{Mem}[\text{Reg}[\text{Rs}] + \text{Const}]$

MIPS I-Format (shown for convenience)

Opcode	Rs	Rt	Const
6 bits	5 bits	5 bits	16 bits

- Draw the datapath and control signals for a single-cycle implementation of the ACCM instruction. Your datapath should show the new components, control signals, multiplexers, and instruction labels. Your illustration must show every logic and memory element on the critical path.
- Identify the critical path for the ACCM instruction. Write the equation (like the lecture slides) for the critical path. For example, use t_{ALU} and t_{MEM} for the latency of ALU and memory. List all your assumptions.

(10 Points)

Q3. Consider the following instruction sequence:

	name	dst	src1	src2
i1:	sub	r4	r1	r0
i2:	lw	r2	16(r4)	
i3:	lw	r1	4(r1)	
i4:	and	r2	r4	r2
i5:	sw	r2	0(r4)	

- A. Suppose the pipelined in-order processor does not implement forwarding or hazard detection. As a programmer, your task is to insert nops to ensure correct execution. Insert nops to ensure correct execution.
- B. Suppose the processor manufacturer forgot to implement hazard detection. The processor still implements forwarding. Explain the consequences of executing the above code on the buggy processor.
- C. Now consider the following scenario: the processor does not implement forwarding. How should we change the hazard detection unit to ensure correct execution? List the conditions for detecting hazards. Explain the new input and output signals we need to add to our hazard detection unit. *Note: Use the above instruction sequence as an example to explain why each input/output signal is required.*
- D. Suppose there is an irrelevant instruction between i4 and i5. Is the store instruction exposed to a hazard? How can we resolve the hazard (if any)?

(10 Points)

Q4. Consider an analytics application running on top of the MIPS processor. A fraction of instructions in this application exposes a specific type of RAW hazard. We identify the type of RAW hazard by the stage that produces the result (EX or MEM) and the instruction that consumes the result (1st following instruction, 2nd instruction that follows, or both). The type of RAW hazard and the fraction of instructions are shown in the table below. Answer the questions below with the following assumptions: (1) A register write happens in the first half of the clock cycle and a register read happens in the second half, (2) CPI of the processor is one if there are no data hazards.

Assume stores are never followed by loads. All other hazards can be resolved by other tricks (RF read/write policy).

EX to 1 st Youngest	MEM to 1 st Youngest	EX to 2 nd Youngest	MEM to 2 nd Youngest	EX to 1 st and MEM to 2 nd	Other RAW Hazards
7%	18%	5%	10%	10%	10%

IF	ID	EX (No FW)	EX (Full FW)	EX (FW from EX/MEM only)	EX (FW from MEM/WB only)	MEM	WB
250 ps	200 ps	220 ps	250 ps	240 ps	230 ps	220 ps	200 ps

- A. What fraction of the cycles does the pipeline stalls with no forwarding?
- B. What fraction of the cycles does the pipelines stalls with full forwarding?
- C. What is the speedup with full forwarding versus no forwarding? *Note: Speedup is defined as the ratio of execution times with and without an optimization.*
- D. To avoid the complexity of large-input multiplexers, we need to decide if it is better to forward only from the EX/MEM pipeline register or the MEM/WB pipeline register. Which option would you choose to minimize data stall cycles? (Show your calculation)

(4 points, 2, 2)

Q5. Find the longest chain of dependent instructions in the following code sequence. If maximizing IPC is the goal, should a microarchitect consider a stall-on-use in-order pipeline over a stall-on-miss in-order pipeline?

	name	dst	src1	src2
i1:	add	r1	r1	r2
i2:	add	r1	r1	r3
i3:	sub	r1	r1	r4
i4:	load	r5	#0	r1
i5:	load	r7	#0	r8
i6:	add	r9	r5	r7

(12 Points)

Q6. Assume that a branch has the following sequence of taken (T) and not-taken (N) outcomes: **T,T,T,N,N,T,T,T,N,N,T,T,T,N,N**

- What is the prediction accuracy for a 2-bit counter (Smith predictor) for this sequence assuming an initial state is strongly taken?
- What is the minimum local history length needed to achieve perfect branch prediction for this branch outcome sequence?
- Draw the corresponding PHT and fill in each entry with one of T (predict taken), N (predict not taken), or X (does not matter).

(6 points)

Q7.

- A. Why does the register read stage must precede the issue stage in an out-of-order (OOO) processor (core) that uses an architectural register file (ARF) plus the reorder buffer to implement register renaming and hardware speculation?
- B. List the reasons for separate dispatch and issue stages in an out-of-order (OOO) processor core that implements dynamic scheduling?

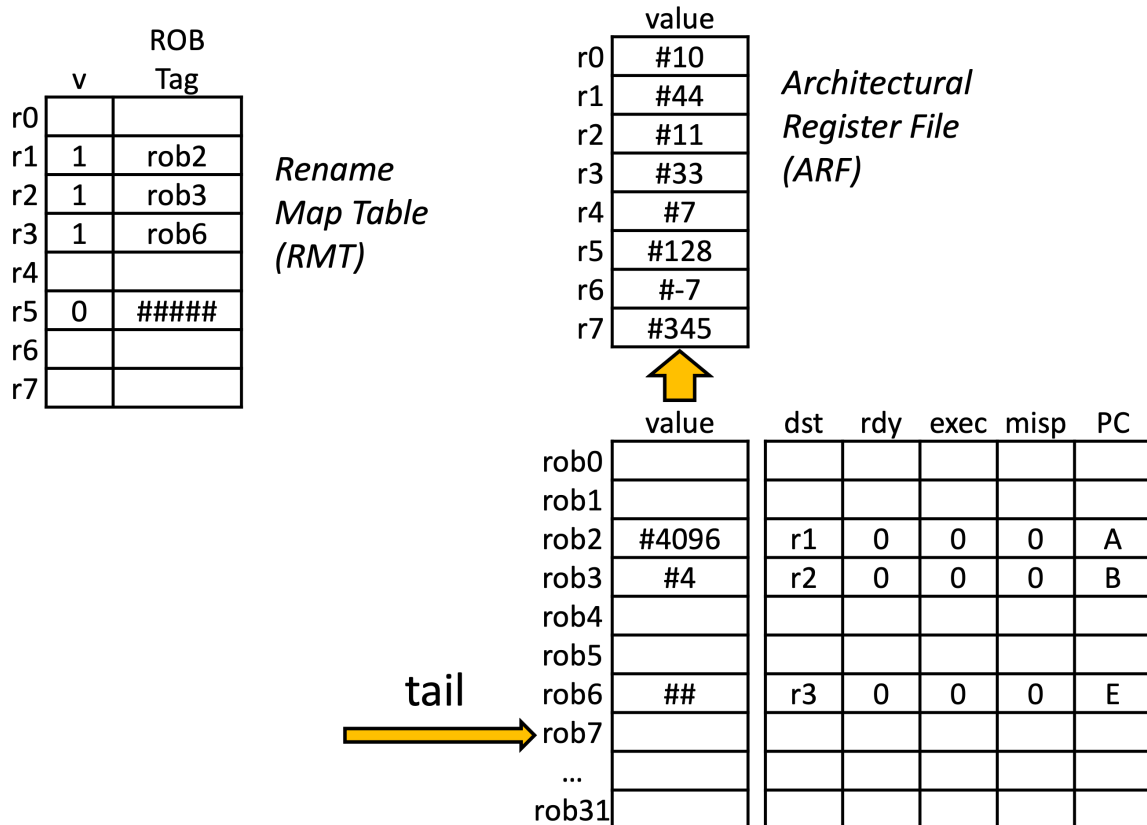
(6 points)

Q8. Indicate dependences and their types in the following instruction sequence. For each of the dependence types, explain the hazards that could result in the following microarchitectures: (1) single-cycle in-order, (2) pipelined in-order, and (3) pipelined out-of-order. *Assumption: No forwarding and hazard detection has been implemented yet.*

	opcode	dst	src1	src2
i1:	and	r3	r4	r5
i2:	and	r4	r3	r6
i3:	and	r3	r3	r4

(6 points)

Q9. In this question, consider an out-of-order pipeline with an architectural register file (ARF) and a reorder buffer (ROB). The ROB has 32 entries. The tail currently points at the eighth entry of the ROB (rob7). The head of the ROB is stalled for an additional 100 cycles. The state of the ARF, the rename map table (RMT), and the ROB are shown below. Rename the destination and source register specifiers in the instruction sequence below. Identify the dependences in the original and the renamed sequence. Draw the state of the RMT after the instruction sequence is renamed.



	name	dst	src1	scr2
i1:	add	r5	r2	r1
i2:	lw	r3	4(r5)	
i3:	lw	r2	0(r2)	
i4:	or	r3	r5	r3
i5:	sw	r3	0(r5)	

(8 points, 2.5, 2.5)

Q10. Briefly explain how we can add the following features to the CDC 6600 scoreboard. (1) Register renaming. (2) Hardware speculation. Start with the scoreboard design as we studied in the lectures and briefly explain the steps required to add the two features.

(8 points, 2, 2, 2)

Q11. The complexity of processor pipelines we have encountered in the lectures vary. We rank three different pipelines with increasing complexity as follows: (1) stall-on-miss (simple) (2) stall-on-use (moderately complex) (3) ARF+ROB (very complex). For each of the following scenarios, pick the simplest pipeline that would likely deliver the highest IPC. The in-order pipelines do not use branch prediction. The OOO pipeline uses a simple one-bit branch predictor.

1. **Scenario 1:** Frequent RAW hazards, infrequent branches, negligible WAR/WAW hazards, infrequent memory operations
2. **Scenario 2:** Infrequent RAW hazards, frequent hard-to-predict branches, frequent independent memory operations, frequency of WAR/WAW hazards is unknown
3. **Scenario 3:** Same as scenario 2, but easy to predict branches, and the frequency of WAR/WAW hazards is known to be very high

(20 points)

Q12. This question has an associated PowerPoint template slide (see the course website) that you need to fill three times (for three scenarios) and attach to your final pdf submission. Consider the following instruction sequence. Suppose we run this sequence on an ARF+ROB pipeline. Your task is to fill the contents of the RMT, the issue queue, ROB, the architectural register file, and the multiple-clock-cycle diagram (bottom of the slide) at three different points in the execution of the code sequence. These structures are marked **XXX** in the template slide.

Questions: Provide the contents of all structures marked **XXX** in the following cycles: **(1)** when the instruction i2 is in the register read stage **(2)** when the instruction i3 is in the issue stage and **(3)** when the instruction i4 is in the retire stage. In which cycle does the branch instruction sets/resets the misprediction bit in the ROB? Provide an itemized list of all the actions that take place in the pipeline during that cycle. (You should answer the last question in text, but you may fill out the PowerPoint template slide one more time for cycle # 13 and attach it to your pdf submission.)

Assumptions: Assume the processor uses the always-untaken branch prediction strategy. Also assume that branch i3 is not taken (on resolution). Assume i2 results in a data cache hit. The data cache hit latency is three cycles. Therefore, the load instruction takes one cycle for address calculation and three cycles for retrieving the value from the data cache. All other operations take one cycle to execute. The initial state of the RMT, the architectural register file, and the head/tail pointers of the ROB is shown on the slide. you should infer "when the instruction is in the issue stage" as when the instruction enters the issue stage. Same for the retire stage.

	name	dst	src1	src2
i1:	add	r2	r3	r4
i2:	load	r5	#16	r2
i3:	bnez		r5	i12
i4:	addi	r2	r7	#13