CS 152

Computer Systems Architecture

Mid-Term Examination

2020-02-10

There are six (6) questions adding up to 110 points
The grade in this exam is a total of 100 out of 110 points.

Name (Last, First)	UCINetID

Problem 1. (15 points)
A compiler designer has to optimized the performance of a compiler on the given machine. This machine uses a non-pipelined processor design. The hardware designer provides the following information.

Instruction class	CPI for this instruction
A	4
B	5
C	9
D	12

The compiler designer has to decide between two different code sequences for a certain high level statement. The table below gives the number of times each instruction class is used in each code sequence:

Sequence	A	B	C	D
Option 1	5	4	6	2
Option 2	6	8	5	1

(a) Find out the total number of clock cycles used for each code sequence **(5 points)**

(b) Find out the average CPI for each code sequences **(5 points)**

(c) Which option gives the better performance, by what ratio? **(5 points)**

Problem 2. (15 points)
Answer the following questions about the RISC-V ISA design

(a) Why does RISC-V have both SRA (Shift Right Arithmetic) and SRL (Shift Right Logical) instruction variants, but only SLL (Shift Left Logical) and no SLA (Shift Left Arithmetic)? **(5 points)**

(b) Why does LB (Load Byte) and LH (Load Half) have unsigned variants (LBU, LHU), but not SB (Store Byte) and SH (Store Half)? **(5 points)**

(c) Why does RISC-V implement a LUI (Load Upper Immediate) instruction? Why can't ADDI (Add immediate) be used for everything? (Hint: "addi x5, x5, 2048" gives an error) **(5 points)**

Problem 3. (20 points)
Consider the following three-stage pipeline design, with separate instruction memory (iMem) and data memory (dMem) interfaces.

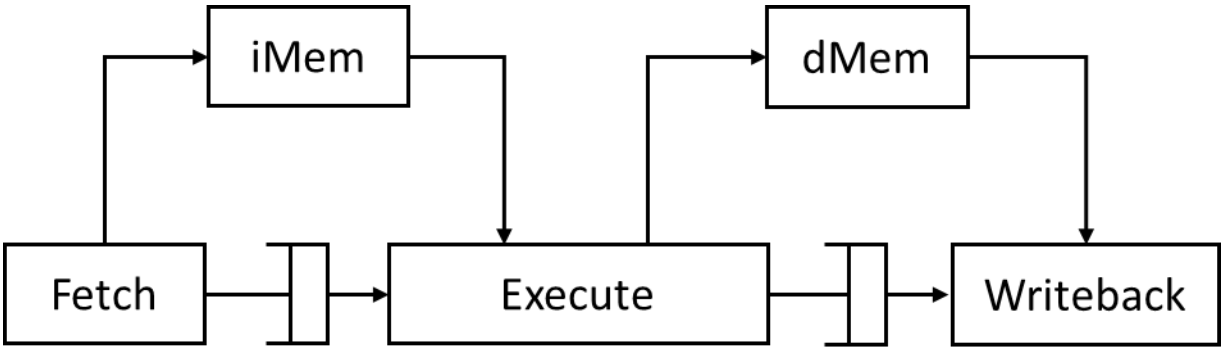

(a) Would you say this is a good pipeline design for performance purposes? Why or why not? **(5 points)**

(b) Can you build a design with a smaller number of pipeline stages? If so, which stages would you merge for performance? If not, why not? **(5 points)**

(c) Compare the above design against a four-stage pipeline which separates the execute stage into two stages: decode and execute. Assuming the following latency for each pipeline stages, which design would have higher sustained performance, and by what ratio? (Assuming ideal CPI of 1) **(10 points)**

Stage	Fetch	Decode	Execute	Decode+Execute	Writeback
Latency (ns)	4	5	2	6	3

Problem 4. (25 points)
Consider the following five-stage pipelined processor design.

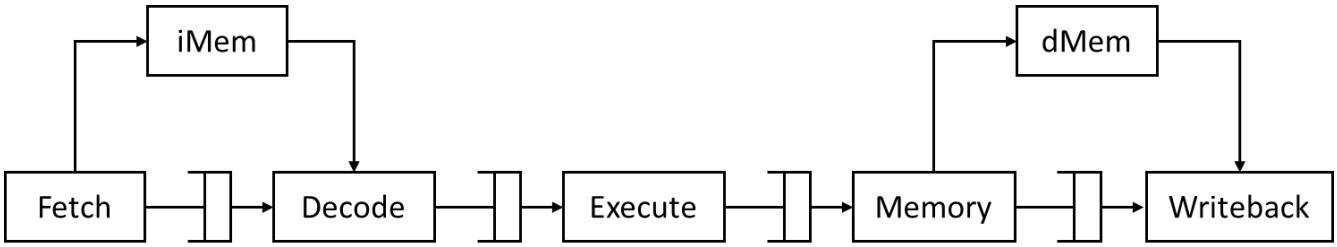

(a) Briefly describe which stage or stages will need to be modified (and how) to add an arithmetic instruction, say "multiplication" which takes two registers as operand and stores it in another register. **(5 points)**

(b) Briefly describe which stage or stages will need to be modified (and how) to handle control hazards using a branch predictor. Assume a dynamic branch predictor will be given. **(5 points)**

(c) This design right now has no data hazard handling. Given this design, what would be the data stored in t3 after executing the following instruction sequence? Is it different from the expected output of the instruction sequence? If so, what is different and why?
(Assume register writes are visible to decode after writeback cycle) **(10 points)**

```
addi t0, zero, 1
addi t1, zero, 1
addi t2, t0, 1
addi zero, zero, 0
addi t3, t0, t2
```

(d) Briefly describe forwarding and its purpose in handling data hazards. Include which stages need to be modified. Is there a reason we won't forward from all of the following stages: "Execute", "Memory", and "Writeback" **(5 points)**

Problem 5. (20 points)
If we change the load/store instructions to only use a register as input (without an offset), these instructions no longer need to use the ALU. As a result, Memory and Execute stages (See pipeline in Problem 4) can be overlapped, and the pipeline has only four stages.
(a) How will the reduction of pipeline stages affect the cycle time? **(10 points)**

(b) Given a generic program (such as the SPEC benchmarks), how will this change effect the overall performance of this processor? **(10 points)**

Problem 6. (15 points)
Bonus questions
(a) Explain the difference between modern CISC processors such as x86 and RISC processor such as RISC-V **(5 points)**

(b) What kind of programs would superscalar processors not be effective on? **(5 points)**

(c) We are running a benchmark program on a new, pipelined processor design. The distribution of instruction types in the benchmark, as well as the CPI of each instruction, is as follows.

Type	ALU	Branch	Load/Store
Percentage	60	25	15
CPI	1	3	5

We cannot calculate the average CPI of this benchmark on this processor just based on this information. Describe why not. **(5 points)**