

**Instructions:**

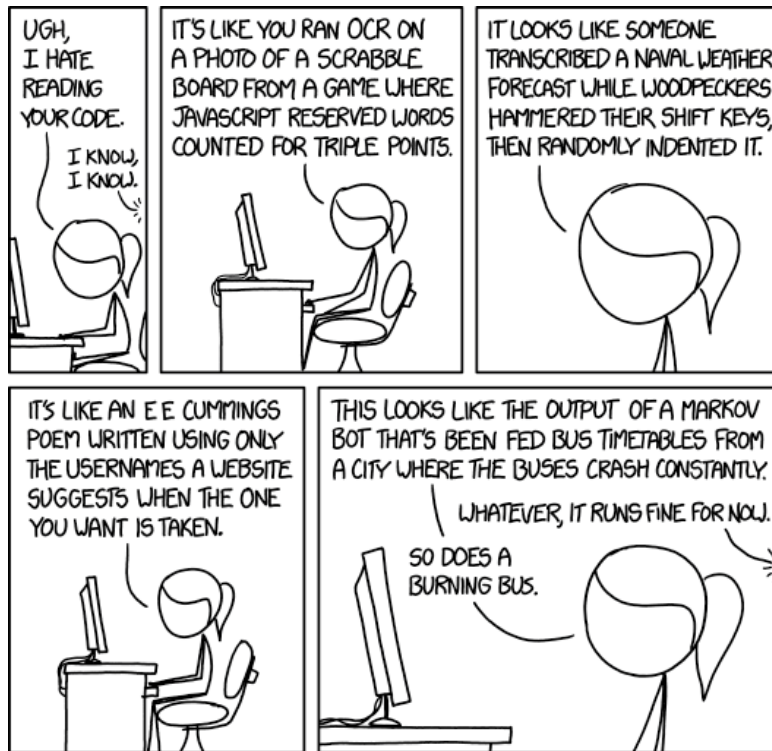
- Print your name in the space provided below.
- This examination is closed book and closed notes, aside from the permitted one-page formula sheet.
- No calculators or other computing devices may be used. The use of any such device will be interpreted as an indication that you are finished with the test and your test form will be collected immediately.
- Answer each question in the space provided. If you need more space to answer a question, you are probably thinking about it the wrong way.
- If you want partial credit, justify your answers, even when justification is not explicitly required.
- There are 6 questions, some with multiple parts, priced as marked. The maximum score is 100.
- When you have completed the test, sign the pledge at the bottom of this page, sign your fact sheet, and turn in the test and fact sheet.
- Note that failing to return this test, and discussing its content with a student who has not taken it are violations of the Honor Code.

**Do not start the test until instructed to do so!**

Name Solution  
printed

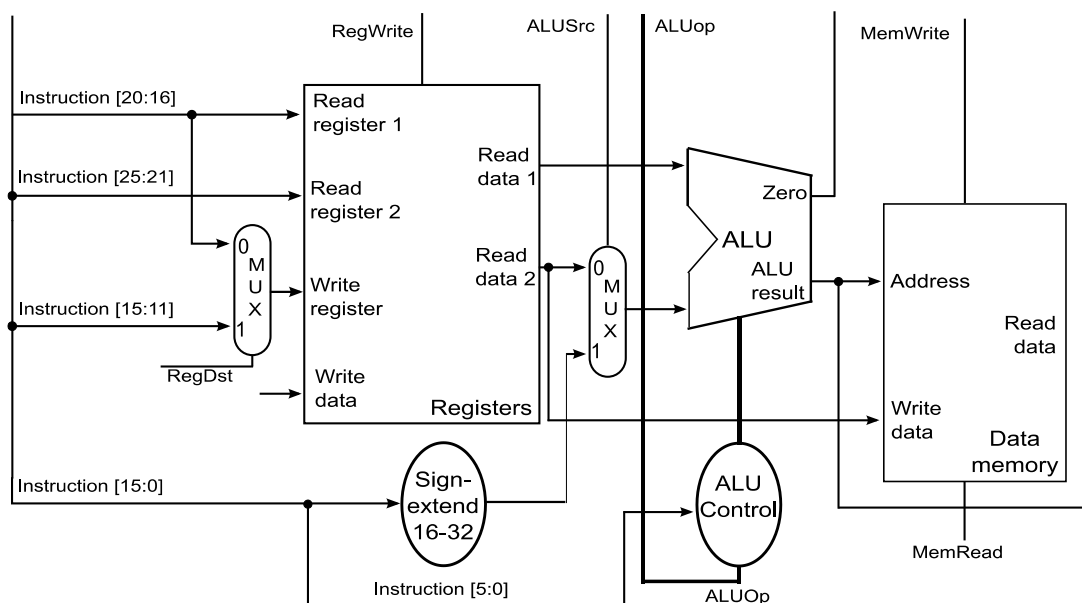
**Pledge:** On my honor, I have neither given nor received unauthorized aid on this examination.

\_\_\_\_\_  
*signed*



xkcd.com

1. Suppose that a buggy implementation of the MIPS data path miswires the handling of the Read register 1 input and the Read Register 2 inputs to the Registers unit as shown below:



The datapath supports execution of any sequence of add, sub, and, or, slt, lw, sw, beq, and j instructions.

Everything else in the datapath is implemented as shown on the diagram in the Supplement. All the control signals would be set as discussed in class. For each question below, we will assume that registers and memory locations hold the values shown in the tables above.

The error has several effects:

- the values taken from the register file are swapped (i.e., Read data 1 is now what should have been the right operand, and Read data 2 is now what should have been the left operand); for an instruction that sends both values to the ALU, this means the ALU inputs have been reversed, and that will break any arithmetic operation that is not commutative; for sw, the value from the RS register will be written to memory, instead of the value from the RT field, and the address will be computed from the RT field, not the RS field
- for lw and sw, the address will be computed using the RS field, not the RT field
- for R-type instructions, the Write register is still specified correctly; for lw, the Write register is now specified by bits 25:21 (i.e., the RS field instead of the RT field), so lw will write to the wrong register (unless the RS and RT fields specify the same register)

- a) [6 points] Which, if any, of the supported instructions would always be executed correctly on the buggy datapath shown above? No explanation is needed.

add, and, or  
beq, j

these operations are commutative R-type  
beq depends on whether the difference is 0, which is unaffected by the reordering; j does not depend on the values in registers; neither writes to a register

- b) [10 points] Suppose that the Registers and Data memory initially hold the values shown in the tables below. How would the contents of the Registers and Data memory be changed, if the following instruction was executed? Justify your answer.

```
lw    $t2, 1000($t1)
```

The instruction becomes: `lw $t2, 1000($t2)`

So:  $\$t2 \leftarrow \text{Mem}[\$t2 + 1000]$   
 $\$t2 \leftarrow \text{Mem}[3000]$   
 $\$t2 \leftarrow 30$

Data memory is unaffected since MemWrite is 0 for lw.

Register	value
\$t0	0
\$t1	1000
\$t2	2000
\$t3	3000

Address	value
0	segfault
1000	10
2000	20
3000	30
4000	40
5000	50
6000	60
7000	70

2. Suppose that when a `beq` instruction is executed the control signals RegDst, MemRead, MemtoReg, MemWrite, RegWrite, ALUOp, and Jump are set to values that would be correct if a `sw` instruction were being executed. And, suppose that the Branch and ALUSrc signals are set properly (for a `beq` instruction). Suppose the following instruction was executed:

```
beq    $t3, $t1, btarg    # btarg is the label for some instruction
```

- a) [8 points] Under what conditions would the branch still be taken? Explain.

Branch is set to 1; therefore the branch will be taken iff the ALU sets Zero to 1.

The ALU will compute  $\$t3 + \$t1$ , since ALUOp is set for sw, which performs an add.

Therefore, the branch will be taken iff  $\$t1 + \$t3 == 0$ .

The question is specifically about the conditions under which the branch will be taken, given the described changes in the way the control signals are set.

- b) [5 points] Could the execution of the `beq` instruction lead to any unintended changes be made to value(s) in Data memory? If so, describe what could happen. If not, explain why not.

MemWrite is set to 1 for sw instructions, so there WILL be a write to Data memory, or a segfault (if the address is not legal).

The ALU computes  $\$t1 + \$t3$ , which is sent to Data memory as the address to be accessed.

The datapath still sends Read register 2,  $\$t1$ , to the Write data input on the Data memory.

So:  $\text{Mem}[\$t1 + \$t3] \leftarrow \$t1$

- c) [5 points] Could the execution of the `beq` instruction lead to any unintended changes be made to value(s) in Registers? If so, describe what could happen. If not, explain why not.

RegWrite is set to 0 for sw instructions, so there cannot be a write to any of the registers.

3. The following questions refer to the single-cycle datapath on the Supplement to the test. They refer only to instructions supported by the datapath.
- a) [9 points] What value should the MemtoReg signal have for each of the instructions supported in this datapath? Be sure to identify any don't-care cases.

MemtoReg	Instructions	Reason
0	R-type instructions (add, sub, and, or, slt)	R-type instructions send a value computed by the ALU to a register
1	lw	lw instructions send a value read from Data memory to a register
D/C	beq, sw, j	For each of these, RegWrite is set to 0, so it doesn't matter what value is sent to Write data on the Registers unit (and neither choice is "correct").

- b) [9 points] For every instruction that is executed, the datapath appears to read the contents of Read register 1 and Read register 2, which are sent to the Read data 1 and Read data 2 lines, respectively.

For which instructions would the data read from Read data 1 not be needed? Explain why.

Only j, which does not depend at all on values from registers.

Every R-type instruction, and beq instructions, send Read data 1 to the ALU for an operation.

lw and sw instructions use the value from Read data 1 to compute an address.

Pick one of the instructions from your answer above, and explain why the fact that this unnecessary data is read, for that instruction, does not lead to any unintended results. Be complete and precise.

For a j instruction:

- RegWrite and MemWrite are set to 0, so j instructions do not store data anywhere.
- MemRead is set to 0, so no location in Data memory is read
- Jump is set to 1, so the branch target address is not used.

4. Assume an instruction set with five categories of instructions and a hardware implementation of this instruction set. In this implementation the instructions in each category take the number of cycles shown in the table below:

Category	A	B	C	D	E
CPI	2	3	1	4	5

- a) [10 points] A certain program executes the following mix of instructions from the five categories:

Category	A	B	C	D	E
CPI	30%	20%	20%	?	?

Note that the percentages of instructions in the D and E categories are unknown. What is the minimum CPI and what is the maximum CPI that this hardware implementation can achieve for the program above? Explain your answers.

The percentages for A, B, and C instructions are set. All we can do is consider the effect of different percentages of D and E instructions; there's no reason to assume that the program MUST use any instructions of either type (so 0% is valid), but the percentages still MUST add up to 100.

To minimize the average CPI, we need to avoid type-E instructions, since they take 5 cycles; hence use 30% type-D and 0% type-E:

$$\text{Min\_CPI} = 0.30 \cdot 2 + 0.20 \cdot 3 + 0.20 \cdot 1 + 0.30 \cdot 4 = .6 + .6 + .2 + 1.2 = 2.6$$

To maximize the average CPI, we need to avoid type-D instructions, since they take 4 cycles; hence use 0% type-D and 30% type-E:

$$\text{Min\_CPI} = 0.30 \cdot 2 + 0.20 \cdot 3 + 0.20 \cdot 1 + 0.30 \cdot 5 = .6 + .6 + .2 + 1.5 = 2.9$$

- b) [6 points] A marketing team from the company that sells the hardware described above claims that for a certain AI program their hardware can achieve a throughput of 1.2 instructions per cycle. Do you believe them? Answer with yes or no and explain why.

Suppose the throughput is 1.2 instructions/cycle.

Then the average CPI (cycles/instruction) must be  $1/1.2$  or  $5/6$  cycles per instruction.

That is a neat trick, since every instruction takes at least 1 cycle.

There's no suggestion this refers to the program described in part a).

5. Consider the following attributes of possible (hypothetical) designs for a pipeline (not the MIPS pipeline):

**A:** 400ps clock cycle, 5 stages

**B:** 300ps clock cycle, 8 stages

Both designs are intended for processors that employ the same machine language design.

There was a LOT of confusion here... as the instructions in part a) say, the latency of an instruction is the time from fetch to completion; in other words, the total time it spends going through the pipeline. This is simply the product of the cycle length and the number of stages.

- a) [4 points] Calculate the ideal instruction latency (time from fetch to completion) for design A.

$$\text{Latency} = 400\text{ps} * 5 = 2000\text{ps}$$

- b) [4 points] Calculate the ideal instruction latency for design B.

$$\text{Latency} = 300\text{ps} * 8 = 2400\text{ps}$$

Throughput is the rate at which instructions are completed; under ideal conditions (no stalls), a pipeline will complete one instruction per clock cycle.

- c) [4 points] Calculate the ideal throughput for design A (assuming an infinite sequence of instructions).

Once the pipeline is full, we complete 1 instruction per 400ps.

- d) [4 points] Calculate the ideal throughput for design B (assuming an infinite sequence of instructions).

Once the pipeline is full, we complete 1 instruction per 300ps.

Performance is dominated by throughput, not latency.

- e) [6 points] Which design offers better performance? Explain.

B yields lower execution time in the long run. Suppose we execute N instructions, where  $N > 7$  (so both pipelines reach fullness):

$$\text{Time}_A = 1600 + 400N \text{ but } \text{Time}_B = 1800 + 300N, \text{ so B is faster if } N > 7.$$

6. [10 points] Suppose that a program executes in 100 seconds, spending 40 seconds on floating-point operations, 20 seconds on integer operations, 25 seconds on branch operations, and 15 seconds on other operations.

Suppose it's possible to use integer operations to simulate floating-point operations, and that integer operations, on average, take half as long to execute as floating-point operations. However, the number of integer operations would be 50% greater than the number of floating-point operations they were replacing.

Use Amdahl's Law to determine how long it would take to execute the program if we replaced all the floating-point operations with integer operations.

$$\text{Time\_after} = \text{Time\_unaffected} + \text{Time\_affected} / \text{Improvement\_factor}$$

$$= 60 + 40 / (2 * 2/3)$$

$$= 60 + 40 / (4/3) = 60 + 30 = 90 \text{ seconds}$$

The question DOES specify that you should use Amdahl's Law, so full credit required some indication of that. There were a number of ways to approach the question.