

**Instructions:**

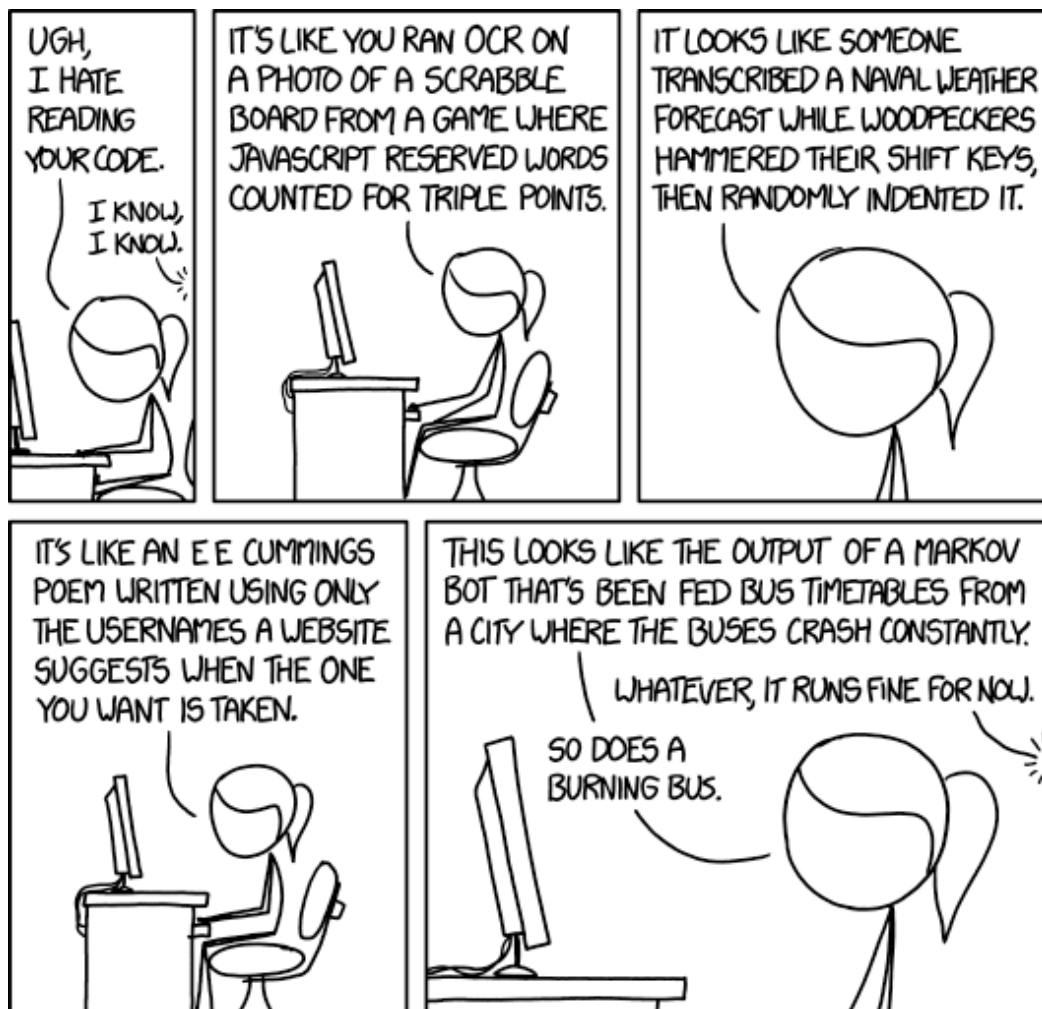
- Print your name in the space provided below.
- This examination is closed book and closed notes, aside from the permitted one-page formula sheet.
- No calculators or other computing devices may be used. The use of any such device will be interpreted as an indication that you are finished with the test 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. The single-cycle MIPS datapath design used one stage with a clock cycle of 800 ps. The MIPS pipeline had five stages, with a clock cycle of 200 ps. The pipelined version achieved better performance than the single-cycle version by improving instruction throughput. Let's consider some alternative outcomes.

Suppose that the MIPS designers managed (somehow) to produce a feasible pipeline design had resulted in ten stages, with a clock cycle of 180 ps.

- a) [6 points] How does the latency for a single instruction on this new design compare to the latency on the original, single-cycle design, and on the five-state pipeline discussed in class?

New design: $\text{AvgCPI} = 10 * 180 = 1800 \text{ ps}$

SCD: $\text{AvgCPI} = 800 \text{ ps}$

Old Pipeline: $\text{AvgCPI} = 5 * 200 = 1000 \text{ ps}$

So, the new design leads to much greater latency than either of the older designs.

- b) [6 points] Assuming an infinite sequence of instructions, what speedup will this design achieve when compared to the original, single-cycle design? Show calculations to support your conclusion.

$$\text{Speedup} = \lim_{n \rightarrow \infty} \frac{800n}{1620 + 180n} = \frac{800}{180} = \frac{40}{9} \approx 4.44$$

- c) [6 points] Assuming an infinite sequence of instructions, what speedup will this design achieve when compared to the five-stage pipeline design discussed in class? Show calculations to support your conclusion.

$$\text{Speedup} = \lim_{n \rightarrow \infty} \frac{800 + 200n}{1620 + 180n} = \frac{200}{180} = \frac{10}{9} \approx 1.11$$

2. Two processors, which support the same ISA, support three classes of instructions, which have the following CPI:

Instruction class	P1	P2
arithmetic/logical	5	6
load/store	80	60
branch	7	6

- a) [10 points] The dynamic instruction count for a particular machine language program includes 40% arithmetic/logical instructions, 30% load/store instructions, and 30% branch instructions. In other words, these percentages reflect the executable from a runtime point of view.

For each processor, what would be the average CPI for this executable? Justify your conclusions precisely.

$$\text{AvgCPI}_1 = 0.40 * 5 + 0.30 * 80 + 0.30 * 7 = 2.0 + 24.0 + 2.1 = 28.1$$

$$\text{AvgCPI}_2 = 0.40 * 6 + 0.30 * 60 + 0.30 * 6 = 2.4 + 18.0 + 1.8 = 22.2$$

- b) [6 points] What do your answers to part a) tell you about the relative performance of P1 and P2, when executing this particular program? Explain.

This tells me absolutely nothing about the relative performance of the program on the two processors. Performance is defined by execution time, and to compute that we would need to know

- the number of instructions to be executed (the actual dynamic count)
- the cycle length for the clock on each processor

3. Pseudo-instructions give MIPS a richer set of assembly language instructions than those supported directly by the hardware. Write an implementation for each of the following pseudo-instructions, using only the MIPS32 instructions listed at the bottom of this page (those are sufficient). You must include comments explaining the logic of your solutions. If you need to use any "extra" registers, you may only use the register \$at.

a) [8 points] `addfm $rd, ($rs), ($rt)`
`# GPR[rd] = Mem[GPR[rs]] + Mem[GPR[rt]]`

```

lw    $rd, ($rs)      # load Mem[GPR[rs]] into rd
lw    $at, ($rt)      # load Mem[GPR[rt]] into at
add   $rd, $rd, $at    # add Mem[GPR[rt]] to rd

```

b) [10 points] `swapifsmaller $rs, $rt`
`# if (GPR[rs] < GPR[rt]) swap the values in GPR[rs] and GPR[rt]`

```

slt    $at, $rs, $rt    # at = (rs < rt)?
beq    $at, $zero, nope  # if no, do nothing
addi   $at, $rs, $zero   # save rs in at
addi   $rs, $rt, $zero   # copy rt into rs
addi   $rt, $at, $zero   # copy backup of rs into rt
nope:

```

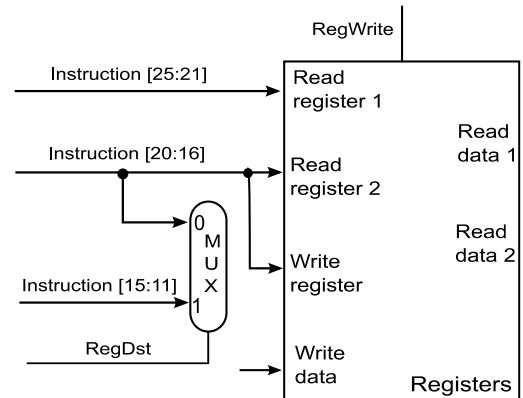
```

add    rd, rs, rt        # GPR[rd] <-- GPR[rs] + GPR[rt]
addi   rd, rs, imm16     # GPR[rd] <-- GPR[rs] + imm16      (sign extended)
and    rd, rs, rt        # GPR[rd] <-- GPR[rs] AND GPR[rt] (bitwise)
beq    rs, rt, label     # if GPR[rs] == GPR[rt], jump to label
bne    rs, rt, label     # if GPR[rs] != GPR[rt], jump to label
lw     rt, imm16(rs)     # GPR[rt] = Mem[GPR[rt]] + imm16
or     rd, rs, rt        # GPR[rd] <-- GPR[rs] OR GPR[rt]  (bitwise)
sub    rd, rs, rt        # GPR[rd] <-- GPR[rs] - GPR[rt]
sll    rd, rt, sa        # GPR[rd] <-- GPR[rs] << sa      (logical shift)
slt    rd, rs, rt        # GPR[rd] <-- (GPR[rs] < GPR[rt] ? 1 : 0)

```

(imm16 denotes a 16-bit immediate value)

4. Suppose that a buggy implementation of the MIPS datapath miswires the selection of the Write register number:



Everything else in the datapath operates normally.

- a) [8 points] Describe, in detail, how this error would affect the execution of the following instruction:

```
add $t0, $t1, $t2
```

The sum of \$t1 and \$t2 is computed correctly, but it is stored to \$t2, not \$t0.

- b) [8 points] Describe, in detail, how this error would affect the execution of the following instruction:

```
sw $s0, 0($s1)
```

Since sw does not write a value to any register, the error has no effect whatsoever.

5. [16 points] Suppose the following sequence of instructions was executed on the preliminary pipeline design derived in class, without any hardware to detect or handle hazards, as shown on the supplement:

```

sub    $t4, $t1, $t2    # 1
slt    $s0, $t4, $t2    # 2
sw     $t4, 0($t2)      # 3
add    $t2, $t4, $t2    # 4
or     $t1, $t2, $t4    # 5
lw     $t2, 0($t2)      # 6

```

Due to the limitations of the preliminary design, a number of different logical errors will occur. These may include possibly incorrect values being used as input to an instruction, possibly incorrect values being generated as output from an instruction, and other kinds of errors.

For each logical error that may occur, clearly identify the relevant instruction whose execution might be incorrect, and identify the register(s), if any, that are involved in the error, and describe the nature of the error.

Instruction #	Description of error
#1	no issues
#2	will read a "stale" value from \$t4 will compute incorrect value for \$s0
#3	will read a stale value from \$t4 will store incorrect value to Mem[GPR[\$t2]]
#4	no issues
#5	will read a stale value from \$t2 will compute incorrect value for \$t1
#6	will read a stale value from \$t2 will read value for \$t2 from incorrect address

6. [16 points] Consider the following proposed instruction for the simplified single-cycle MIPS32 datapath discussed in class:

`sas $rd, $rt, $rs # GPR[rd] = Mem[rs + rt]`

sas	rs	rt	rd	ignored	ignored
31 30 29 28 27 26	25 24 23 22 21	20 19 18 17 16	15 14 13 12 11	10 9 8 7 6	5 4 3 2 1 0

The instruction adds the values in registers `$rs` and `$rt` to obtain an address, and stores the contents of that address in register `$rd`.

Of course, this could already be accomplished by a sequence of supported instructions, but that would require more than one clock cycle. Supporting this new instruction would also require modifying the internals of the **Control** unit to recognize the `opcode` field for the new instruction, but assume that's easily accomplished (and that there is a currently-unused opcode that we can use for the new instruction).

Haskell Hoo IV, who proposed the new instruction, insists that it can be added to the current datapath design (as shown on the datapath diagram) with no changes other than to the internals of the **Control** unit (to recognize the new opcode). That is, there will be no need to add any new hardware to the datapath, nor will there be a need for any new control signals.

If Haskell Hoo IV is correct, explain how the eight existing control signals (excluding `ALUOp`) would need to be set. Be sure to consider (and indicate) if any of the signals are don't-cares.

If Haskell Hoo IV is incorrect, describe at least one hardware modification that must be made to the existing datapath in order to support `sas`, and explain why that modification is necessary. (Do not show control signal settings in that case.)

Signal	Value
RegDst	1
Jump	0
Branch	0
MemRead	1
MemtoReg	1
MemWrite	0
ALUSrc	0
RegWrite	1