Virginia IIII Tech

Instructions:

• Print your name in the space provided below.

Solution

- 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

printed

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

sígned



xkcd.com

CS 2506 Computer Organization II

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 three stages, with a clock cycle of 300 ps.

a) [6 points] How does the <u>latency</u> for a single instruction on this new design compare to the latency on the five-stage pipeline discussed in class?

The latency for the new design is 3 * 300 = 900 ps, versus 1000 ps for the five-stage design.

Latency is the time from the beginning of the fetch until the completion of execution.

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.

The single-cycle design completes one instruction per 800 ps; the three-stage design completes one instruction per 300 ps. So, the speedup is

800 ps / 300 ps = 8 / 3 or about 2.67

More precisely:

$$\lim_{n \to \infty} \frac{800n}{600 + 300n} = \lim_{n \to \infty} \frac{800}{300} = \frac{8}{3}$$

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.

Versus the five-stage pipeline, the speedup will be 200 ps / 300 ps = 2 / 3 or about 0.67.

(That's actually slower, of course...)

More precisely:

$$\lim_{n \to \infty} \frac{800 + 200n}{600 + 300n} = \lim_{n \to \infty} \frac{200}{300} = \frac{2}{3}$$

CS 2506 Computer Organization II

- 2. Use Amdahl's Law to justify your answers to the following questions.
 - a) [8 points] Suppose that a program spends 40% of its execution time performing integer arithmetic operations, 35% performing floating point operations, and 25% performing I/O operations. Given the advanced state of hardware for arithmetic computations, it is not likely we can make much improvement with respect to that. However, we may be able to speed up the I/O operations. Theoretically, what is the limit on the speedup that could be achieved by improving the I/O hardware? (Not that we are claiming that limit can be achieved.)

We aren't given the actual execution time, but we can assume any convenient value, so let's say it was 100 seconds. So, the program would spend 75 seconds on operations that are not affected by our improvements, and 25 seconds on operations that are affected:

T_{after} = 75 + 25 / (I/O improvement)

The theoretical limit would be if we produced an infinite speedup for I/O operations (so, no, we won't actually achieve the limit). In that case, the time after the improvement is 75 seconds, and so the limiting speedup would be 100 / 75 or 4/3 (about 1.33).

b) [8 points] On certain hardware, integer multiplication takes 4 times as many cycles as integer addition, which takes 3 times as many cycles as simple integer operations like shifting and incrementing. A particular program spends 60 seconds on integer multiplication, 90 seconds on integer addition, 20 seconds on simple integer operations, and 30 seconds on other operations.

A code review reveals that each of the integer multiplication operations can actually be replaced with a sequence of three shift operations and a single addition operation. What speedup would the program achieve if the integer multiplications were replaced as described?

Each integer multiplication will be replaced by three simple operations and one addition. If we say that one simple operation takes 1 time unit, then an addition takes 3 time units and a multiplication takes 12 units. The substitution would replace each multiplication instruction with a sequence that takes 6 time units, taking 1/2 as long.

So, the substitution would effectively speed multiplication operations up by a factor of 2. By Amdahl's Law, the execution time of the revised program would be:

T_after = T_unaffected + T_affected / SpeedupFactor = 140 + 60 / 2 = 170 seconds

The speedup would be T_old / T_new = 200 / 170 = 1.12

The key was to determine what improvement factor would be achieved if we made the instruction substitutions that were described; it resulted in a factor of 2 improvement for multiplication instructions.

3. [14 points] Pseudo-instructions give MIPS a richer set of assembly language instructions than those supported directly by the hardware. Write an implementation for the following pseudo-instruction, using only the MIPS32 instructions listed at the bottom of this page (those <u>are sufficient</u>). You must include comments explaining the logic of your solution. If you need to use any "extra" registers besides those used in the instruction itself, you may use the the t-registers \$t0, \$t1, etc.

```
mov_lt $rd, ($rs), ($rt)
    # if ( Mem[GPR[rs]] < Mem[GPR[rt]])
    # GPR[rd] = GPR[rs]
    # else
    # GPR[rd] = GPR[rt]
    #
    #
    # $rs and $rt should not be modified</pre>
```

The instruction requires fetching two values from memory, seeing if the first is less than the second, and copying a value if it is:

```
lw
           $t0, ($rs)
                                # load left operand to $t0
      lw
           $t1, ($rt)
                                # load right operand to $t1
      slt $t2, $t0, $t1
                                # $t2 = 1 iff $t0 < $t1
      beq $t2, $zero, else
                                # skip to else if $t0 >= $t1
      add
           $rd, $rs, $zero
      j.
           fi
else:
      add $rd, $rt, $zero
fi:
```

```
add
      rd, rs, rt
                        # GPR[rd] <-- GPR[rs] + GPR[rt]</pre>
addi
      rd, rs, imm16
                        # GPR[rd] <-- GPR[rs] + imm16
                                                            (sign extended)
and
      rd, rs, rt
                        # GPR[rd] <-- GPR[rs] AND GPR[rt] (bitwise)</pre>
                       # if GPR[rs] == GPR[rt], jump to label
beq
      rs, rt, label
                       # if GPR[rs] != GPR[rt], jump to label
bne
      rs, rt, label
lw
      rt, imm16(rs)
                       # GPR[rt] = Mem[GPR[rt] + imm16]
      rt, imm16(rs)
                       # Mem[GPR[rt] + imm16] GPR[rt] =
SW
      rd, rs, rt
                        # GPR[rd] <-- GPR[rs] OR GPR[rt]</pre>
or
                                                            (bitwise)
      rd, rs, rt
                       # GPR[rd] <-- GPR[rs] - GPR[rt]</pre>
sub
sll
      rd, rt, sa
                       # GPR[rd] <-- GPR[rs] << sa</pre>
                                                            (logical shift)
slt
      rd, rs, rt
                       # GPR[rd] <-- (GPR[rs] < GPR[rt] ? 1 : 0)
(imm16 denotes a 16-bit immediate value)
```



Everything else in the datapath is implemented as shown on the diagram in the Supplement.

Suppose that, initially, \$t0=0x3000, \$t1=0x2000, and \$t2=0x1000. Consider the execution of the following code in this buggy datapath. Determine the final values of the \$t0, \$t1, and \$t2 registers.

SW	\$t1,	0x0(\$t0)
SW	\$t2,	0x1000(\$t0)
addi	\$t0,	\$t0, 0x1000
lw	\$t1,	0x0(\$t0)
add	\$t2,	\$t1, \$t2

Register	Final Value
\$t0	
\$t1	
\$t2	

#	Mem[0x3000] = 0x3000
#	Mem[0x4000] = 0x3000
#	t0 = 0x4000
#	t1 = Mem[0x4000] = 0x3000
#	t2 = 0x3000 + 0x1000 = 0x4000

5. [15 points] Suppose the following sequence of instructions was executed on the preliminary pipeline design shown on the Supplement, which has no hardware to detect or handle hazards:

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

Due to the limitations of the preliminary datapath design, a number of different logical errors are likely to occur. These may include possibly incorrect values being used as input to an instruction, known as *read-after-write* hazards.

Identify each of the read-after-write hazards that occur in the given sequence of instructions. For each, clearly identify the writing instruction, the reading instruction, and the register that is involved in the hazard.

Writer#	Reader#	Register(s)	
#3	#5	\$t2	< example, may be wrong
#1	#2	\$t4	
#1	#3	\$t4	
#2	#4	\$t4	
#2	#4	\$10	
#4	#5	\$t2	
#4	#6	\$t2	

- 6. Suppose we would like to design the new MIPS64 architecture using 64-bit addresses, 64-bit instructions, and 64 (generalpurpose integer) registers. We will keep the same number of instructions as in the original MIPS32 architecture, so that we can use 6-bit opcodes, and we still require that machine instructions be aligned on 8-byte multiples.
 - a) [5 points] For I-format instructions, how many bits could we use for the immediate field? Explain.

Register numbers take 6 bits (2⁶ registers now), opcodes still take 6 bits, so with 64-bit instructions, we have 46 bits for the immediate field.

The existing MIPS32 datapath contains the following hardware elements to support the beq instruction:



The new MIPS64 datapath will involve essentially similar hardware, from a logical point of view, but some details will be different. For example, the change to the item labelled \mathbf{A} in the diagram below was explained in your answer to part \mathbf{a}).



b) [5 points] Explain <u>how</u> item **B**, in the new design, differs from the corresponding element in the original MIPS32 design, and <u>why</u>. (The MIPS32 diagram may provide hints.)

This specifies the value to be added to the PC if we don't branch. Since instructions are now 8 bytes wide, this must now be 8.

c) [5 points] Explain how item **C**, in the new design, differs from the corresponding element in the original MIPS32 design, and why.

The PC holds the instruction that was just fetched, so it must now be a 64-bit register.

d) [5 points] What is the logical connection between **C** and **D**? (The answer would be the same for both the old and the new designs.)

Since instructions are 8 bytes long, they are stored at addresses that are multiples of 8.

D shifts the immediate to compensate for the bits that would have been dropped from the offset between PC + 8 and the branch target address. Since instruction addresses always have 000 for their low bits, so we will drop 3 bits from the offset, and the shifter must now shift left by 3 bits.

e) [5 points] Explain what item **E** should be, in the new design, and <u>why</u>. (The MIPS32 diagram may provide hints.)

The immediate field for jump is now 58 bits wide (64 bits minus the opcode width). The shifter produces a 61-bit value.

So, E completes the 64-bit address by supplying bits 63:61 from PC + 8.