

EECS 151/251A Homework 4

Due Friday, Oct 2nd, 2020

Midterm Practice [1 pt]

Before you start the rest of this homework assignment, please practice the mechanics of the midterm exam. Everything must be uploaded by Friday, October 2 (same time as this HW).

1. Make sure you read through the [Exam Policy](#). Set up your workspace and Zoom to conform to the rules in the policy.
2. Submit your Zoom meeting ID that you will use for the exam to [this form](#).
3. For this mock exam, you will just solve the practice problems on the next page only. Set up your recording and workspace according to the policy, solve the problem as you would on the actual exam (give yourself 5 minutes), then end the recording.
4. Upload your completed problems to Gradescope, under the assignment "Practice Exam". This will not be graded for correctness, just for completion.
5. Submit your recording link to [this form](#). GSIs will contact you before the midterm if there are any issues with your recording.
6. Done! Now, proceed with the HW.

Practice Problems

1. Fill in the 5 stages of the RISC-V datapath (abbreviations are fine):

Stage 1: _____

Stage 2: _____

Stage 3: _____

Stage 4: _____

Stage 5: _____

2. A key part of the RISC-V base (32-bit) datapath is the register file, which provides the source operands and destination for most computations.

(a) Draw the conceptual diagram of the register file, labeling the input and output signals with their bitwidths.

(b) What is the value in register x0? _____

(c) During which operation is the input clock actually used: read or write? _____

Reading

In addition to reviewing the RISC-V ISA and datapath lectures, skim through the [RISC-V ISA spec](#). In particular, focus on the Introduction, Chapter 2 (RV32I Base Integer Instruction Set), and the tables on pages 129 and 130.

Problem 1: RISC-V Manual Assembly [4 pts (1 each)]

Manually construct the binary instruction for the following assembly instructions. Submit all of the following for each instruction:

- The 32-bit binary number for the instruction
- The core instruction format it belongs to
- Delineate the 32 bits into the subfields of the instruction format and label each field with the opcode/registers/immediate/offset etc. specified by the instruction

Note: we highly encourage you to do this by hand from the ISA spec, but it is possible to assemble them using RISC-V GCC or [venus](#).

- (a) `sra x1, x2, x3`
- (b) `andi x1, x2, 100`
- (c) `sh x1, 4(x2)`
- (d) `bne x6, x8, 1024`

Problem 2: RISC-V Assembly Programs [4 pts (1 each)]

Write down the values of the specified registers after the following programs have run. Show your work by annotating the what happens/changes after each instruction. Note that some instructions are pseudo-instructions, such as `li` for load immediate. Refer to Table 25.2 in the RISC-V spec for a list of pseudo-instructions and their base implementations.

(a) `li x0, 100`
`li x1, 200`
`sub x2, x1, x0`
`addi x2, x2, 100`

`x0` = _____, `x1` = _____, `x2` = _____

(b) `li x1, 0xdead`
`li x2, 0xbeef`
`li x3, 0x128`
`sb x2, 0(x3)`
`sra x2, x2, x3`
`sb x2, 1(x3)`
`sb x1, 2(x3)`
`sra x1, x1, x3`
`sb x1, 3(x3)`
`lw x4, 0(x3)`

`x1` = _____, `x2` = _____, `x4` = _____

(c) `li x1 -1`
`li x2 1`
`f1: sll x3 x1 x2`
`sub x1 x1 x3`
`addi x2 x2 1`
`blt x1 x2 f1`

`x1` = _____, `x2` = _____, `x3` = _____

(d) Assume the following instructions start at address `0x0`.

```

    li x1, 0
    jalr x2, x0, 12
    addi x1, x1, 100
    jalr x2, x2, 0
    addi x1, x1, 100
    jalr x2, x2, 0
    j more
more: jalr x2, x2, 0
      j end
end:  nop

```

`x1` = _____, `x2` = _____

Problem 3: RV64M Multiplication ALU [5 pts (1 per inst.)]

Refer to Chapter 7 ("M" Standard Extension for Integer Multiplication and Division, Version 2.0) in the RISC-V spec and the corresponding instruction set listing on page 131. Using the template given here, implement an ALU that supports the subset of the **64-bit "M" extension** that performs multiplication. Don't worry about any synthesizability/performance issues of your implementation in this exercise.

Pay attention to which bits are important in each variant of the multiply, as well as the signed-ness of the inputs to the multiply (*) operator. You may want to make sure you are comfortable with Verilog [signed arithmetic](#) and the [concatenation operator](#) first.

```
`define ALU_MUL 0
`define ALU_MULH 1
`define ALU_MULHSU 2
`define ALU_MULHU 3
`define ALU_MULW 4
module rv64mult_alu(
    input [63:0] a,
    input [63:0] b,
    input [2:0] op, // op is one of the values `define'd above
    output [63:0] c,
);

    // Your implementation

endmodule
```

Problem 4: Extending RV32I Branching for ReLU [10 pts]

The Rectified Linear Units, or ReLU function is very useful in machine learning. It is the most popular [activation function](#) and is defined as $y = \max(0, x)$. Graphically, it looks like this:




Figure 1: ReLU activation function

In RV32I, this function can be implemented with a couple instructions, using a branch. Let's try to extend RV32I to do a ReLU in a single instruction by using what exists in the datapath.

- (a) (1 pt) Fill in the assembly below needed to do a ReLU in RV32I, using a `bge` instruction. Assume register `x1` already stores x and we want register `x2` to get y .

```

mv
bge
mv
end: nop

```

- (b) (4 pts) To guide you toward a solution, first complete the Verilog implementation of `ALUSel` in the Control Logic for RV32I. For this part, write down the logical expressions for the output control signals in terms of an instruction's `opcode`, `funct3`, and `funct7` bits, as well as the `BrEq` and `BrLT` signals from the branch comparator. You can also simplify the expressions by comparing the `opcode` field against the constants in Table 24.1 in the RISC-V ISA manual. For example: `assign sig = (opcode == OP-32) || (opcode == LOAD);`.

```

// The rest of the Control Logic module is not implemented here.
// ALUSel here is a 4-bit output, and should take the constants
// defined below for each type of operation the ALU is to do.
parameter ALU_ADD = 4'b0000;
parameter ALU_SUB = 4'b0001;
parameter ALU_SLT = 4'b0010;
parameter ALU_SLTU = 4'b0011;
parameter ALU_AND = 4'b0100;
parameter ALU_OR = 4'b0101;
parameter ALU_XOR = 4'b0110;
parameter ALU_SLL = 4'b0111;
parameter ALU_SRL = 4'b1000;

```

```
parameter ALU_SRA = 4'b1001;

always @(*) begin
    ALUSel = ALU_ADD; // default

    // Complete your implementation here

end
```

(c) (2 pts) We now want to implement ReLU as an R-type instruction. For this problem, the instruction format would be like `relu rd rs1 rs2` such that:

- We must specify register `x0` for `rs2` so as to check against the value 0 for ReLU. Any other value for `rs2` would give us an invalid result.
- `rs1` should be the value of x and `rd` would be the value of y .

In essence, we want the branch comparator to help us with the ReLU function. Submit the following:

- (i) Refer to table 24.1 and the following description in the ISA spec. How can we extend our ISA to enable adding our ReLU instruction?
- (ii) With the scheme, we do not need to update the datapath. Instead, what do we need to change?

(d) (2 pts) Now, update your Verilog code from b) to enable your new ReLU instruction.

(e) (1 pt) Could we also implement our new ReLU instruction as an I-type instruction instead? If so, how would its implementation be different from the R-type version?

Problem 5: Extending RV32I for BNNs [10 pts]

We want to add another, more complex instruction to this ISA that may help us with certain neural network applications: a bitwise multiply-and-accumulate (MAC) operation. This operation can also be considered as a bitwise dot product in linear algebra speak. Here are the specifications:

- Each 32-bit datum in our architecture represents a 32-element long vector of binary values
- Bits translate to binary values as follows: a "0" bit represents the value "-1" and a "1" bit represents the value "+1"
- The MAC is therefore calculated as follows, shown for a 4-bit vector example for understanding.

Given input binary vectors $x = \{1, -1, 1, -1\} = 1010$ and $y = \{-1, 1, 1, -1\} = 0110$:
 $\text{mac}(x, y) = (+1 * -1) + (-1 * +1) + (+1 * +1) + (-1 * -1) = 0$.

Background: [Binary Neural Networks \(BNNs\)](#) are neural networks with weight and activation matrices quantized to ± 1 . This extreme degree of quantization is desired for applications where data movement must be kept to an absolute minimum because it compresses one dimension of the weight and activation matrices. Furthermore, BNNs promise to dramatically save computational resources, because a MAC can use bitwise computations rather than integer multiplication and addition.

Your task: Implement the algorithm that accomplishes this bitwise MAC, called [XNOR-Popcount](#). In this algorithm, the bitwise XNOR is taken, followed by a popcount, and finally a scaling adjustment. In pseudocode:

```
a = xnor(x,y)
b = popcount(a)
c = len(a) // this is known (32)
mac(x,y) = 2 * b - c
```

The [popcount](#) operation here (otherwise known as Hamming weight) is also very useful for other applications such as communications and cryptography, but is missing from the RISC-V ISA spec. Due to popular demand, it is currently draft proposed in the "B" or [bitmanip extension](#).

The popcount of a binary number is the sum of all the 1's in the bitstring. For example, the popcount of 10110100 is 4. The naïve method of implementing popcount using an integer instruction set would be check each bit one-by-one, incrementing a counter if we see a 1. In a C implementation for 32 bits, it would look like this:

```
int pop(unsigned x) {
    int count = 0;
    for (i = 0; i < 32; i++) { // loop through all bits
        if ((x & 1) == 1)     // mask LSB and check for 1
            count++;
    }
    return count;
}
```



```

        count++;
    x >>= 1;           // right shift by 1
}
return count;
}

```

Clearly, there are an extremely large number of RV32I instructions needed to implement this. There is a much faster algorithm using an integer instruction set, as described in [Hacker's Delight](#):

```

int pop(unsigned x)
{
    x = x - ((x >> 1) & 0x55555555);
    x = (x & 0x33333333) + ((x >> 2) & 0x33333333);
    x = (x + (x >> 4)) & 0x0F0F0F0F;
    x = x + (x >> 8);
    x = x + (x >> 16);
    return x & 0x0000003F;
}

```

To accelerate the MAC operation for BNNs, we would like to implement XNOR-Popcount as a single instruction. To do so, you'll need to add an accelerator (new ALU). For these questions, assume that in your new accelerator, you can implement any number of new combinational logic operations (no matter how complex), and that the input binary vector x is already stored in the register file.

- (a) (3 pts) Fill in the assembly instructions below that are needed to calculate the popcount using the faster popcount algorithm. The destination registers have been given to you, the original number is already in register $x1$, and the result should also be in register $x1$ at the end:

```

srli x2
li x3
and x2
sub x1
li x3
and x2
srli x4
and x4
add x1
srli x2
add x1
li x2
and x1
srli x2
add x1
srli x2
add x1
andi x1

```

How many RV32I instructions is this?

- (b) (2 pts) For this faster popcount, if you could extend RV32I with one fused (i.e. 2+ instructions combined) I-type instruction to minimize the number of instructions, what logic function should it implement? How many instructions would it save compared to part a)?
- (c) (3 pts) It turns out at the hardware level, we can do a popcount without all this fancy shifting and masking, and can therefore do an XNOR-Popcount purely combinatorially. Explain how and write the Verilog implementation of this.
- (d) (1 pt) Is the algorithm you used in part c) practical, especially as the length of the input vectors grows (e.g. we put this in a 64-bit ISA instead)? Explain.
- (e) (1 pt) Let us now expand this operation to perform a binary MAC on vectors larger than 32 bits on our 32-bit architecture. How would you do this over multiple instructions without changing the RV32I datapath, control logic, or your new ALU from part d)?