

CS61C MIDTERM 2

<i>Last Name (Please print clearly)</i>		
<i>First Name (Please print clearly)</i>		
<i>Student ID Number</i>		
<i>Circle the name of your Lab TA</i>	Damon	Jonathan
	Sean	Sruthi
	Suvansh	Sukrit
<i>Name of the person to your: Left Right</i>		
<i>All my work is my own. I had no prior knowledge of the exam contents nor will I share the contents with others in CS61C who haven't taken it yet. (please sign)</i>		

Instructions


- This booklet contains **10** pages including this cover page. **The back of the exam contains scratch paper and a handout that can be used for scratch work, but will not be graded** (i.e. not even scanned into Gradescope).
- Please turn off all cell phones, smartwatches, and other mobile devices. Remove all hats, headphones, and watches. Place *everything* except your writing utensil(s), cheat sheet, and beverage underneath your seat.
- You have 80 minutes to complete this exam. The exam is closed book: no computers, tablets, cell phones, wearable devices, or calculators. You are allowed one page (US Letter, double-sided) of *handwritten* notes.
- There may be partial credit for incomplete answers; write as much of the solution as you can.
- Please write your answers within the boxes and blanks provided within each problem!
- For multiple choice questions, please **bubble** in the circle. Failure to do so may impact our ability to grade your exam.

Question	1	2	3	4	5	Total
Possible Points	11	12	20	22	15	80

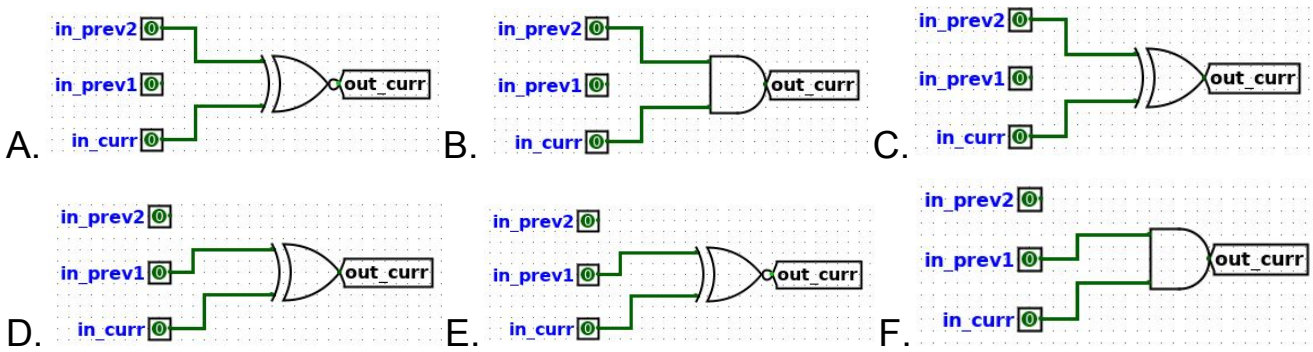
If you have the time, **feel free to doodle on this front page!**

Question 1: FSMSF (11 pts)

- 1) Construct an FSM which outputs a 1 when the most recent 3 bits in the input stream are a palindrome. A palindrome is a sequence which is the same when displayed in reverse order. An example of a palindrome is 010, whereas 001 is not. This FSM should continually be updating the output when the input stream changes. For example, when the bitstream is 0010010101, the output is **01001111. The first two outputs are undefined and you should ignore them. Assume the bitstream is always initialized to 00 for you. Your FSM should use the fewest possible number of states. As a reminder you should represent transitions with an arrow, where the tail of the arrow is the origin state and the tip is the destination state. You should also label each transition arrow as <input> / <output>. (e.g: 0 / 1)



Select the circuit diagram which produces the correct output for the FSM. in_curr is the current input to the FSM ($in[i]$), in_prev1 is the previous input ($in[i - 1]$), and in_prev2 is the input before the previous input ($in[i - 2]$).



- 2) Ⓐ Ⓑ Ⓒ Ⓓ Ⓔ Ⓕ

Question 2: Simple Democratic Selection (12 pts)

As the semester is reaching a close, Steven, Nick, and Damon are busy determining the difficulty of the final exam. All three will vote on whether the final should be easy or hard, but the final decision will always be made based on the following rules:

- Rule 1.** If the vote is unanimously hard or unanimously easy, then it will be hard or easy, respectively.
- Rule 2.** If Damon disagrees with Steven and Nick, then Damon's vote will be chosen.
- Rule 3.** If Steven and Nick differ, then the minority vote will be chosen.
- Else** In all other situations, the outcome can be either easy or hard (i.e. they can be anything)


We will represent Steven's vote with the variable **S** which takes on values of 0 (easy) and 1 (hard). Similarly, Nick's vote is represented as **N** and Damon's vote is represented as **D**. For each rule, write out the **simplest** boolean logic expression using these three binary inputs that outputs whether or not the final exam will be easy or hard. Note: the symbol for XOR is \oplus .

Rule 1: _____ **Rule2:** _____ **Rule3:** _____

Below is a boolean algebra expression that models this problem. Simplify it into as few gates as possible:

$\bar{S}(N+D) + S(S + (N \oplus D)) + \bar{D}$ Simplified: _____

For the next 2 questions, refer to the circuit diagram below which models the unsimplified boolean expression:



You are given that all logic gates have a propagation delay of 5ps, the register clk-to-q delay is 3ps, the register hold time is 11ps, and the setup time is 8ps. What is the critical path delay of this circuit?

Critical path delay: _____ ps

Using the assumptions above, what is the smallest value that the clk-to-q can be and **not** cause a hold-time violation?

clk-to-q: _____ ps

Question 3: Maddi-o and Lui-p (20 pts)

As discussed in lecture, the standard RISC-V approach for loading in a 32-bit immediate into a register is by using the following `lui-addi` pair of instructions:

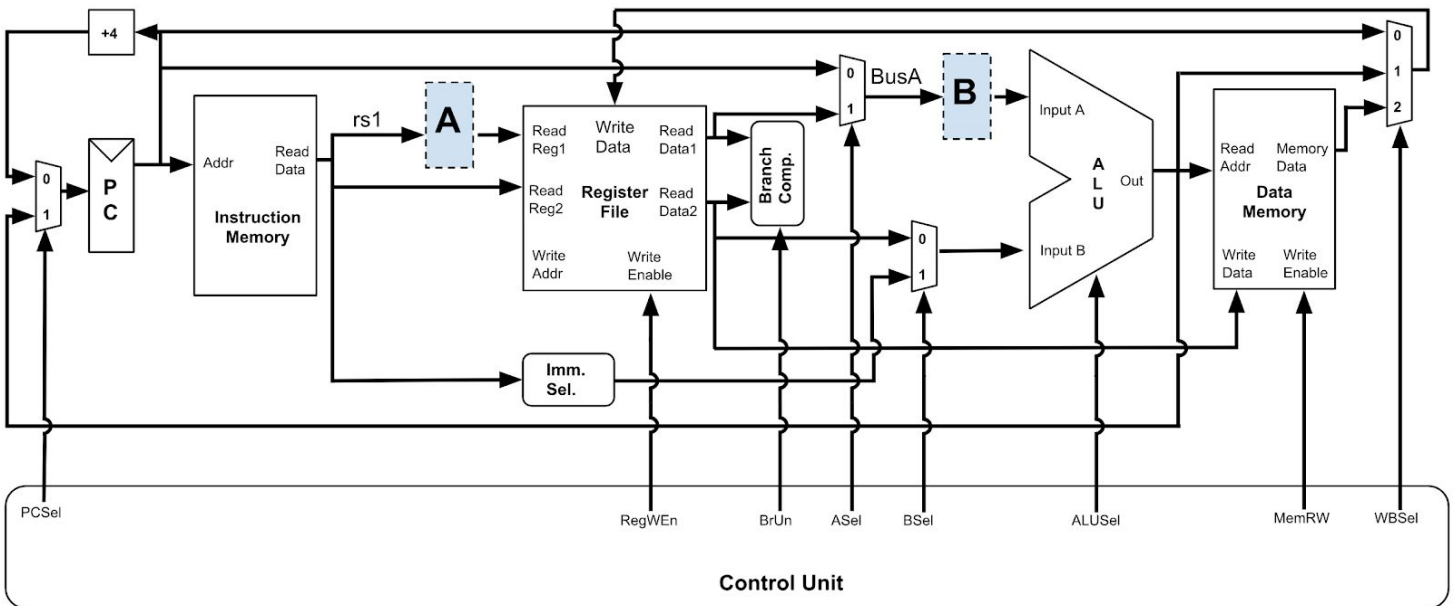
```
lui rd HIGH_IMM
addi rd rd LOW_IMM
```

We will explore an alternative method of storing a 32-bit immediate into a register by implementing a new **U-type** instruction named `lui-p`, which will load a 20-bit immediate into the upper 20 bits of the `rd` register and **preserve** the lower 12 bits of the `rd` register. This way, we can use an `addi-lui-p` pair to load a 32-bit immediate into a register, as shown on the left. On the right is a description of what `lui-p` does:

```
addi rd x0 LOW_IMM
lui-p rd HIGH_IMM
```

$$R[rd] = R[rd][11:0] + \{imm, 12b'0\}$$

Below is the existing RISC-V datapath as presented in lecture, except there are two boxes labeled “A” and “B” which are areas in the datapath that you will fill in with the necessary hardware to implement `lui-p`.



For Parts A and B, please **clearly circle** the correct hardware that would fill their respective dashed boxes above. For Part C, **circle** all the ALUSel operations that could be used by the `lui-p` instruction in the ALU unit.

A.

B.

C.

AND	OR
ADD	SUB
XOR	BSEL

D. Fill out the following control signals that must be generated by the control unit in order to correctly implement **luip**. Fill in the bubble for “X” if the exact value of the control signal does not matter (e.g. a value of 0 or 1 is incorrect when the signal could instead be X).

- 0 Signal = 0 1 Signal = 1 2 Signal = 2
- R Write Disabled W Write Enabled X Don't Care

luip	PCSel	RegWEn	BrUn	ASel	BSel	MemRW	WBSel
1	0 1 X	R W X	0 1 X	0 1 X	0 1 X	R W X	0 1 2 X

E. Now that we've added **luip** to our instruction set architecture, we need to give it an opcode. We now have 3 U-type instructions: **auipc**, **lui**, and **luip**. The opcode for **auipc** is 0b001 0111 and the opcode for **lui** is 0b011 0111. For this question, we refer to the most significant (leftmost) opcode bit as bit 6 and the least significant (rightmost) opcode bit as bit 0.

You are given the following truth table whose two inputs are two of the opcode bits, and the output is which U-type instruction is being identified by the two opcode bits. Fill out all possible opcode bits that could be the input to the truth table in order to generate the specified outputs (**an output of X means there is no defined instruction with that opcode in our ISA**).

(Hint: the other five bits of the opcode are compared to a **constant** value which identifies the opcode as being a U-type opcode—your task is figuring out which bits can identify which U-type instruction it is.)

Opcode bit(s): 6 5 4 3 2 1 0	Opcode bit(s): 6 5 4 3 2 1 0	Instruction
0	0	X
0	1	luip
1	0	auipc
1	1	lui

Question 4: RISC IS HAZARDOUS (22 pts)

Consider a RISC-V pipeline in 3 stages with the following specification:

Stage 1	Stage 2	Stage 3
Instruction Fetch/Instruction Decode (IFD)	Execute (EX)	Memory/Write Back (MWB)

It also has the following details:

- No Forwarding
- No reading and writing to the same register in the same cycle
- Reading and writing to the same register is considered a data hazard

Consider the following piece of code

```

1.      add t1 x0 x0
2.      add t2 x0 x0
3.      addi a0 x0 2
4.      slli a0 a0 2
5. L2:  bge t1 a0 End
6.      add t3 sp t1
7.      lw t3 0(t3)
8.      add t2 t2 t3
9.      addi t1 t1 4
10.     j L2
.
.
.
End:
    
```

- 1) Assume that instead of branch prediction, the CPU always stalls to resolve a control hazard. How many stall(s) are necessary for each control hazard **each time** it is encountered? You may not need all boxes.

Line Number	# Stalls/Encounter

2)

- a) Where are the data hazards that produce stall(s)? You should describe each hazard as a tuple, (A, B), where instruction A is the the instruction that triggered a data hazard in instruction B. Place A's line number in the left box and B's line number in the right box You may not need all boxes.

Line of Instruction that first accesses the data	Line of instruction that must be stalled

- b) How many **total cycles** will it take to complete the code above? Assume the pipeline is cleared upon reaching End. In addition, assume that there is perfectly accurate branch and jump prediction in the IFD stage. Thus, **ONLY CONSIDER STALLS DUE TO DATA HAZARDS. ASSUME THERE ARE NO STALLS FOR STRUCTURAL OR CONTROL HAZARDS.** Remember to **calculate the total number of cycles for fully executing the code.** A workspace was provided with this exam for you to show your work, which will only be graded if your final answer is not correct.

Total Cycles: _____

Now assume that our pipeline has full forwarding along with perfect branch/jump prediction done in the IFD stage.

- 3) What are the data hazards that remain? How many cycles must be stalled **each time** this hazard is encountered? If there are two instructions with the same data hazard and both are resolved by the first stall then only list the first instruction as needing to be stalled. The left box(es) should be a subset of the value(s) you had in the rightmost box(es) in 2a. You may not need all boxes.

Line of instruction that must be stalled	# Stalls/Encounter

Question 5: C.R.E.A.M. (15 pts)

- 1) A machine has an 8 way set associative cache with 512 B of data. The size of each block is 16 B and there are 8 MiB of main memory. How large are the tag, index, and offset fields for an address on this machine using this cache?

Tag: _____ Index: _____ Offset: _____

- 2) Now we have a different machine with two caches, an L1 and an L2 cache. Both caches are direct mapped caches. The L1 cache can hold 256 B of data and the L2 cache can hold 4 KiB of data. Assume the following code is run on this machine:

```
#define ARR_SIZE 2048
```

```
uint16_t sum (uint16_t *arr) {
    total = 0;
    for (int i = 0; i < ARR_SIZE; i++) {
        total += arr[i];
    }
    return total;
}
```

This produces a hit rate (HR) of $\frac{7}{8}$ for the L1 cache and $\frac{3}{4}$ for the L2 cache. Given that arr is a block aligned address and `sizeof (uint16_t) == 2`:

- A. What is the blocksize of the L1 cache **in bytes** that produces its hit rate?

L1_BLOCKSIZE: _____

- B. Use the variable Y to represent the answer to part A. What is the blocksize of the L2 cache **in bytes** that produces its hit rate? Express your answer as a function of Y and NOT as a single number.

L2_BLOCKSIZE: _____

Recall that the L1 HR is $\frac{7}{8}$ and the L2 HR is $\frac{3}{4}$. If the L1 Cache has a hit time of 2 cycles, the L2 cache has a hit time of 8 cycles, and main memory has a hit time of 96 cycles:

- 3) How many total cycles are spent accessing memory on this piece of code? Express your answer in the form $C * 2^i$, where C is an integer not divisible by 2.

TOTAL_CYCLES: _____

- 4) If we change the L1 cache from being direct mapped to being fully associative with LRU, how does its HR change on the same code?

- (A) Increases (B) Decreases (C) No Change (D) Impossible to tell