

Solutions last updated: Wednesday, December 20th, 2023

PRINT your name: \_\_\_\_\_

PRINT your student ID: \_\_\_\_\_

You have 170 minutes. There are 11 questions of varying credit (100 points total).

|           |   |    |   |    |    |   |   |   |   |    |    |       |
|-----------|---|----|---|----|----|---|---|---|---|----|----|-------|
| Question: | 1 | 2  | 3 | 4  | 5  | 6 | 7 | 8 | 9 | 10 | 11 | Total |
| Points:   | 7 | 13 | 6 | 11 | 13 | 7 | 9 | 9 | 8 | 16 | 1  | 100   |

For questions with **circular bubbles**, you may select only one choice.

- Unselected option (completely unfilled)
- Only one selected option (completely filled)
- Don't do this (it will be graded as incorrect)

For questions with **square checkboxes**, you may select one or more choices.

- You can select
- multiple squares
- (completely filled)

Anything you write outside the answer boxes or you ~~cross-out~~ will not be graded. If you write multiple answers, your answer is ambiguous, or the bubble/checkbox is not entirely filled in, we will grade the worst interpretation. For coding questions with blanks, you may write at most one statement per blank and you may not use more blanks than provided.

If an answer requires hex input, you must only use capitalized letters (**0xDEADBEEF** instead of **0xdeadbeef**). For hex and binary, please include prefixes in your answers unless otherwise specified, and do not truncate any leading 0's. For all other bases, do not add any prefixes or suffixes.

**Write the statement below in the same handwriting you will use on the rest of the exam.**

I have neither given nor received help on this exam (or quiz), and have rejected any attempt to cheat; if these answers are not my own work, I may be deducted up to 0x0123 4567 89AB CDEF points.

---

---

---

---

---

---

SIGN your name: \_\_\_\_\_

Clarifications made during the exam:

Q4: You may assume that `i` is stored in a register.

Q4: You may assume that the cache starts cold.

Q5.10: Lines 6 and 7 should say

```
uint32_t evens = 0;
uint32_t odds = 0;
```

Q7: "48-bit address space" refers to a virtual address space.

Q9: The written description and given code is erroneously inverted. Assume that the code is correct.

Q8: The function `f` returns `f(x)` in `a0`. On line 5, `t7` should be `t6`.

Q9: The argument to `solve_orion` should say `uint32_t(*orion)(uint32_t)`.

Q9.1: You may use functions from standard libraries, but may not use any other external functions.

**Q1 Potpourri****(7 points)**

Q1.1 (1.5 points) Using multithreading is guaranteed to speed up all programs.

- True                       False

**Solution:** Using multithreading to add together two numbers is not faster than using a single thread due to the overhead.

Q1.2 (1.5 points) Virtual memory provides memory isolation across all threads.

- True                       False

**Solution:** Threads within the same process share the same address space, so memory is not isolated.

Q1.3 (2 points) You currently have a memory system with an L1 cache and DRAM with the following hit times and hit rates:

|          | Hit Time | Local Hit Rate |
|----------|----------|----------------|
| L1 Cache | 5ns      | 20%            |
| DRAM     | 250ns    | 100%           |

You want to add an L2 Cache to improve the average memory access time of this system. Assume that the hit time for this L2 Cache would be 15ns. What is the local hit rate that the L2 Cache would need to make the average memory access time of this system 25ns? Express your answer as a percentage.

**Solution:** 96%.

$$\begin{aligned}
 t_{L1} + MR_{L1} \times (t_{L2} + MR_{L2} \times t_{DRAM}) &= 25ns \\
 5ns + 0.8 \times (15ns + MR_{L2} \times 250ns) &= 25ns \\
 0.8 \times (15ns + MR_{L2} \times 250ns) &= 20ns \\
 15ns + MR_{L2} \times 250ns &= 25ns \\
 MR_{L2} \times 250ns &= 10ns \\
 MR_{L2} &= \frac{1}{25} = 4\%
 \end{aligned}$$

Since the miss rate for the L2 cache is 4%, the hit rate must be 96%.

**Grading:** This problem was graded on an all-or-nothing basis.

(Question 1 continued...)

Q1.4 (2 points) You have a program that takes 20 seconds to run, but you've found a way to make 80% of the code 4 times faster at the cost of some overhead. After rerunning your program, it takes 13 seconds. What is the overhead (in seconds)?

**Solution:** 5 seconds

80% of the 20 second program would take 16 seconds, and would take 4 seconds if we make it 4 times faster. The remaining 20% of the program, or 4 seconds, remains the same. If there was no overhead, it would take 8 seconds (4 seconds from non-parallelizable section and 4 seconds from the parallelized section). Since the overall runtime was 13 seconds, this must mean the overhead is 5 seconds.

**Grading:** This problem was graded on an all-or-nothing basis.

**Q2 Eddy Needs a Project 3 Extension****(13 points)**

For this problem, assume that we're working with the single-cycle datapath presented on the reference card.

Suppose Eddy wants to support the following instruction:

`baleq rd rs1 rs2 offset` (branch and link if equal)

```
if (rs1 == rs2) {  
    rd = PC + 4  
    PC = PC + offset  
}
```

For each of the following control signals, indicate the value it should have for `baleq`. If the control signal is not constant, select "None of the above".

Q2.1 (1.5 points) PCSel

- |                              |  |
|------------------------------|--|
| <input type="radio"/> PC + 4 | <input type="radio"/> Doesn't matter               |
| <input type="radio"/> ALUOut | <input checked="" type="radio"/> None of the above |

**Solution:** The value of PCSel is determined by the output of the branch comparator, so it is not a constant, therefore "None of the above".

Q2.2 (1.5 points) RegWEn

- |                                      |  |
|--------------------------------------|--|
| <input type="radio"/> Write disabled | <input type="radio"/> Doesn't matter               |
| <input type="radio"/> Write enabled  | <input checked="" type="radio"/> None of the above |

**Solution:** The value of RegWEn is determined by the output of the branch comparator (since we only update `rd` if the branch is taken), so it is not a constant, therefore "None of the above".

Q2.3 (1.5 points) ASel

- |                                     |   |
|-------------------------------------|---|
| <input type="radio"/> rs1           | <input type="radio"/> Doesn't matter    |
| <input checked="" type="radio"/> PC | <input type="radio"/> None of the above |

**Solution:** Since the ALU needs to compute `PC + offset`, ASel must be PC. This is identical to branch instructions.

(Question 2 continued...)

Q2.4 (1.5 points) BSEL

- rs2
- imm
- Doesn't matter
- None of the above

**Solution:** Since the ALU needs to compute  $PC + \text{offset}$ , BSEL must be `imm`. This is identical to branch instructions.

Q2.5 (1.5 points) BRUN

- Signed
- Unsigned
- Doesn't matter
- None of the above

**Solution:** Since we only check if `rs1 == rs2`, it doesn't matter if the branch comparator makes a signed or unsigned comparison.

Q2.6 (1.5 points) WBSSEL

- PC + 4
- ALUOut
- MemReadData
- Doesn't matter
- None of the above

**Solution:** We need to write `PC + 4` back to `rd`, so we should select `PC + 4` for WBSSEL.

(Question 2 continued...)

Our RISC-V datapath currently does not support an instruction that atomically loads from and stores to memory. Suppose we introduce the following instruction:

`alsw rd rs2 imm(rs1)` (atomic load and store word)

```
rd = *(rs1 + imm)
*(rs1 + imm) = rs2
```

Q2.7 (4 points) What additional changes would we need to make to our datapath in order for us to implement `alsw` (with as few changes as possible)? Select all that apply.

- Create a new instruction type and update the ImmGen
- Add a new output to the RegFile for a third register value
- Add another WriteData and WriteIndex input to the RegFile
- Add a new input to the AMux and update any relevant selector/control logic
- Add a new input to the BMux and update any relevant selector/control logic
- Add a new ALU operation and update any relevant selector/control logic
- Add a third input into ALU and update any relevant selector/control logic
- Allow the ALU to send out more than 1 output and update any relevant selector/control logic
- Allow the DMEM to be able to read and write at the same clock cycle and update any relevant selector/control logic
- Add a new input to the WBMux and update any relevant selector/control logic
- None of the above

**Solution:** Since this instruction has `rd`, `rs1`, `rs2`, and `imm`, we need to create a new instruction type (as no current instruction type support all of these fields). We also need to create a new immediate type for this new instruction type since we can't share an immediate type with another instruction type.

This instruction also performs one read and one write memory operation in the same cycle, which is currently not supported by our DMEM (as it only has one port). Therefore, we need to allow the DMEM to read and write in the same cycle.

The remaining control signals remain the same as a load instruction, so we don't need to add any additional inputs/outputs to other components.

**Q3  $\pi$ -pelining**

**(6 points)**

For this problem, assume that we're working with the five-stage pipelined datapath presented on the reference card, and that the CPU will always predict that branches are not taken.

Consider the following RISC-V code:

```

1  beq x0 x0 Label
2  addi x0 x0 3
3  addi x0 x0 1
4  addi x0 x0 4
5  addi x0 x0 1
6  addi x0 x0 5
  Label:
7  addi t0 x0 9
8  addi t1 t0 2
9  xori t1 x0 6
    
```

Suppose that the IF stage of the `beq` on line 1 occurs during cycle 1.

Q3.1 (3 points) If all hazards are resolved through stalling (no double pumping or forwarding paths), during which cycle does the `xori` on line 9 execute its WB stage?

For each hazard, write down the hazard, the instructions involved, and the number of stalls required. We will only look at this box if you request a regrade.

**Solution:** There is a control hazard from line 1 to line 7, as the branch will always be taken. Then, from line 7 to line 8, we have a data hazard from writing to `t0` to accessing it again in the next line. This leads to the below timing diagram:

| Instruction            | 1  | 2  | 3  | 4   | 5  | 6  | 7  | 8   | 9  | 10 | 11 | 12  | 13  | 14 |
|------------------------|----|----|----|-----|----|----|----|-----|----|----|----|-----|-----|----|
| 1. beq x0 x0 Label     | IF | ID | EX | MEM | WB |    |    |     |    |    |    |     |     |    |
| 2. addi x0 x0 3 -> nop |    | IF | X  | X   | X  | X  |    |     |    |    |    |     |     |    |
| 3. addi x0 x0 1 -> nop |    |    | IF | X   | X  | X  | X  |     |    |    |    |     |     |    |
| 4. addi x0 x0 4 -> nop |    |    |    | IF  | X  | X  | X  | X   |    |    |    |     |     |    |
| 7. addi t0 x0 9        |    |    |    |     | IF | ID | EX | MEM | WB |    |    |     |     |    |
| 8. addi t1 t0 2 -> nop |    |    |    |     |    | IF | X  | X   | X  | X  |    |     |     |    |
| 8. addi t1 t0 2 -> nop |    |    |    |     |    |    | IF | X   | X  | X  | X  |     |     |    |
| 8. addi t1 t0 2 -> nop |    |    |    |     |    |    |    | IF  | X  | X  | X  | WB  |     |    |
| 8. addi t1 t0 2        |    |    |    |     |    |    |    |     | IF | ID | EX | MEM | WB  |    |
| 9. xori t1 x0 6        |    |    |    |     |    |    |    |     |    | IF | ID | EX  | MEM | WB |



(Question 3 continued...)

Q3.2 (3 points) If we implement double pumping and all forwarding paths, during which cycle does the `xori` on line 9 execute its WB stage?

For each hazard, write down the hazard, the instructions involved, and the number of stalls required. We will only look at this box if you request a regrade.

**Solution:** There is still the control hazard from line 1 to line 7, as the branch will always be taken. The data hazard from line 7 to line 8, for register `t0` is resolved via forwarding. This leads to the below timing diagram:

| Instruction                            | 1  | 2  | 3  | 4   | 5  | 6  | 7  | 8   | 9   | 10  | 11 |
|--|----|----|----|-----|----|----|----|-----|-----|-----|----|
| 1. <code>beq x0 x0 Label</code>        | IF | ID | EX | MEM | WB |    |    |     |     |     |    |
| 2. <code>addi x0 x0 3 -&gt; nop</code> |    | IF | X  | X   | X  | X  |    |     |     |     |    |
| 3. <code>addi x0 x0 1 -&gt; nop</code> |    |    | IF | X   | X  | X  | X  |     |     |     |    |
| 4. <code>addi x0 x0 4 -&gt; nop</code> |    |    |    | IF  | X  | X  | X  | X   |     |     |    |
| 7. <code>addi t0 x0 9</code>           |    |    |    |     | IF | ID | EX | MEM | WB  |     |    |
| 8. <code>addi t1 t0 2</code>           |    |    |    |     |    | IF | ID | EX  | MEM | WB  |    |
| 9. <code>xori t1 x0 6</code>           |    |    |    |     |    |    | IF | ID  | EX  | MEM | WB |

**Q4 Piracy: arr!****(11 points)**

For Q4.1, assume that we have a 32-bit address space with a 32KiB, 8-way associative cache with a block size of 512B.

Q4.1 (3 points) Calculate the TIO bits with this cache setup.

**Solution:** 20/3/9

The offset equals  $\log_2(\text{block size}) = 9$ . Then, the index equals  $\log_2\left(\frac{\text{cache size}}{\text{block size} \cdot \text{associativity}}\right) = 3$ . Finally, the tag is equal to address length  $-$  index  $-$  offset = 20.

For Q4.2 to Q4.5, assume that we have a 16 byte, fully associative cache with 4B blocks. For hit rates, please express your answer as a simplified fraction.

```
#define ARRAY_SIZE 6
int main() {
    int32_t arr[ARRAY_SIZE];

    for(int i = 0; i < ARRAY_SIZE; i++) {
        arr[i] += arr[0];
        arr[i] += arr[1];
        arr[i] += arr[2];
        arr[i] += arr[3];
    }
}
```

Q4.2 (1.5 points) What is the hit rate for the **first** iteration of the **for** loop, using an **LRU** replacement policy?

**Solution:** 2/3. We have 12 accesses total per iteration after splitting up the line `arr[i] += arr[0]` into `arr[i] = arr[i] + arr[0]`, and likewise with the other lines. The first access to `arr[0]`, `arr[1]`, `arr[2]` and `arr[3]` result in compulsory misses, but the other accesses hit, meaning that we have a total of 8 hits out of 12 accesses, resulting in a 2/3 hitrate.

Q4.3 (1.5 points) What is the hit rate for the **first** iteration of the **for** loop, using an **MRU** replacement policy?

**Solution:** 2/3, for the same reason as Q4.2.

Q4.4 (2.5 points) What is the hit rate for the **last** iteration of the **for** loop, using an **LRU** replacement policy?

**Solution:** 7/12. The key insight here is to notice that once we start accessing `arr[4]` and greater, we begin to have to evict members. For ease of explanation, we consider the 5th iteration of the loop below. Listing out the 12 accesses, we have the below cache state at the start of the iteration:

| Data                | LRU |
|---------------------|-----|
| <code>arr[0]</code> | 3   |
| <code>arr[1]</code> | 2   |
| <code>arr[2]</code> | 1   |
| <code>arr[3]</code> | 0   |

1. Read `arr[4]`. This evicts `arr[0]`, resulting in a miss.
2. Read `arr[0]`. This evicts `arr[1]`, resulting in a miss.
3. Write to `arr[4]`. This results in a hit.
4. Read `arr[4]`. This results in a hit.
5. Read `arr[1]`. This evicts `arr[2]`, resulting in a miss.
6. Write to `arr[4]`. This results in a hit.
7. Read `arr[4]`. This results in a hit.
8. Read `arr[2]`. This evicts `arr[3]`, resulting in a miss.
9. Write to `arr[4]`. This results in a hit.
10. Read `arr[4]`. This results in a hit.
11. Read `arr[3]`. This evicts `arr[0]`, resulting in a miss.
12. Write to `arr[4]`. This results in a hit.

Counting our hits, we have 7 hits and 5 misses, leading to our above hitrate.

Q4.5 (2.5 points) What is the hit rate for the **last** iteration of the **for** loop, using an **MRU** replacement policy?

**Solution:** 3/4. Once again, the key insight here is to notice that once we start accessing `arr[4]` and greater, we begin to have to evict members. For ease of explanation, we consider the 5th iteration of the loop below. Listing out the 12 accesses, we have the below cache state at the start of the iteration:

| Data                | LRU |
|---------------------|-----|
| <code>arr[0]</code> | 3   |
| <code>arr[1]</code> | 2   |
| <code>arr[2]</code> | 1   |
| <code>arr[3]</code> | 0   |

1. Read `arr[4]`. This evicts `arr[3]`, resulting in a miss.
2. Read `arr[0]`. This results in a hit.
3. Write to `arr[4]`. This results in a hit.
4. Read `arr[4]`. This results in a hit.
5. Read `arr[1]`. This results in a hit.
6. Write to `arr[4]`. This results in a hit.
7. Read `arr[4]`. This results in a hit.
8. Read `arr[2]`. This results in a hit.
9. Write to `arr[4]`. This results in a hit.
10. Read `arr[4]`. This results in a hit.
11. Read `arr[3]`. This evicts `arr[4]`, resulting in a miss.
12. Write to `arr[4]`. This evicts `arr[3]`, resulting in a miss.

Counting our hits, we have 9 hits and 3 misses, leading to our above hitrate.

**Q5 Even Greater Odds****(13 points)**

Consider the following C function, which returns `true` if the given array has more even elements than odd elements, or `false` otherwise. You may assume that the array only contains strictly positive integers, that the variables `even_count` and `odd_count` will not overflow, and that all necessary C library header files are included.

```
1 bool more_even(uint32_t* array, uint32_t n) {
2     uint32_t even_count = 0;
3     uint32_t odd_count = 0;
4     for (uint32_t i = 0; i < n; i++) {
5         if (array[i] % 2 == 0) {
6             even_count++;
7         } else {
8             odd_count++;
9         }
10    }
11    return even_count > odd_count;
12 }
```

You have access to the following SIMD operations. A vector is a 128-bit vector register capable of holding 4 32-bit unsigned integers.

- `vector vec_load(uint32_t* A)`: Loads 4 integers at memory address `A` into a vector
- `vector vec_setnum(uint32_t num)`: Creates a vector where every element is equal to `num`
- `vector vec_and(vector A, vector B)`: Computes the bitwise AND between each pair of corresponding vector elements in `A` and `B`, and returns a new vector with the result
- `vector vec_or(vector A, vector B)`: Computes the bitwise OR between each pair of corresponding vector elements in `A` and `B`, and returns a new vector with the result
- `vector vec_xor(vector A, vector B)`: Computes the bitwise XOR between each pair of corresponding vector elements in `A` and `B`, and returns a new vector with the result
- `vector vec_add(vector A, vector B)`: Adds `A` and `B` together elementwise, and returns a new vector with the result
- `uint32_t vec_sum(vector A)`: Adds all elements of the vector together, and returns the sum

(Question 5 continued...)

Fill in the blanks below to finish the implementation of `more_even_simd`. Assume that your code for this subpart is only required to work on inputs where `n` (the length of `array`) is a multiple of 4. You may only use up to one SIMD operation per blank.

```
1 bool more_even_simd(uint32_t* array, uint32_t n) {
2     vector counts = vec_setnum( 0 );
                                   Q5.1
3     vector mask = vec_setnum(1);
                                   Q5.2
4     for (uint32_t i = 0; i < n / 4 * 4; i+=4) {
                                   Q5.3
5         vector temp = vec_load(array + i);
                                   Q5.4
6         vector masked = vec_and(temp, mask);
                                   Q5.5
7         counts = vec_add(counts, masked);
                                   Q5.6
8     }
9     return (n / 2) >
                                   Q5.7
10         (vec_sum(counts));
                                   Q5.8
11 }
```

(Question 5 continued...)

For Q5.9 to Q5.10, consider the following implementations of `more_even` that use thread-level parallelism. Assume that there are no syntax errors.

For each implementation, determine whether the implementation is...

- ...correct and faster than the naive `more_even`
- ...correct and slower than the naive `more_even`
- ...incorrect.

A correct implementation is one that will always return the same value as the naive `more_even` function.

You should evaluate the performance of each implementation as the array size approaches infinity (in other words, when the array is really really large). You may assume that the machine has 16 cores and OpenMP uses 16 threads.

If you choose “correct and slower” or “incorrect”, please justify your answer. We will only read the first 15 words of each justification.

Q5.9 (2 points)

```
1 bool more_even_pragma(uint32_t* array, uint32_t n) {
2     uint32_t even_count = 0;
3     uint32_t odd_count = 0;
4     #pragma omp parallel for
5     for (uint32_t i = 0; i < n; i++) {
6         if (array[i] % 2 == 0) {
7             even_count++;
8         } else {
9             odd_count++;
10        }
11    }
12    return even_count > odd_count;
13 }
```

- Correct and faster than the naive `more_even`
- Correct and slower than the naive `more_even`
- Incorrect

**Solution:** There is a race condition through multiple threads accessing `even_count` and `odd_count` without critical sections or reductions.

Q5.10 (2 points)

```
1 bool more_even_pragma(uint32_t* array, uint32_t n) {
2     uint32_t even_count = 0;
3     uint32_t odd_count = 0;
4     #pragma omp parallel
5     {
6         uint32_t evens;
7         uint32_t odds;
8         for (uint32_t i = 0; i < n; i++) {
9             if (array[i] % 2 == 0) {
10                evens++;
11            } else {
12                odds++;
13            }
14        }
15        #pragma omp critical
16        {
17            even_count = evens;
18            odd_count = odds;
19        }
20    }
21    return even_count > odd_count;
22 }
```

- Correct and faster than the naive `more_even`
- Correct and slower than the naive `more_even`
- Incorrect

**Solution:** There is no declaration of parallelizing the `for` loop via a `#pragma omp parallel for` directive. However, since each thread completes all of the work, the direct assignment on lines 17 and 18 give us the correct answer.



**Q6 Mixed Messages****(7 points)**

A new data center has to perform  $2N$  tasks using its  $C$  cores. These tasks, indexed  $0$  to  $2N - 1$  inclusive, are independent, except that each task  $N + i$  must be done after task  $i$  has been completed. We decide to use the manager-worker approach with one process per core, where the manager keeps track of a queue `task_queue` containing all unassigned tasks that can be started. We'll use process  $0$  as the manager.

Throughout this question, you must minimize the resources consumed by the data center:

- Do not send unnecessary messages.
- Terminate unused processes as soon as possible.

Assume that each process enters the main loop simultaneously, and that processes will not crash during execution. We implement this task with the following messages:

- `EXECUTE(x)`, where  $0 \leq x < 2N$ .
- `EXIT`
- `DONE(x)`, where  $-1 \leq x < 2N$ . `DONE(-1)` indicates that a worker is ready but has not performed a task.

A worker should have the following behavior:

Q6.1 (1 point) Before the main loop...

- Perform task  $0$
- Cleanup and exit process
- Send the `DONE(-1)` message to process  $0$
- Do nothing

**Solution:** Since `DONE(-1)` is defined to indicate that a worker is ready but has not performed a task, we send this message when a worker is about to start.

Q6.2 (1 point) Within the main loop, if we receive an `EXECUTE(x)` message...

- Perform task  $x$
- Perform task  $x$ , then cleanup and exit process
- Perform task  $x$  and send the `DONE(x)` message to process  $0$
- Perform task  $x$  and send the `DONE(x)` message to process  $x$
- Do nothing

**Solution:** We must include the task number in the done message in order to be able to handle prerequisites.

(Question 6 continued...)

Q6.3 (1 point) Within the main loop, if we receive an **EXIT** message...

- Perform task 0
- Cleanup and exit process
- Send the **DONE(-1)** message to process 0
- Do nothing

**Solution:** We clean up and exit upon receiving the **EXIT**.

A manager should have the following behavior:

Q6.4 (1 point) Before the main loop, initialize **task\_queue** to be...

- An empty list
- A list containing tasks 0 through  $N - 1$  in order
- A list containing tasks 0 through  $2N - 1$  in order

**Solution:** We must initialize the task list to only have the tasks with no prereqs.

Q6.5 (1 point) Within the main loop, if we receive a **DONE(x)** message from process **p**...

- Perform task **x**
- Add task **x** to **task\_queue**
- If  $0 \leq x < N$ , add task  $N + x$  to **task\_queue**
- Do nothing

**Solution:** Once a task is completed, we can add the task that depended on it to the queue.

Q6.6 (1 point) Immediately after Q6.5, if the **task\_queue** is not empty...

- Remove the first task **t** in **task\_queue**, send **EXECUTE(t)** to process 0
- Remove the first task **t** in **task\_queue**, send **EXECUTE(t)** to process **p**
- Remove the first task **t** in **task\_queue**, send **EXECUTE(N+t)** to process 0
- Remove the first task **t** in **task\_queue**, send **EXECUTE(N+t)** to process **p**
- Send **EXIT** to process **p**
- Do nothing

**Solution:** We assign the first task available to the now ready thread.

(Question 6 continued...)

Q6.7 (1 point) Immediately after Q6.5, if the `task_queue` is empty...

- Send **EXIT** to process **p**
- If we have assigned all  $2N$  tasks, send **EXIT** to process **p**, otherwise do nothing
- Do nothing

**Solution:** Since we must terminate cores as soon as possible, we must immediately send the core an **EXIT** message once the queue is empty. For example, if  $N=5$  and task 1 was some long-running task, then task 6 would not show up in the queue until task 1 completed. If we only terminated cores once all tasks have been assigned, then all cores would have to wait for 1 to finish so that task 6 can be assigned.

Alternate explanation: We can immediately ask process **p** to exit immediately, since it is guaranteed that the process that completed task  $N$  can immediately complete task  $2N$  afterwards, so there's no need to keep around processes when the queue is empty.

Once all workers have received an **EXIT**, the manager should exit.

**Q7 The Fault in Our Pages**

**(9 points)**

Q7.1 (3 points) Suppose we have a 48-bit address space with 32 GiB of physical memory and a 2 MiB page size. How many bits are in our page offset, physical page number, and virtual page number?

**Solution:** 21/14/27

Our offset comes from  $\log_2(\text{page size}) = 21$ . Then, the length of a physical address comes from  $\log_2(\text{physical memory size}) = 35$ . Recalling that the structure of a VA is **VPN|Offset** and the structure of a PA is **PPN|Offset**, we get that there are 27 VPN bits and 14 PPN bits.

Regardless of your answers to Q7.1, now assume we have a 48-bit address space using 24-bit page offsets, 24-bit virtual page numbers, and 12-bit physical page numbers. The system also has a TLB. The TLB and a subset of the page table are shown below. You may assume that the next physical page to be allocated has PPN 0x123, and that all accesses are independent of each other.

| TLB   |       |           |       |
|-------|-------|-----------|-------|
| Dirty | Valid | VPN       | PPN   |
| 1     | 0     | 0x00 0000 | 0x203 |
| 0     | 1     | 0x00 0003 | 0x168 |
| 0     | 1     | 0x00 0002 | 0x164 |
| 1     | 0     | 0x61 C002 | 0x727 |
| 1     | 1     | 0x61 B001 | 0xC8E |
| 0     | 0     | 0x0F 100F | 0xE02 |
| 0     | 1     | 0x61 C001 | 0xAF4 |
| 1     | 1     | 0x00 0001 | 0x162 |

| Page Table |       |       |       |
|------------|-------|-------|-------|
| Index      | Dirty | Valid | PPN   |
| 0x00 0000  | 1     | 1     | 0x161 |
| 0x00 0001  | 1     | 1     | 0x162 |
| 0x00 0002  | 0     | 1     | 0x164 |
| ...        |       |       |       |
| 0x61 C000  | 0     | 0     | 0x625 |
| 0x61 C001  | 0     | 1     | 0xAF4 |
| 0x61 C002  | 1     | 0     | 0x727 |
| ...        |       |       |       |

For each of the following virtual addresses, translate it into a physical address and determine what will happen if we access this address.

Q7.2 (2 points) 0x0000 00AB ACAB

**Solution:** We split the address into VPN: 0x00 0000 and Offset: 0xAB ACAB. Then, we look in the TLB and extract the following entry:

| Dirty | Valid | VPN       | PPN   |
|-------|-------|-----------|-------|
| 1     | 0     | 0x00 0000 | 0x203 |

Since this is not valid, we go to the page table and extract the following entry:

| Index     | Dirty | Valid | PPN   |
|-----------|-------|-------|-------|
| 0x00 0000 | 1     | 1     | 0x161 |

Since this is valid, we have a page table hit, and grabbing the PPN, we have 0x161ABACAB.

- TLB hit
- TLB miss and page table hit
- Page fault

Q7.3 (2 points) 0x61C0 02B1 ADE2

**Solution:** We split the address into VPN: 0x61 C002 and Offset: 0xB1 ADE2. Then, we look in the TLB and extract the following entry:

| Dirty | Valid | VPN       | PPN   |
|-------|-------|-----------|-------|
| 1     | 0     | 0x61 C002 | 0x727 |

Since this is not valid, we go to the page table and extract the following entry:

| Index     | Dirty | Valid | PPN   |
|-----------|-------|-------|-------|
| 0x61 C002 | 1     | 0     | 0x727 |

Since this is also not valid, we have a page fault, and assign the next PPN, 0x123. Thus, we have 0x123B1ADE2.

- TLB hit
- TLB miss and page table hit
- Page fault

Q7.4 (2 points) 0x61B0 01FE 3121

**Solution:** We split the address into VPN: 0x61 B001 and Offset: 0xFE 3121. Then, we look in the TLB and extract the following entry:

| Dirty | Valid | VPN       | PPN   |
|-------|-------|-----------|-------|
| 1     | 1     | 0x61 B001 | 0xC8E |

Since this is valid, we have a TLB hit, and we grab our PPN and create the address 0xC8EFE3121.

- TLB hit
- TLB miss and page table hit
- Page fault

**Q8 Cumulative: Unconditional RISC****(9 points)**

Assume we have a function,  $f$ , that takes in a 32-bit unsigned integer,  $x$ , as an argument.  $f(x)$  is defined as follows:

$$f(x) = \begin{cases} x + 9 & \text{if } x \% 4 == 0 \\ x * 2 & \text{if } x \% 4 == 1 \\ x & \text{if } x \% 4 == 2 \\ x // 8 & \text{if } x \% 4 == 3 \end{cases}$$

Jero wants to write this function in RISC-V, but he couldn't get his CS61CPU's branch instructions to work! As a result, you may use any RV32I instruction **except** branch instructions.

Write a function,  $f$ , which accepts one argument  $x$  in  $a0$ , and returns  $f(x)$ . You may assume that there is no overflow.

```
1 f:
2  andi t3 a0 3
   Q8.1
3  slli t3 t3 3
   Q8.2
4  auipc t7 0
   Q8.3
5  add t3 t3 t7
6  jalr x0 t3 12
   Q8.4
7  addi a0 a0 9
   Q8.5
8  j f_end
9  slli a0 a0 1
   Q8.6
10 j f_end
11 nop
   Q8.7
12 j f_end
13 srli a0 a0 3
   Q8.8
14 f_end:
15 ret
```

**Q9 Cumulative: Orion's Hell****(8 points)**

Having failed to completely tame the CS61Cerberus, Heracles has been stuck in Hades for the past two months. The goddess of knowledge Athena decides to help, by letting Heracles cast a spell on Orion.

**Warning: This question is significantly harder than any other problem Heracles has faced so far.**

Recall from the midterm the following:

- Orion has a favorite number, represented as an  $n$ -bit integer. In order to tame Orion, Heracles must determine Orion's favorite number (through his `solve_orion` function) by calling the `orion` function on various numbers, and observing the results.
- Originally, the `orion` function works by performing a bitwise OR on the input and Orion's favorite number, and returning 1 (true) if the result is 0, and 0 (false) otherwise.

The below is the compiled (RISC-V RV32I) code of Orion's `orion` function:

```
orion: # Input is received in a0, and the result is outputted in a0
    li a1 ORNUM # Orion's favorite number omitted
    or a0 a0 a1
    bne a0 x0 FalseCase
    addi a0 x0 0
    jr ra
FalseCase:
    addi a0 x0 1
    jr ra
```

Athena's spell can change **one bit** in the assembled bytecode of the `orion` function; Heracles can then use the modified `orion` function instead in his `solve_orion` function. The spell only lasts for a short time, so `solve_orion` will now have stricter asymptotic runtime requirements.

After the single bit change, the `orion` function must still be valid RISC-V code that observes calling convention, and must still work regardless of any other code (e.g. you can't load/store to unknown memory, specify the value of any register except `a0`, or jump out of the `orion` function in undefined manners).

(Question 9 continued...)

Q9.1 Select one bit in the `orion` function to flip, which will allow Heracles to determine Orion's favorite number. In addition, write the `solve_orion` function (in C), which will work given your modified `orion` function.

For full credit, your `solve_orion` must run in  $O(1)$  time relative to the number of bits in Orion's favorite number. For 75% credit, your solution may run in  $O(n)$  time instead.

Flip bit  of the instruction  
to turn it into the instruction .

```
uint32_t solve_orion(uint32_t (*orion)(uint32_t)) {
    // Your code here

}
```

Q9.2 Briefly explain your solution. We will only look at this box if you request a regrade.

**Solution:** Note: This solution assumes familiarity with the solution to CS61Cerberus, found on the Fall 2023 midterm. Multiple solutions were possible. In ascending order of efficiency:

- Flip bit 13 of `or a0 a0 a1` to turn it into `xor a0 a0 a1`. This reduces this problem to solving Luxor (see the midterm), which takes  $\Theta(2^N)$  time.
- Flip bit 12 of `or a0 a0 a1` to turn it into `and a0 a0 a1`. This reduces this problem to solving Andy (see the midterm), which takes  $\Theta(N)$  time.
- Flip bit 14 of `or a0 a0 a1` to turn it into `slt a0 a0 a1`. This changes `orion` so that it returns 1 if and only if its input is less than Orion's number. We can binary search this to determine Orion's number in  $\Theta(N)$  time. An example code solution is:

```
uint32_t solve_orion(uint32_t (*orion)(uint32_t)) {
    int32_t min = 0x80000000;
    int32_t max = 0x7FFFFFFF;
    while((uint32_t) (max-min) > 0) {
        int32_t avg = ((uint64_t) min+max)/2;
        if(avg < n) {
            min = avg+1;
        }
        else {
            max = avg;
        }
    }
    return min;
}
```



- Flip bit 7 or 9 of `addi a0 x0 1` to turn it into `addi a1 x0 1` or `addi a4 x0 1`. Either way, this changes `orion` so that it never sets `a0` to 1 when the `or` returns a nonzero value. Therefore, the function will return the bitwise OR of its input and Orion's number. We can thus simply return `orion(0)` to solve this problem.
- Flip bit 9 of `bne a0 x0 FalseCase`, which changes the immediate to jump 2 instructions instead of 3 instructions. This also prevents the `addi a0 x0 1` line from running, thus allowing us to return Orion's number by returning `orion(0)`.
- Flip no bit at all. Since we now know how Orion's function is defined, we can access the first 64 bits of the `orion` function to get the `li` instruction directly, and thus extract ORNUM directly. It is also possible to flip bit 3 of `addi a0 x0 1` to turn it into `auipc a0 256`, which causes Orion to return  $PC+(2^{20})$ , and use that as a relative offset. An example code solution is:

```
uint32_t solve_orion(uint32_t (*orion)(uint32_t)) {
    uint32_t firstInstruction = ((uint32_t*)(orion))[0];
    if((firstInstruction&127)==0x33) {
        return firstInstruction>>20;
    }
    uint32_t secondInstruction = ((uint32_t*)(orion))[1];
    return ((firstInstruction>>12)<<12)+(secondInstruction>>20);
}
```

**Q10** *Cumulative: Art Class*

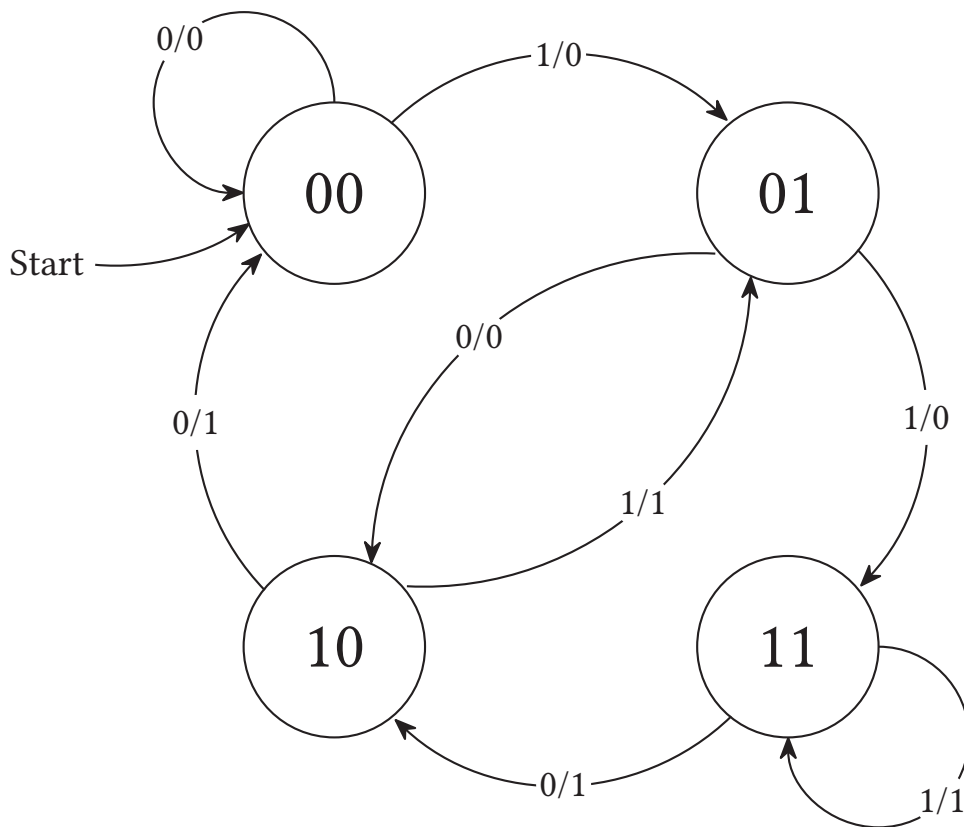
**(16 points)**

You wish to create an FSM whose output at time step  $N$  is 0 for the first two time steps, and the input of time step  $N - 2$  otherwise.

For example, if the input to this FSM was 0b01 1011 1011 1000 1001,

the output should be 0b00 0110 1110 1110 0010.

Q10.1 (8 points) Complete the FSM below. You may not add additional states. Note that you must also label the state transitions we have provided for you.



(Question 10 continued...)

Q10.2 (4 points) Fill in the circuit diagram below to implement this FSM. For full credit, your circuit must have the minimum possible clock period, assuming the following component delays:

$$t_{\text{AND gate}} = 12\text{ps}$$

$$t_{\text{Register clk-to-q}} = 10\text{ps}$$

$$t_{\text{OR gate}} = 15\text{ps}$$

$$t_{\text{Register setup}} = 15\text{ps}$$

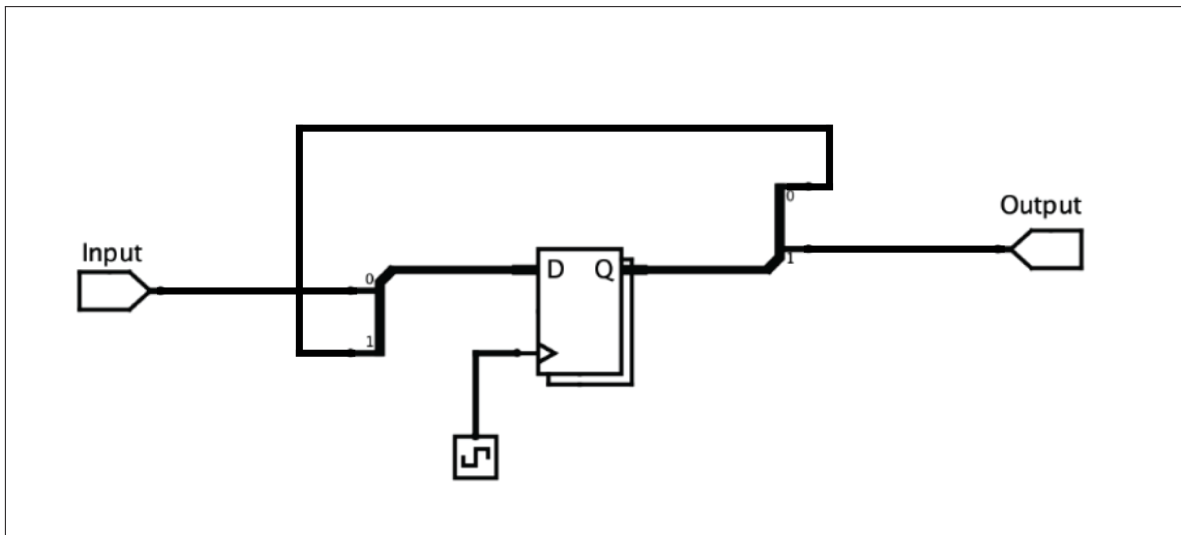
$$t_{\text{NOT gate}} = 4\text{ps}$$

$$t_{\text{Bit splitter}} = 0\text{ps}$$

$$t_{\text{XOR gate}} = 31\text{ps}$$

$$t_{\text{Wire}} = 0\text{ps}$$

You may not use any other components. You may assume that the input and output connect directly to registers (for the purpose of determining the clock period), and that the register stores 2 bits. Your circuit does not need to "match" the states you use in your answer to Q10.1; it will be considered correct if its behavior matches the intended behavior described above.



Unsatisfied with just delaying the input by 2 cycles, you decide to create an FSM that delays the input by 12 cycles (and outputs 0 for the first 12 cycles).

For example, if the input to the new FSM was 0b01 1011 1011 1000 1001,  
you should output 0b00 0000 0000 0001 1011.

Q10.3 (2 points) What is the fewest number of states that an FSM solving this problem can have? Your answer must be an exact **integer**.

**Solution: 4096.** To get this answer, notice that to reference something from 12 inputs ago in an FSM, we must keep a record of that for at least 12 states. In other words, we must "encode" 12 bits of memory into our FSM states, requiring  $2^{12} = 4096$  states.

(Question 10 continued...)

Q10.4 (2 points) What is the minimum clock period of any circuit that solves this problem (assuming the register is expanded to sufficiently many bits without increasing clk-to-q and setup times)?

**Solution: 25ns.** Since the circuit will look exactly the same as the solution above, but extended to 12 bits, (Input  $\rightarrow$  Reg[0], Reg[0]  $\rightarrow$  Reg[1], ..., Reg[11]  $\rightarrow$  Output) the minimum combinatorial path will be 0s, (either from Input to D via a splitter, or Q to Output via a splitter) so our minimum clock cycle equals  $t_{\text{clk-to-q}} + 0 + t_{\text{setup}} = 25\text{ns}$ .

**Q11** *The Finish Line*

**(1 points)**

Everyone will receive credit for this question, even if you leave it blank.

Q11.1 (1 point) How long does it take to “do nothing”?

**Solution:** ~45 minutes plus or minus an eternity.

Q11.2 (0 points) If there’s anything else you want us to know, or you feel like there was an ambiguity in the exam, please put it in the box below.

For ambiguities, you must qualify your answer and provide an answer for both interpretations. For example, “if the question is asking about A, then my answer is X, but if the question is asking about B, then my answer is Y”. You will only receive credit if it is a genuine ambiguity and both of your answers are correct. We will only look at ambiguities if you request a regrade.

**Solution:** ^\\_(\^)\\_/\_^