

CS61C FINAL

<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 Emaan 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 **24** pages including this cover page.
- 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 170 minutes to complete this exam. The exam is closed book: no computers, tablets, cell phones, wearable devices, or calculators. You are allowed three pages (US Letter, double-sided) of *handwritten* notes.
- There may be partial credit for incomplete answers; write as much of the solution as you can and show work in the white spaces provided.
- Please write your answers within the blanks provided within each problem!
- Please clearly bubble in or circle the correct answer choice(s) for multiple choice questions!

Question	1	2	3	4	5	6	7	8	9	10	11	12	13	14	Total
Possible Points	16	8	9	9	8	16	16	6	10	10	10	10	6	6	140

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

Question 1: Main [Memory] Stacks (16 pts)

Recall the idea of the Stack section in our memory hierarchy: when functions are called, their function frames are **pushed** onto the stack; when a function returns, it is **popped** off the stack. We will be implementing the Stack section of memory in software by representing each function frame on the stack as a struct, defined below. Assuming we are running on a **32-bit machine** with **32-bit integers**.

```
typedef struct stack {
    void *frame;           // Pointer to the space allocated for this function frame
    struct stack *prev;    // Pointer to previous stack node
    char *func_name;       // Name of the stack frame's function
    int threads;           // The number of threads below this stack frame
    int in_use;            // The number of threads currently pointing to this stack frame
} StackNode;
```

```
StackNode *sp = NULL;
```

1) What would sizeof (StackNode) return?

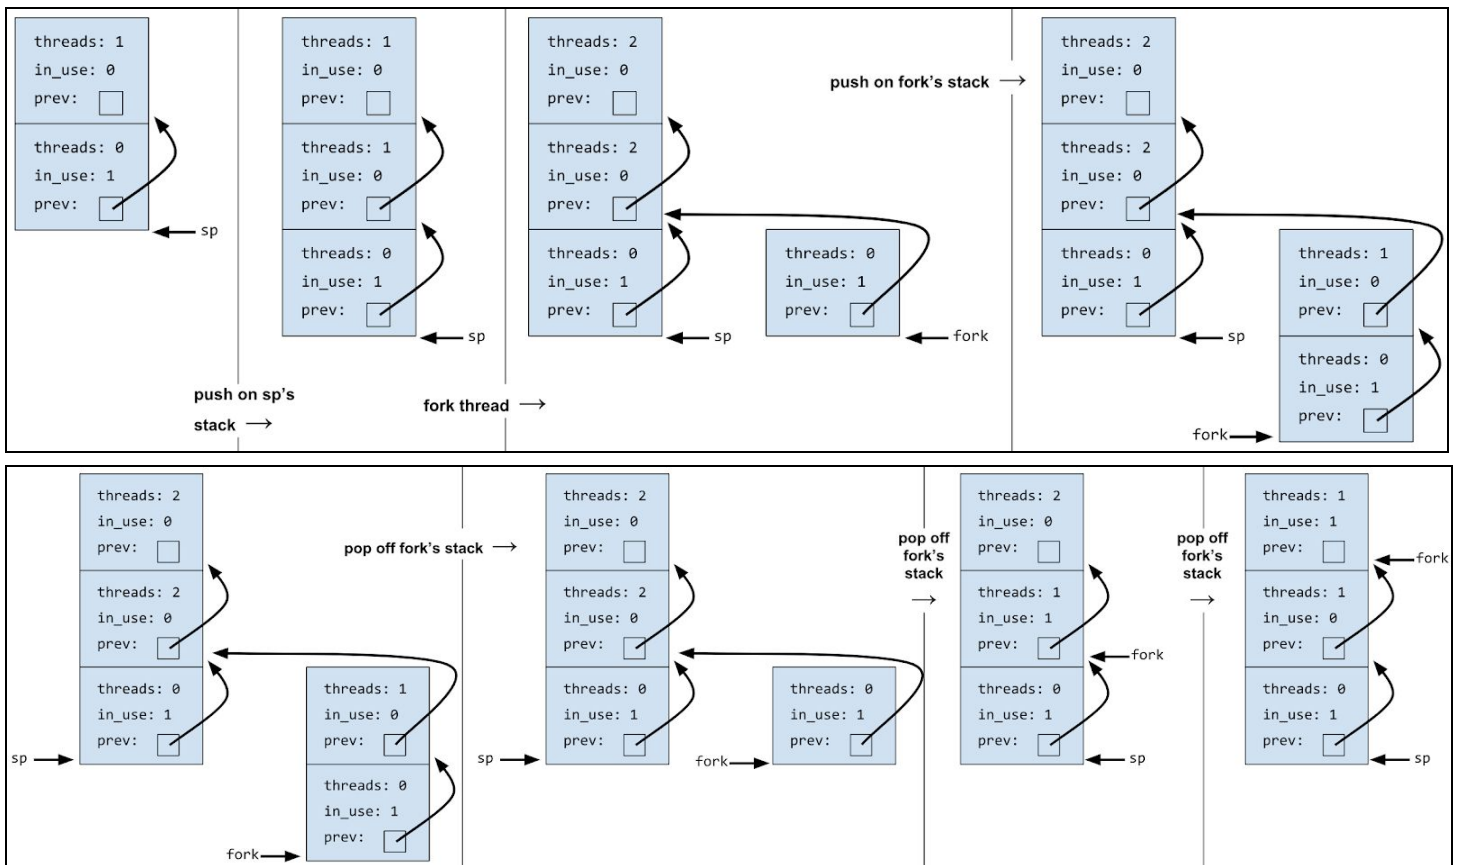
Suppose that the staff has implemented a function frame_size which takes in a C function's name as a string and outputs how many **bytes** that the function's local variables would need in the function frame:

```
int frame_size (char *func_name) {
    \\ Assume this function has been correctly implemented for you.
}
int main_stacks (int argc, char *argv[]) {
    char arr[] = "61C rox";
    return 0;
}
```

2) What would frame_size ("main_stacks") return?

Recall that for a single thread, its Stack section is essentially a linked list where the stack pointer **sp** points to the end of the stack. Also recall that when you **fork** a thread, the new thread will have their own copy of the stack. Instead of actually copying the function frames, we just point to the "shared" stack frames.

On the next page is a visual of the stack that we will be designing. The first set is an example of pushing function frames on the stack along with creating a new fork. The second set is continuously popping off from the fork's stack.



Fill in the following code sequence that will push the function frame onto the stack and modify the sp passed in. Assume sp is never NULL and func_name is always stored in static data—don't allocate space for it.

```

void push (char *func_name, StackNode **sp) {
    StackNode *temp = (StackNode *) malloc(sizeof(StackNode));
    temp->frame = malloc(_____);
    temp->prev = _____;
    temp->func_name = _____;
    temp->threads = _____;
    temp->in_use = _____;
    if (*sp == NULL) {
        _____;
        return;
    }
    if ((*sp)->in_use > 0) {
        (*sp)->threads += 1;
        (*sp)->in_use -= 1;
    } else { /* We are making a fork. */
        for ( _____ ) {
            _____
        }
    }
}

```

Fill out the following code sequence that will pop the function frame off the stack, modifying the `sp` passed in. Also assume that `sp` is never NULL. Since our stack frame is possibly shared by multiple “threads” you will want to consider when we should free the memory for a frame we pop off.

```
void pop (StackNode **sp) {
    printf("%s has been returned\n", (*sp)->func_name);
    (*sp)->in_use--;
    bool free_frame = false;
    if (_____ ) {
        free_frame = true;
    }
    StackNode *temp = (*sp)->prev;
    if (free_frame) {
        _____;
        _____;
    }
    _____;
    if ((*sp) != NULL) {
        _____;
        _____;
    }
}
```

Question 2: Main [Memory] Stacks Management (8 pts)

In this question, we will continue the Stack implementation from Question 1, assuming it functions correctly. Consider the following program (and feel free to draw the stack in the space to the right :D):

```
StackNode *sp = NULL;
int main(int argc, char *argv[]) {
    push("main", &sp);
    push("foo", &sp);
    push("bar", &sp);
    StackNode *fork = sp;
    push("orig", &sp);
    push("split", &fork);
    return 0;
}
```

Each of the following values on the next page evaluate to an address in the C code above. Select the region of memory that these addresses point to right before the main function returns.

- | | | | | |
|------------------|---------|--------|----------|--------|
| 1. main | Ⓐ Stack | Ⓑ Heap | Ⓒ Static | Ⓓ Code |
| 2. &sp | Ⓐ Stack | Ⓑ Heap | Ⓒ Static | Ⓓ Code |
| 3. sp | Ⓐ Stack | Ⓑ Heap | Ⓒ Static | Ⓓ Code |
| 4. *sp | Ⓐ Stack | Ⓑ Heap | Ⓒ Static | Ⓓ Code |
| 5. sp->func_name | Ⓐ Stack | Ⓑ Heap | Ⓒ Static | Ⓓ Code |
| 6. &fork | Ⓐ Stack | Ⓑ Heap | Ⓒ Static | Ⓓ Code |
| 7. argv | Ⓐ Stack | Ⓑ Heap | Ⓒ Static | Ⓓ Code |

Suppose we had a simple recursive function defined as follows:

```
long factorial(long n):
    if (n == 1):
        return 1;
    else:
        return n * factorial(n-1)
```

Assume that function frames only require space for the local variables (i.e. the return value of `frame_size("factorial")`). You are given the following specifications:

Stack and Heap: 16 KiB

Static: 12 KiB

Code: 4 KiB

`frame_size("factorial") = 8`

`sizeof(StackNode) = 56`

Suppose we call the `factorial` function on some number `N` **using our Stack data structure from Question 1** (note: we allocated data for our `StackNode` structs):

```
StackNode *sp = NULL;
int main () {
    push("factorial", &sp);
    push("factorial", &sp);
    push("factorial", &sp);
    ...

    // the N'th call to factorial
    push("factorial", &sp);

    // the first return from factorial
    pop(&sp);
    ...
}
```

What is the smallest value of `N` that will cause a maximum recursion depth error (meaning no more function frames can be created)? Ignore the stack space required for the `main` function. If convenient, put your answer as a power of 2.

`N = _____`

Question 4: Bloating Point (9 pts)

Being the feisty floating point fanatic you are, you devise a new scheme to interpret 32-bit floats. You keep the breakdown of the Sign/Exponent/Significand fields the same. However, there is no implicit 1. or 0. before the significand bits; also, you evaluate the 23 significand bits as an **unsigned number**, which you then multiply with $2^{Exp - Bias}$ and $(-1)^{Sign}$ as you normally would.

$$(-1)^{Sign} \times Significand_{unsigned} \times 2^{Exp - Bias}$$

Exponent	Significand	Value
Largest	Zero	±Infinity
Largest	Nonzero	NaN

where $Bias = 127$.

There are no denormalized numbers.

We walk through one way of representing 3 in this scheme, using a Significand of 0b00...011 = 3
 $(-1)^0 \times 3 \times 2^{127 - Bias} = 1 \times 3 \times 2^0 = 3$ **Sign:** 0, **Exponent:** 127 = 0x7F, **Significand:** 3 = 0x000003

Answer the following questions comparing the standard IEEE Floating Point and your new scheme (let's call it Bloating Point).

- 1) Bloating Point represents a larger amount of unique numeric values than Floating Point. (T) (F)

- 2) 0x3D80001E is a valid representation of 1.875 in Bloating Point. (Hint: $1.875 = 15/8$) (T) (F)

- 3) There exists a 32-bit number whose Floating Point value and Bloating Point value are the same. (T) (F)

4) How many Bloating Point numbers evaluate to $+2^{-124}$?

5) What is the smallest positive number divisible by 5 that Bloating Point cannot represent? You may leave all or part of your answer as a power of 2.

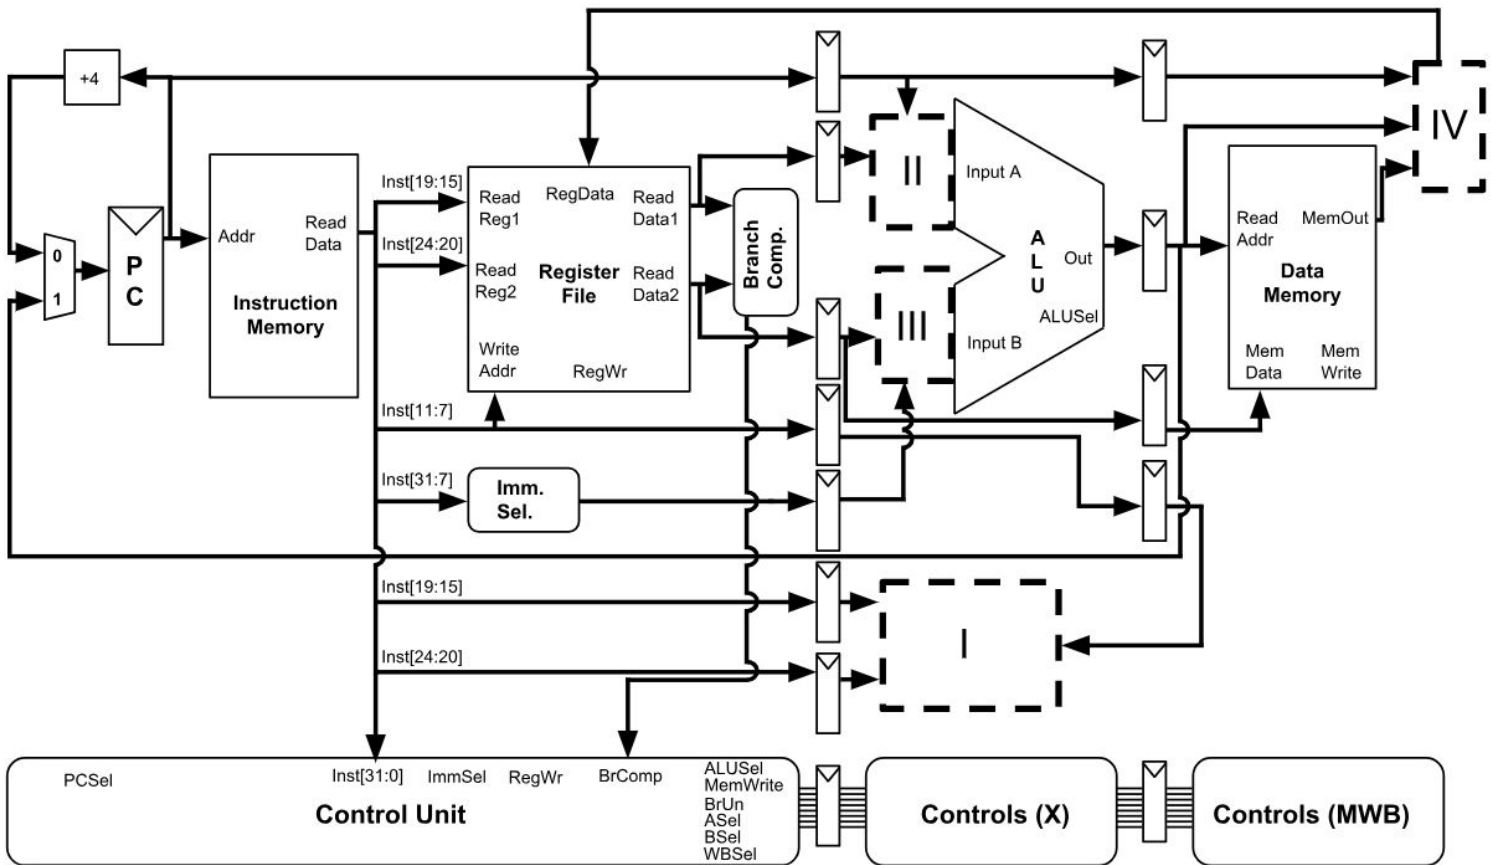
Question 5: DamonPath (8 pts)

Below we have implemented 3-stage CPU with the stages IFD, EX, and MWB. We're interested in implementing forwarding from the output of the MWB stage to the input of the ALU in the EX stage. This datapath should still implement the regular RISC-V instruction set as well as forwarding.

Make sure to read through all parts of the question (I, II, III, IV) as some variables may be defined in different parts.

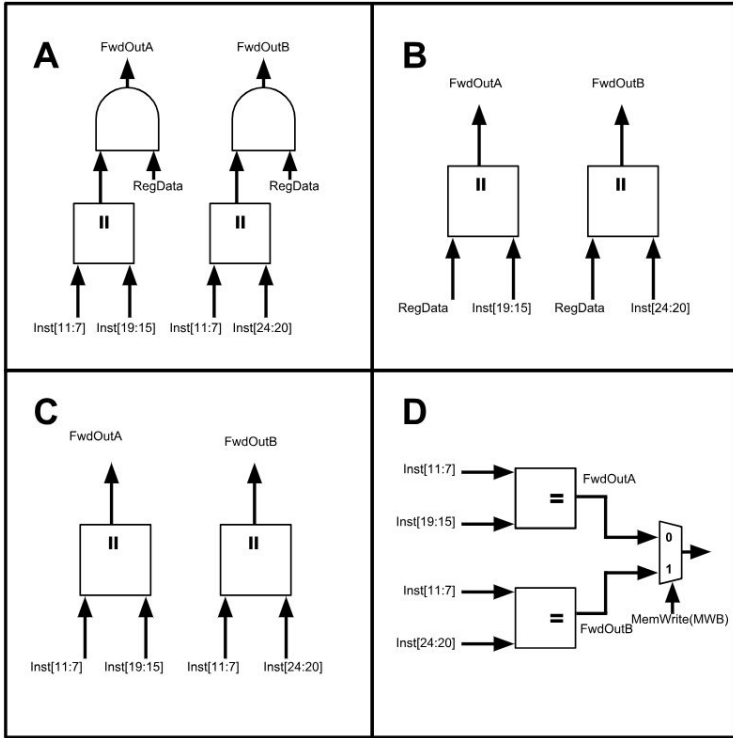
Assume that the control values generated by the control unit are driving their respective control inputs in the datapath.

(Note: In cases that may be ambiguous, we have marked certain values with the stage that they come from. For example, if we write $ASel(X)$ this means the $ASel$ coming from Controls(X) unit.)

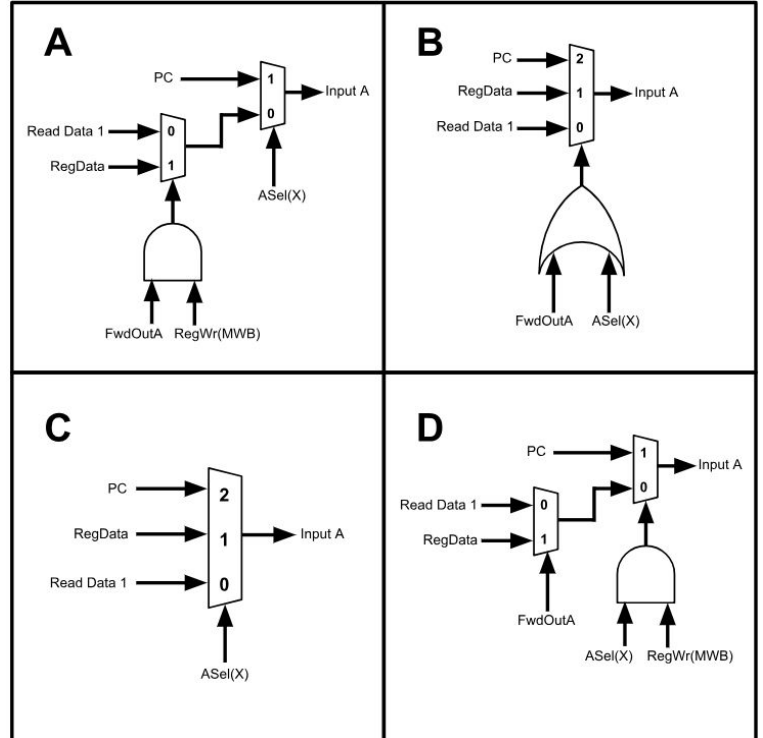


Choose the correct implementation for I - IV from the options below:

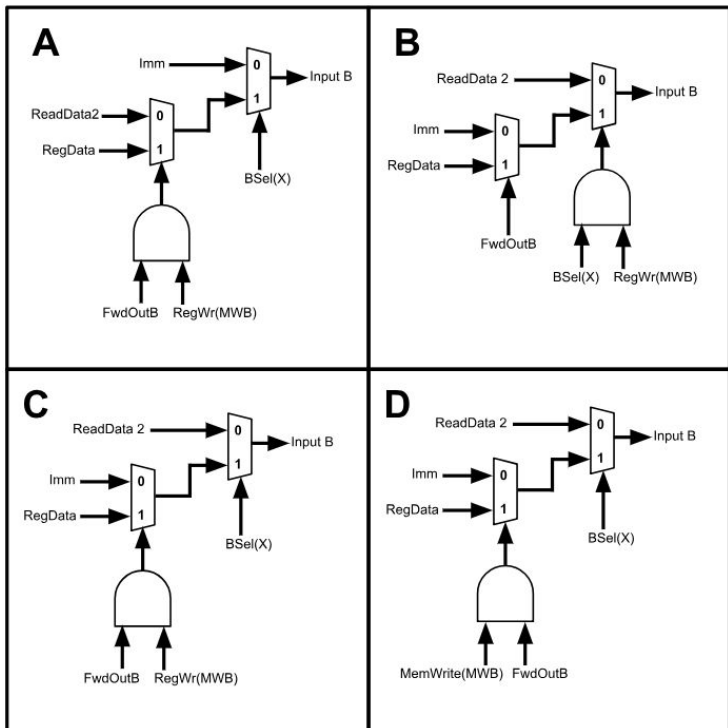
I.



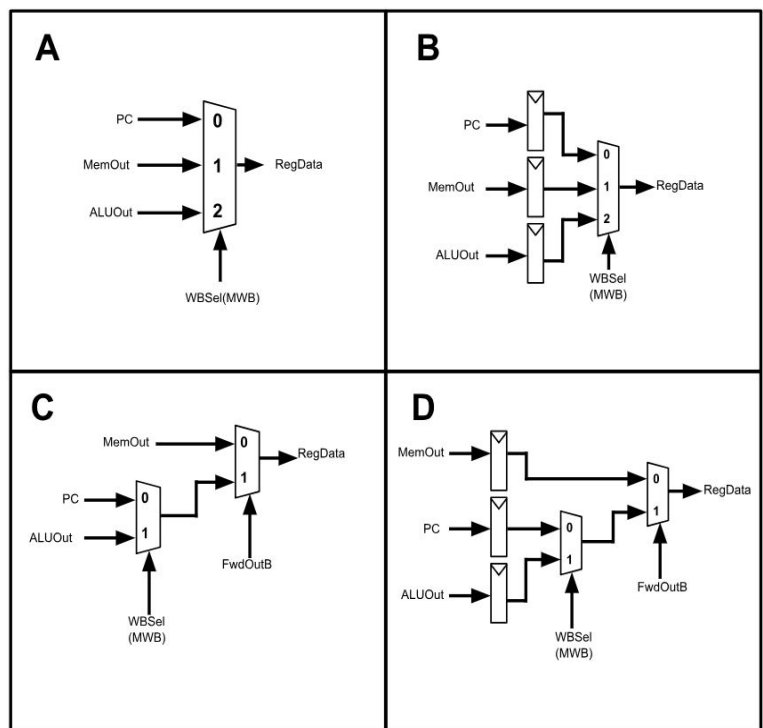
II.



III.



IV.



Question 6: Cache These Hands (16 pts)

A machine with a 19 bit address space has a single 256 B cache. The cache is 4-way set associative with 8 total entries.

1. Determine the number of bits in Tag, Index, and Offset fields for an address on this machine.

Tag: _____ Index: _____ Offset: _____

The following piece of code is executed on the aforementioned machine. This code computes an outer product of a $\mathbf{N} \times \mathbf{1}$ vector A and a $\mathbf{1} \times \mathbf{N}$ vector B, placing the result in a $\mathbf{N} \times \mathbf{N}$ matrix C. Use this code to answer the follow questions about the hit rate the code was produced.

For all questions assume the following:

- `sizeof (double) == 8`
- `A = 0x10000`
- `B = 0x20000`
- `C = 0x30000`
- The cache begins cold before each question.
- Code is executed from left to right.

```
#define N 16
```

```
void outer_product (double *A, double *B, double *C) {
    for (int i = 0; i < N; i++) {
        for (int j = 0; j < N; j++) {
            C [j + i * N] = A[ i ] * B[ j ];
        }
    }
}
```

2. What is the hit rate for executing this code if it uses LRU replacement and is a write back cache with write allocate on a miss? Fill in all blanks for credit.

HR for accesses to A: _____

HR for accesses to B: _____

HR for accesses to C: _____

OVERALL HR: _____

3. What is the hit rate for executing this code if it uses LRU replacement and is a write back cache with no write allocate on a miss? Fill in all blanks for credit.

HR for accesses to A: _____

HR for accesses to B: _____

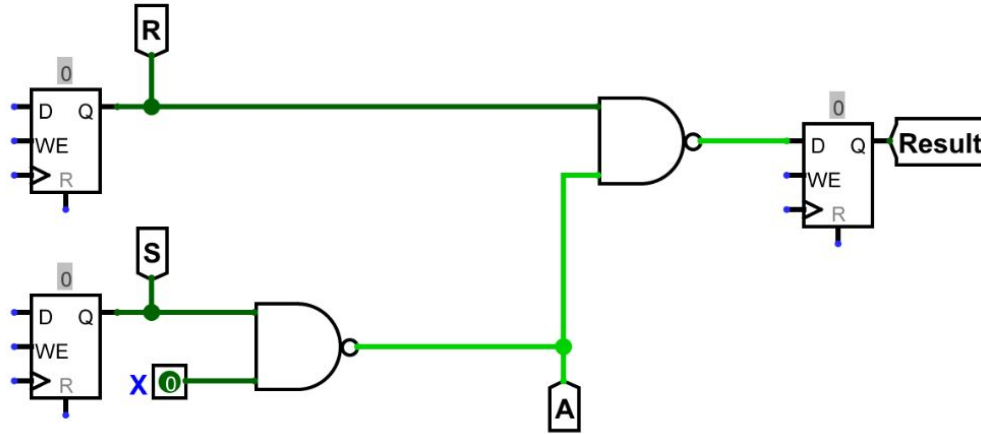
HR for accesses to C: _____

OVERALL HR: _____

Question 7: NANDerson (16 pts)

Part 1: We have three inputs: R, S, and X. R and S are the outputs of the registers shown below, while X appears **instantaneously** at the rising edge of the clock. We also know the following:


- CLK → q = 5ns
- Setup Time = 5ns
- NAND Gate Delay = 3ns
- Assume that hold time is negligible




a) Given the following waveform diagram, list all the times that A changes its value. For example, if A flipped at 20, 40, and 60 ns the answer would be "20, 40, 60"

Anything drawn will be considered as scratch work, and **may** be used for partial credit. (The dotted lines show the rising edge, and each tick is 2.5ns)

Answer: _____




b) Now, we've generated an arbitrary A and Result waveform shown below.




Choose all of the answers below that may be the waveform for **R** such that the correct **Result** waveform is output.

I II III IV V




Part 2: We now connect the “X” output to the output of our second NAND gate with the hope of making something useful. The corresponding diagram is shown below.



- a) In order to determine the functionality of this circuit, we decide to model it as an FSM. In this FSM, **we must flip one of either R or S (and not both) in each time step.**

Represent this as an FSM using the fewest states below. Each transition should be labeled with the bit that you’re flipping, and the output of the FSM where the output is “Result”. It should be of the form: <R or S> / <Result>. (e.g. “R/1”).

We arbitrarily start at state “00”. The first bit represents R and the second bit represents S.



Question 8: Mr. MOESI (6 pts)

For this question you will be asked to determine which cache coherence scheme(s) can fulfill tasks efficiently based upon assumptions of what task our machine must perform. For each question there will be two parts:

- **Expected Behavior:** the behavior that your scheme must perform efficiently.
- **Necessary Behavior:** the behavior that doesn't need to be performed efficiently but must be supported.

You should select **all** schemes that can handle the expected behavior with maximum efficiency.

For all scenarios, we have the following assumptions:

- The machine has multiple cores
- **Writing** to memory is very very slow (a performance bottleneck)

If this is unclear, consider an example. If the expected behavior involves writing and **MSI** and **MOESI** can do fewer writes to memory than **MESI**, then you would select **MSI** and **MOESI** and not **MESI**.

1. **Expected Behavior:** Processing completely read only data.
Necessary Behavior: Nothing additional.

(A) MSI

(B) MESI

(C) MOESI

2. **Expected Behavior:** A single program with threads for different purposes. A single thread will write in the information about a user. Then after this write, the other three threads will in parallel perform different computations based upon the user's data (each written to different, unique location).
Necessary Behavior: Writing to a final shared location may be necessary to accumulate results once all threads have completed.

(A) MSI

(B) MESI

(C) MOESI

3. **Expected Behavior:** Processing a unique program on each core (with its own memory space) and quick reading from memory shared by programs.
Necessary Behavior: Writing to data that is shared across programs.

(A) MSI

(B) MESI

(C) MOESI

Question 9: SIMD Once Told Me (10 pts)

In lecture we showed you how to use Intel Intrinsics to perform an efficient Matrix Multiplication. For this question you are going to perform a similar operation, an inner product between an $m \times n$ matrix \mathbf{A} and an $n \times 1$ vector \mathbf{b} . This formula can be expressed as:

$$\mathbf{C}_i = \sum_{j=1}^n A_{i,j} b_j$$

To make this question more interesting, rather than working with the 128 bit registers we showed in lecture, we will instead work with 256 bit registers. You can select from the following SIMD functions, each of which has a brief explanation.

<code>__m256d _mm256_setzero_pd ()</code>	Produce a <code>__m256d</code> with all 0s.
<code>__m256d _mm256_loadu_pd (double *)</code>	Loads the doubles at the ptr's address into a <code>__m256d</code> .
<code>__m256d _mm256_load_pd (double *)</code>	Loads the doubles at the ptr's address into a <code>__m256d</code> . This is faster, but the ptr passed in must always be divisible by 32 (or else a segfault will occur).
<code>void _mm256_storeu_pd (double *, __m256d)</code>	Stores the contents of the <code>__m256d</code> in the memory location pointed to.
<code>void _mm256_store_pd (double *, __m256d)</code>	Stores the contents of the <code>__m256d</code> in the memory location pointed to. This is faster, but the ptr passed in must always be divisible by 32 (or else a segfault will occur).
<code>__m256d _mm256_fmadd_pd</code> <code>(__m256d A, __m256d B, __m256d C)</code>	Returns $(A * B) + C$ by performing operations on packed doubles.

You may assume that all pointers passed into the function and any stack variables allocated are at addresses divisible by 32. You may also assume `sizeof (double) = 8`. Fill in the function below. You must use the **fastest load and store** whenever possible.

/* Computes an inner product between **SM**, which is a **M x N** matrix, and **SV**, which is
 * is a **N x 1** vector. This result is stored in **D**, which is a **M x 1** vector. */

```
void inner_product_simd (double *d, double *sm, double *sv, unsigned m, unsigned n) {
    __m256d matrix_part;
    __m256d vector_part;
    __m256d total;
    double arr[4];
    for (int i = 0; i < m; i += 1) {
        total = _mm256_setzero_pd ();
        for (int j = 0; j < _____; j += _____) {
            if (_____) { // Check if the address is aligned.
                matrix_part = _____;
            } else {
                matrix_part = _____;
            }
            vector_part = _____;
            total = _____;
        }
        _____;
        d[i] = arr[0] + arr[1] + arr[2] + arr[3];
    }
    if (_____) { // Check if there is a tail case.
        for (int i = 0; i < m; i += 1) {
            for (int j = _____; j < _____; j += _____) {
                _____;
            }
        }
    }
}
```

Question 10: Hashtag Pragma (10 pts)

In OMP we have critical sections which need to be resolved via atomic instructions in assembly. When trying to execute this OpenMP code in multiple threads:

```
#pragma omp critical
{
    x++
}
```

a student produces this assembly to implement it:

```
try: lw t0 0(s0) #load in the value of x from the address
      addi t2 t0 1
      amoswap.w.aq t1, t2, (s0)
      bne t1 t2 try #check if the value of x was the one previously loaded
End:
```

1. This code DOES NOT ensure the critical section works properly. Select **all of the following** options that accurately describes the code the student produced.
 - Ⓐ The code increments x every time it attempts store. This problem can be solved by restoring x to the value of t0 before trying again.
 - Ⓑ Amoswap successfully changes the value at the address of s0 and returns the old value in memory in a single, uninterrupted operation.
 - Ⓒ If instead of a lw, addi, and amoswap, a single amoadd instruction was used to add 1 to x, then the code would work properly.
 - Ⓓ The code can still work properly in some cases.

When working with fully associative caches, you learned that a **cache will search each block in parallel**. Let's envision how this could be simulated in software using openmp. Imagine you have a cache struct that looks like this:

```
typedef struct cache {
    uint32_t *blocks;
    unsigned num_blocks;
} cache;
```

Additionally imagine a cache block stores both data and metadata in a **single 32 bit value** and does so **big endian**. It has the following details:

- Bit 31 (MSB): Valid Bit
- Bit 30: Dirty Bit
- Bits 29 - 16 (14 total bits): Tag
- Bits 15 - 0 (16 total bits): Data

2. Complete the function for handling the **parallel search of each cache block**. It takes in a cache struct, a pointer to data (where data should be stored), and a tag. It should update two shared variables; the first a boolean value indicating if the data was stored in the cache; the second a pointer whose contents should be updated with the data. For this portion **assume we will have as many hardware threads as we do cache blocks**.

```
bool search_cache (cache *cache, uint16_t *data_ptr, uint16_t tag) {
    bool found = false;
    unsigned num_blocks = cache->num_blocks;

    _____

    {
    for (int i = 0; i < num_blocks; i++) {
        uint32_t contents = cache->blocks[i];
        uint32_t valid = contents & (_____);
        if (valid) {
            uint32_t block_tag = (contents >> 16) & (_____);
            if (block_tag == tag) {
                uint32_t data = contents & (_____);
                #pragma omp critical
                {
                    _____;
                    _____;
                }
            }
        }
    }
    }
    return found;
}
```

3. Now let's **assume we may have fewer hardware threads than cache blocks**. Which of the follow statements about the code above represent ways the code can be optimized? Select **all** that are true:
- (A) When a thread finds that it contains the data, it can break out of the loop.
 - (B) As we increase our number of threads we will see little to no speedup because a critical section is included.
 - (C) If our cache works properly, no matter how many threads exist in hardware we can completely remove the critical pragma.
 - (D) If the number of hardware threads is near the number of cache blocks, then we run the risk of having false sharing occur in our code.

Question 11: TL;Br (Too Long; But read) (10 pts)

Consider a machine with 4 KiB pages, a 32-bit virtual address space with 256 MiB of DRAM for main memory. It has a single level of page table and a TLB containing 4 entries which is fully associative.

1. How many bits are there for the VPN? How many bits are there for the offset of the virtual address?

VPN: _____ Offset: _____

2. How many bits are there for the PPN? How many bits are there for the offset of the physical address?

PPN: _____ Offset: _____

Next let's identify how translations of various **virtual addresses** will be resolved. For each translation **identify if the result is a TLB hit, a Page Table Hit, or a Page Fault. Assume each access restarts from the original layout of the TLB and Page Table.** Assume any page table entries not shown have a valid bit of 0.

TLB

VPN	PPN
0x6	0x15
0x4	0x31

PAGE TABLE

VPN	Valid Bit	PPN
0x0	1	0x3
0x1	1	0x7
.....
0x4	1	0x31
0x5	0	0x3
0x6	1	0x15
0x7	1	0x11
.....

3. 0x7ABC
 (A) TLB Hit (B) Page Table Hit (C) Page Fault
4. 0x3000
 (A) TLB Hit (B) Page Table Hit (C) Page Fault
5. 0x6423
 (A) TLB Hit (B) Page Table Hit (C) Page Fault
6. 0x5221
 (A) TLB Hit (B) Page Table Hit (C) Page Fault
7. 0x20282
 (A) TLB Hit (B) Page Table Hit (C) Page Fault

Finally let's consider how the TLB and Page Table change (or don't) as a result of memory accesses. Assume we have **256 B Pages, 2 TLB entries, 4 physical pages and 8 virtual pages in our machine**. These initially appear as shown below. After each access fill in the new contents of the TLB and PT. Assume we evict from main memory and the TLB by evicting the smallest VPN. **Once again assume each access restarts from the original layout of the TLB and Page Table**. If a row in either the TLB or the Page Table does not change from the original, you can either fill it in again or **leave it blank** in the same location.

TLB

VPN	PPN
0x1	0x2
0x5	0x3

PAGE TABLE

VPN	Valid Bit	PPN
0x0	0	0x3
0x1	1	0x2
0x2	0	0x1
0x3	1	0x0
0x4	0	0x0
0x5	1	0x3
0x6	1	0x1
0x7	0	0x2

SID: _____

8. 0x608

TLB

VPN	PPN

PAGE TABLE

VPN	Valid Bit	PPN

9. 0x2F4

TLB

VPN	PPN

PAGE TABLE

VPN	Valid Bit	PPN

Question 12: Go Ham[ming] (10 pts)

We will consider a system that works with 7-bit codewords encoded with Single Error Correction (SEC) parity bits. The Hamming code table is provided for your reference to the right.

Bit	1	2	3	4	5	6	7
Data	P1	P2	D1	P4	D2	D3	D4
P1	X		X		X		X
P2		X	X			X	X
P4				X	X	X	X

Suppose we read a codeword of 0b1010101 from our disk. What are the data bits that were encoded?

0b_____

Suppose two pesky alpha particles struck our disk (darn!) and changed the first two bits of our codeword so that now, we retrieve 0b**01**10101 from our disk. What data bits will be interpreted?

0b_____

What is the Hamming distance between two valid 7-bit SEC codewords? Hint: you can solve this with or without the previous two codewords.

We will now add in a fourth parity bit p_8 to allow our disk to also detect two errors (DED). We read 0b0110101 p_8 from our disk, where p_8 is the DED parity bit for the 7-bit string. What **must** be the value of p_8 such that our 8-bit codeword contains a single correctable error?

$p_8 = \textcircled{1} \textcircled{0}$

What is the Hamming distance between two valid 8-bit DED codewords?

Suppose we have a 5-disk RAID 3 system. Compared to a 5-disk RAID 5 system that is **byte striped**, fill in all the bubbles for the data request patterns / characteristics where RAID 5 **strictly outperforms** RAID 3.

- | | | |
|---|--|--|
| <input type="radio"/> (A) long sequential reads | <input type="radio"/> (B) long sequential writes | <input type="radio"/> (C) recovering from 1 disk failure |
| <input type="radio"/> (D) small random reads | <input type="radio"/> (E) small random writes | <input type="radio"/> (F) capacity |

Consider adding a DMA controller to facilitate bringing in data from our external disk/devices into main memory. The DMA controller can either make transfers in **burst mode**, where it stalls the CPU and writes 32B of data per clock cycle, or **cycle-stealing mode** where it doesn't stall the CPU (i.e. it runs normally) and writes 4B of data per clock cycle. The CPU returns to normal instruction execution after the transfer is completed. Which mode would be better for the following situations? Assume every instruction uses 1 clock cycle with no stalls.

- | | | |
|--|---------------------------------|--|
| • Dealing with page fault; the page size is 4KiB | <input type="radio"/> (A) burst | <input type="radio"/> (B) cycle stealing |
| • Playing a 200MB 5-minute video | <input type="radio"/> (A) burst | <input type="radio"/> (B) cycle stealing |
| • Processing 10 keyboard strokes per second | <input type="radio"/> (A) burst | <input type="radio"/> (B) cycle stealing |

Question 13: South Spark (6 pts)

In this question you will write Spark to find the mode of a list of values and how often it occurs. As a refresher the mode is the number that appears most often. If there is a tie you can select any of the options. **Fill in the blanks for the Python code below.** Use the following Spark Python functions when necessary: **map, flatmap, reduce, reduceByKey**. Here is a sample input and output:

#Input: [1, 2, 1, 2, 3, 4, 5, 6, 4, 2, 1, 3, 3, 1, 1, 2, 2, 1]
 #Output: (1, 6)

```
def output_data (val):
    return _____
```

```
def compute_count (a, b):
    return _____
```

```
def find_max_occurrence (a, b):
    _____
    _____
    _____
    _____
```

```
#values = list (numbers)
modeData = sc.parallelize (values)
```

```
modeData._____ (_____)
._____ (_____)
._____ (_____)
```

Question 14: CALLs well that ends well (6 pts)

For each question select all options which correctly answers the question. The options are

- A. Compiler
- B. Assembler
- C. Linker
- D. Loader
- E. None of the other options

- | | |
|--|---------------------|
| 1. Output includes pseudoinstructions | (A) (B) (C) (D) (E) |
| 2. Is itself a program | (A) (B) (C) (D) (E) |
| 3. Determines the virtual address space for the static section | (A) (B) (C) (D) (E) |
| 4. Initializes the physical address space for the static section | (A) (B) (C) (D) (E) |
| 5. Resolves pseudoinstructions | (A) (B) (C) (D) (E) |
| 6. Fills in the immediate for jump instructions | (A) (B) (C) (D) (E) |