

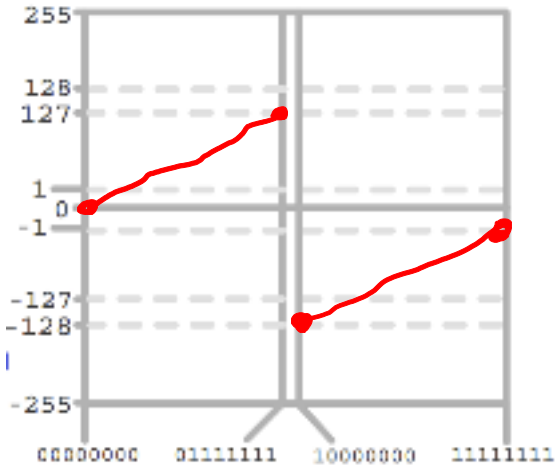
M1-1: I smell a potpourri section covering midterm one... (9 points)

a) Which of the following number representations give $0xFFFFFFFF$ the most positive value when converted to decimal?

Explanation on NEXT Page

- A) Bias (with standard bias) **B) Unsigned** C) Two's complement D) Sign and Magnitude

b) Consider a plot that shows the mapping between 8-bit two's complement binary numbers and their decimal equivalents (i.e. binary is on the x-axis and decimal is on the y-axis). Fill in the plot to the left and answer the following questions.



Translate the numbers on the x-axis to decimal

- i) Fill in the plot to the left.
 ii) Describe (in binary) where discontinuities occur in the plot, if any:
The discontinuity occurs between 01111111 to 10000000, as shown by the graph.

iii) What are the most positive and most negative decimal values that this representation can store?

*Most Positive: 127
 Most Negative: -128*

c) Consider the C code below. Indicate where the values on the right live in memory (using (S)ack, (H)eap, s(T)atic, or (C)ode). Assume no registers are used:

```
#define a 10
int b = 0;

int main(int argc, char** argv) {
    int c = a;
    char d[10];
    int* e = malloc(sizeof(int));
}
```

Explanation on next Page

- a: C
 b: T
 *d: S
 *e: H
 e: S

d) Convert the following instructions from TAL to hex or vice versa. Use register names when possible.

i) `lw $s0, 0($a0)`

→ consider lw s0, 0(a0) I-Type

- opcode: 000 0011

- rs1: a0 → 10 → 01010

- rd : s0 → 8 → 01000

- imm: 0

- funct 3: 010

ii) ~~0x02021021~~



0000 0000 0000 0101 0010 0100 0000 0011

→ 0x00052403

*IGNORE,
old question*

1A) In both signed and two's complement, the number $0x\text{FFFFFFFFE}$ is negative.

In unsigned format, $0x\text{FFFFFFFFE} = 2^{32} - 2$

In Biased format, the bias will be

$$2^{31} - 1, \text{ so } 0x\text{FFFFFFFFE} = (2^{32} - 2) - (2^{31} - 1).$$

This is less than our unsigned representation, yielding our solution.

1C)

a → this is a macro, so the symbol "a" is replaced by 10 by the compiler. Thus, "a" is stored in the code section.

b → this is a global variable, so it is stored on static

*d → d is an array that is locally allocated (within the "main" function), so d refers to a block of memory on the stack

*e → e is a local variable, but its value is an address on the heap. Thus, when we dereference e, the block of memory e refers to is on the heap.

e → e itself is a local variable, so it is located on the stack

M1-2: I'll believe it when I C it (9 points)

Your friend wants to take 61C next semester, and is learning C early to get ahead. They try to implement the ROT13 function as practice, but the code they wrote has some bugs, so you've been called in, as the C expert, to help debug their program, reproduced below:

```

1 /* Applies the ROT13 cipher. rot13("happytimes") == "uncclgvzrf" */
2 void rot13(char *str) {
3     while (*str) {
4         if (str >= 'a' && str <= 'z') {
5             *str = (*str + 13) % 26;
6         }
7         str++;
8     }
9 }
10
11 int main(int argc, char *argv[]) {
12     char a[] = "happy";
13     char b[] = "times";
14     char *s = "XXXXXXXXXXXX"; // 12 X's
15     // Apply cipher to a and b.
16     rot13(a);
17     rot13(b);
18     printf("%s%s\n", a, b);
19
20     // Concatenate and place in s.
21     int i = 0;
22     for (int j = 0; a[j]; ) s[i++] = a[j++];
23     for (int j = 0; b[j]; ) s[i++] = b[j++];
24
25     printf("%s\n", s);
26 }

```

Handwritten annotations:

- Line 4: *we are comparing an address to a number; these should be *str instead*
- Line 5: *we want this to be a letter, so the ascii should be in the range ['a', 'z']. Taking the number mod 26 will not achieve this.*
- Line 14: *s is a pointer to a static string, so it cannot be mutated*
- Line 16: *rot13(a); → makes a refer to "unccl" (if rot13 works)*
- Line 17: *rot13(b); → makes b refer to "gvzrf" (if rot13 works)*
- Line 22: *ERROR: mutating the string in read-only static memory*
- Line 25: *If s is longer than len(a) + len(b), then there will be residue characters here unless we manually insert a NUL byte*

- a) You want to impress your friend, so you predict the result of executing the program as it is written, just by looking at it. If the program is guaranteed to execute without crashing, describe what it prints, otherwise explain the bug that may cause a crash.

The program will crash on line 23 when it tries to modify the string literal referred to by s.

- b) Now, fix all the errors in the program so that it executes correctly. Fill in the corrections you made in the table below. You may not need all the rows.

Line #	Insert Before / Replace / Delete	Change (Explanation or Code)
4	Replace	<code>*str >= 'a' && *str <= 'z'</code>
5	Replace	<code>*str = (*str - 'a' + 13) % 26 + 'a'</code>
14	Replace	<code>char s[] = "XXXXXXXXXXXX"</code>
25	Insert	<code>s[i] = '\0'</code>

M1-3: I don't want to MIPS a thing (9 points)

The following C code recursively sums the elements in an array of length n.

```
int32_t sum_arr(int32_t *arr, size_t n) {
    if (n) {
        return sum_arr(arr + 1, n - 1) + arr[0];
    }
    return 0;
}
```

Translate `sum_arr` into MIPS below. Your code must follow all function calling conventions, and you may not use any pseudoinstructions. You may not need every blank.

`sum_arr:` _____ `non_zero`

`addu $v0, $0, $0`

`jr $ra`

`non_zero:` _____

`addiu $s0, $a0, 0`

`# Store arr into $s0`

_____ `sum_arr`

`# Make the recursive call`

`# Then add arr[0] to the result`

`jr $ra`

RISC-V answers on next page

sum_arr: bne a0, x0, non_zero

addi a0, x0, 0
jr ra } base case

non_zero: addiu SP, SP, -8

sw s0, 0(SP)
sw ra, 4(SP) } Backup saved register & return addr,
as we are making a function call

addi s0, a0, 0 #back up arr

addi a0, a0, 4
addi a1, a1, -1 } Prepare parameters (arr is an int array,
so we do +4)

jalr ra, sum_arr # recursive call

lw t0, 0(s0) # get arr[0]

add a0, a0, t0 # add output of recursive call to arr[0],
put in return value register

lw s0, 0(SP)

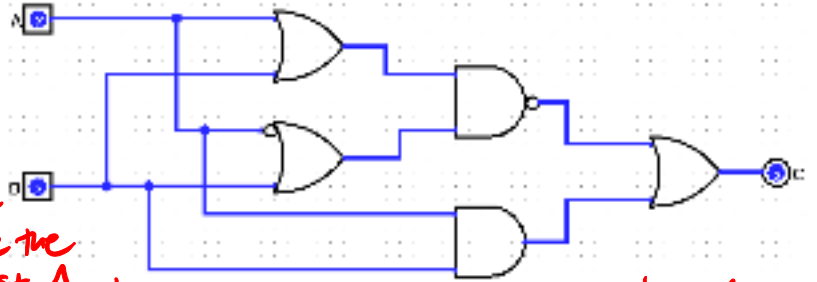
lw ra, 4(SP) } Restore stack and
addi SP, SP, 8 } get back return address

jr ra

M2-1: I couldn't come up with a clever title for SDS. (10 points)

a) Give the simplest Boolean expression for the following circuit in terms of A and B, using the minimum number of AND, OR, and NOT gates:

$(A+B)(\bar{A}+B) + (A\bar{B})$



$= \overline{A+B} + \overline{\bar{A}+B} + AB$ (De Morgan's)

$= \bar{A}\bar{B} + A\bar{B} + AB$ (De Morgan's)

$= \bar{B}(\bar{A}+A) + AB$

$= \bar{B} + AB = \overline{A+\bar{B}}$

This is because if B is 1, the expression is just A, but if B is 0, then the statement is 1

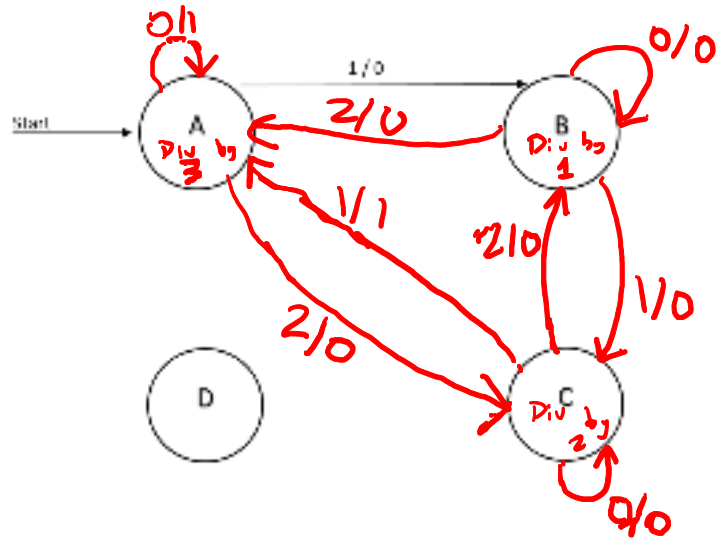
$C = \overline{A+\bar{B}}$

(You must show your work above to earn points.)

b) Using as few states as possible, complete the transition table for an FSM that takes an input with 3 values: 0, 1, or 2. The machine will output a 1 when the sum of the inputs seen so far is divisible by 3. Otherwise it should output a 0.

Assume you have seen no digits at the start state. You might not need all of the states, and you should not draw additional states. You must represent your FSM using the table to the left, **the table is the only part that will be graded.** The first transition has been filled in for you.

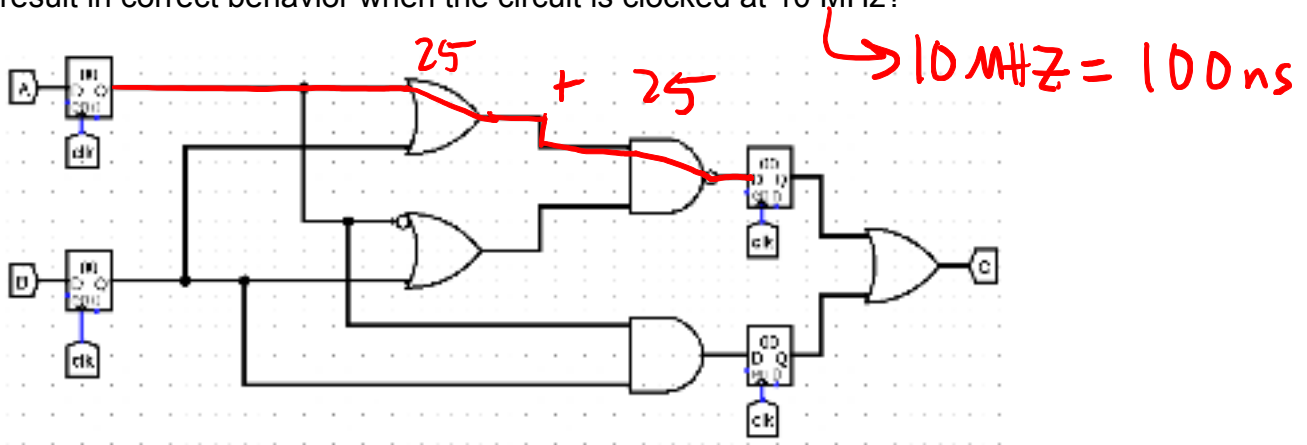
Current State	Input/Output	Next State
A	1/0	B
A	0/1	A
A	2/0	C
B	0/0	B
B	1/0	C
B	2/1	A
C	0/0	C
C	1/1	A
C	2/0	B



A means: sum is divisible by 3
 B means: sum is divisible by 1
 C means: sum is divisible by 2
 D means: unused

M2-1: (continued)

c) Suppose we add registers to the unoptimized circuit in part A to increase the clock rate (this modification is shown below). What is the longest clock-to-Q that the registers on inputs A and B can have that will result in correct behavior when the circuit is clocked at 10 MHz?



Assumptions:

- Assume that clock-to-Q > hold time
- All registers have a setup time of 2 ns
- All logic gates have a delay of 25 ns
- Bubbles on gates do not introduce additional delay

$t_{CLK} = 100 \text{ ns}$

$t_{LOGIC} = 50 \text{ ns}$ b/c critical path has 2 logic gates (as shown on path above)

$t_{CLK} \geq t_{CLK-to-Q} + t_{LOGIC} + t_{setup}$

$100 \geq t_{CLK-to-Q} + 50 + 2$

$t_{CLK-to-Q} \leq 48 \text{ ns} \rightarrow \text{max is } 48 \text{ ns}$

Answer: 48 ns

M2-2: Float like a butterfly and sting like an IEEE (4 points)

Let's take another look at the IEEE754 standard for single-precision floating-point numbers. [x, y) represents a range where x is included and y is not.

a) How many floats are representable in the interval [0.5, 1)? Answer: $1 \cdot 2^{23}$

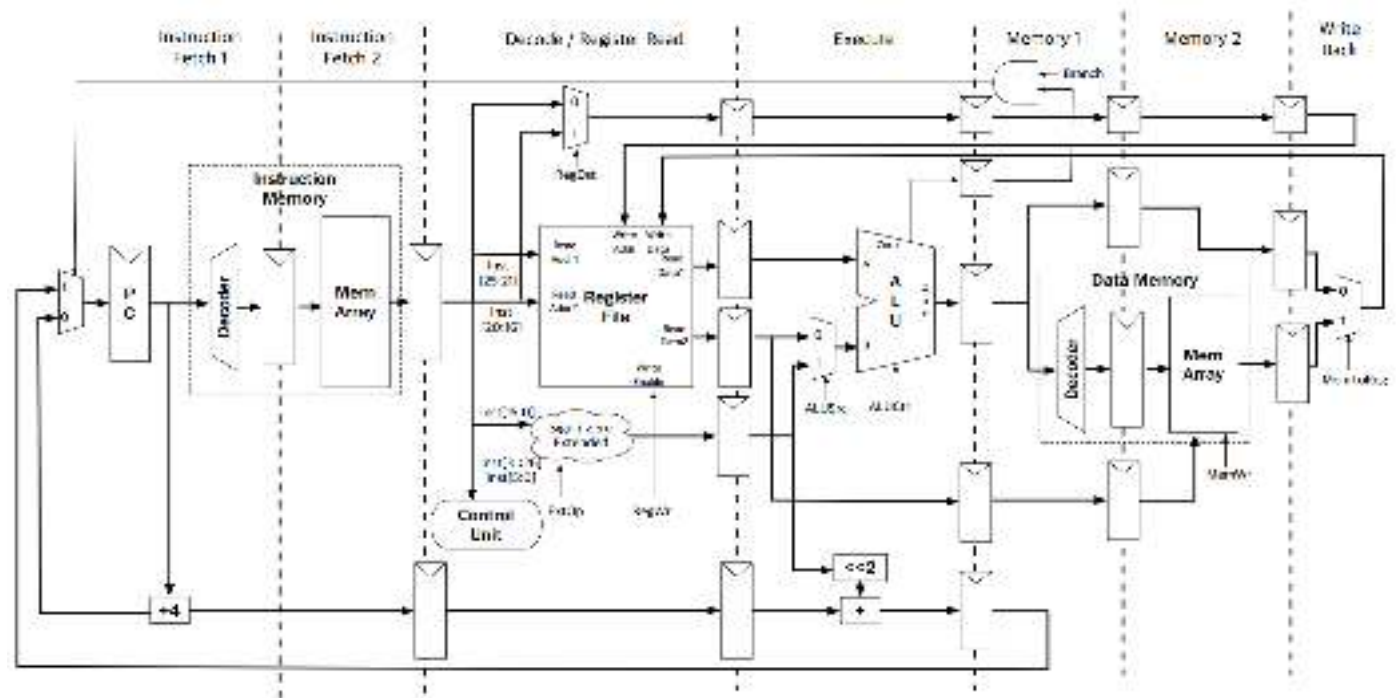
All numbers like $0.1 \underbrace{XXXX \dots X}_{23 \text{ significant field bits}} \rightarrow$ each 'X' can be 0 or 1, so we have 2^{23} in total

b) How many floats are representable in the interval [0, 0.5)? Answer: $126 \cdot 2^{23}$

Denorm numbers $\rightarrow 2^{23}$ of these for each bit of significant field being 0 or 1
 Any normalized float of form $\rightarrow 1.XXXX \dots X \cdot 2^n$ for $n \geq 2$
 As maximum exponent is 127, we have 125 such values of n, and each has 2^{23} value (each significant bit can be 0 or 1)
 Thus, we have $(1 \cdot 2^{23}) + (125 \cdot 2^{23}) = 126 \cdot 2^{23}$

M2-3: If this exam were a CPU, you'd be halfway through the pipeline (8 points)

We found that the instruction fetch and memory stages are the critical path of our 5-stage pipelined MIPS CPU. Therefore, we changed the IF and MEM stages to take **two** cycles while increasing the clock rate. You can assume that the register file is written at the falling edge of the clock.



Assume that no pipelining optimizations have been made, and that branch comparisons are made by the ALU. Here's how our pipeline looks when executing two add instructions:

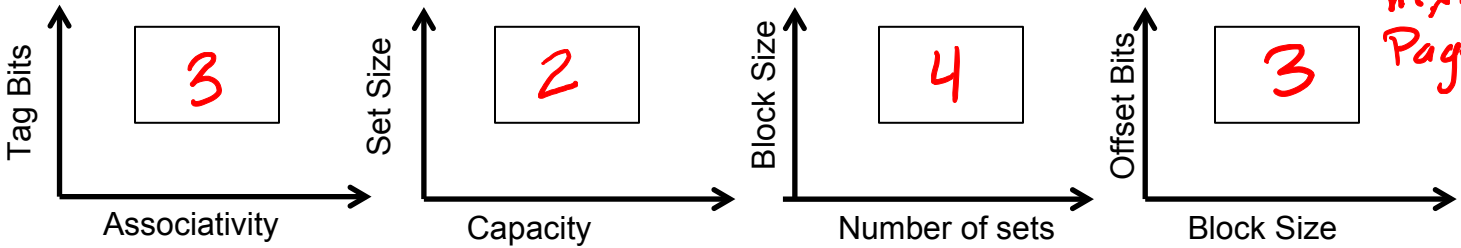
Clock Cycle #	1	2	3	4	5	6	7	8
add \$t0, \$t1, \$t2	IF1	IF2	ID	EX	MEM1	MEM2	WB	
add \$t3, \$t4, \$t5		IF1	IF2	ID	EX	MEM1	MEM2	WB

Make sure you take a careful look at the above diagram before answering the following questions:

- How many stalls would a data hazard between back-to-back instructions require?
- How many stalls would be needed after a branch instruction?
- Suppose the old clock period was 150 ns and the new clock period is now 100ns. Would our processor have a significant speedup executing a large chunk of code...
 - Without any pipelining hazards? Explain your answer in 1-2 sentences.
 - With 50% of the code containing back-to-back data hazards? Explain your answer in 1-2 sentences.

M2-4: Some say there's nothing better than cold, hard cache (12 points)

a) What shape do the following trade-off curves have? Select a shape and enter its number into the box for each of the graphs. Unless they are the parameters being varied, assume that associativity, capacity and block size are constant. You should assume that the axes are linear. *Explanations on next Page*



1: Constant	2: Linear	3: Logarithmic	4: Reverse linear	5: Constant - Logarithm

b) Consider a system with inclusive L1 and L2 caches with 4B cache block size. Assume we have 1 MiB of on-chip memory available and want to determine how much of this memory we should give to the L1 cache and how much to the L2 cache. We will try to minimize the AMAT to do so.

Assume both caches are fully associative with LRU replacement. Their combined capacity is 1MiB (excluding tags and meta-data). **You can consider all miss rates approximate.**

Say you are running the following program starting from cold L1 and L2 caches:

```
#define ARRAY_SIZE 256*1024
int a[ARRAY_SIZE];
int sum = 0; // assume sum, i, and j are stored in registers

for (int i = 0; i < 100000; i++) {
    for (int j = 0; j < ARRAY_SIZE; j++) sum += a[j];
    for (int j = ARRAY_SIZE-1; j >= 0; j--) sum += a[j];
}
```

1) How would we compute AMAT if we had the local L1 miss rate (“L1Miss”), the local L2 miss rate (“L2Miss”) and the memory access time (“Memory”)”? Use “H1” and “H2” to represent the L1 and L2 hit times respectively. (We will compute these quantities later in the question)

$$AMAT = H1 + L1Miss(H2 + LZMiss(Memory))$$

Just use the original AMAT formula with the variables here

a)

Tag Bits - Associativity: The number of tag bits increases by 1 every time the associativity doubles, as doubling the associativity halves the # of sets, thus decreasing the # of index bits by 1. This decrease leads to 1 more tag bit. This relationship is logarithmic.

Capacity - Set Size: Assuming the # of sets (rows) remains constant, doubling the cache capacity will make each set twice as large because we cannot add sets to our cache. This is a linear relationship.

Number of Sets - Block Size: If we double the number of sets, but keep the overall capacity and assoc. the same, we must have half as much data per set. This equates to halving the block size, creating a neg. linear relationship.

Offset Bits - Block Size: We have $2^{\text{offset bits}} = \text{Block Size}$ by definition, so $\text{offset bits} = \log_2(\text{Block Size})$

M2-4: (continued)

- 2) For the program above, express the local miss rate for the L1 cache in general terms as a function of the L1 cache size (write L1 for the size of L1 in bytes). Hint: The miss rate is 0 for a 1 MiB cache, 0.5 for a 0.5 MiB cache and 1 for a 0 MiB (i.e., no) L1 cache.

$$L1\text{ Miss} = 1 - \frac{L1}{1\text{MiB}}$$

If we have 1 MiB L1 cache, the all memory is L1, so we always hit. As we decrease this size, our miss rate increases proportional to the %age of the 1 MiB memory we have in L1.

- 3) What is the global miss rate for the L2 cache as a function of the L1 cache size? Hint: Start by expressing the global miss rate as a function of the L2 cache size.

$$\text{Global Miss 2} = 1 - \frac{L2}{1\text{MiB}} = \frac{1\text{MiB} - L2}{1\text{MiB}} = \frac{L1}{1\text{MiB}}$$

Global miss rate is %age of ALL accesses that miss at L2. Same as frac of mem NOT in L2. $1\text{MiB} - L2 = L1$ b/c $L1 + L2 = 1\text{MiB}$

- 4) What is the local miss rate for the L2 cache as function of the L1 and L2 sizes? Hint: Use your results from questions 2 and 3.

$$\text{Local Miss 2} = \frac{\text{Global Miss 2}}{L1\text{ Miss}} = \frac{L1/1\text{MiB}}{1 - L1/1\text{MiB}} = \frac{L1}{1\text{MiB} - L1} = \frac{L1}{L2}$$

all accesses that miss \rightarrow L2
all accesses that get to L2 \rightarrow L1 + L2 = 1 MiB

- 5) Assume the hit time of the L1 cache is 10 cycles, the hit time of the L2 cache is 20 cycles and the memory access time is 100 cycles. Using the formula from question 1, what is the AMAT for this system as a function of only the L1 size?

$$\begin{aligned} \text{AMAT} &= 10 + \left(\frac{L2}{1\text{MiB}} \right) \left(20 + \frac{L1}{L2} \cdot 100 \right) \quad (\text{from parts 2 \& 4}) \\ &= 10 + \frac{20 \cdot L2}{1\text{MiB}} + \frac{L1}{1\text{MiB}} \cdot 100 = 10 + \frac{20\text{MiB} - 20L1}{1\text{MiB}} + \frac{100L1}{1\text{MiB}} \\ &= 10 + 20 + \frac{100L1 - 20L1}{1\text{MiB}} = \boxed{30 + \frac{80L1}{1\text{MiB}}} \end{aligned}$$

- 6) What sizes of L1 and L2 caches should we pick to minimize the AMAT? (assume the caches have non-zero size, i.e., both of them exist)

We want $30 + \frac{80L1}{1\text{MiB}}$ to be small, so L1 must be as small as possible. Thus, the answer is L1: 4B, L2: 1 MiB - 4B, as neither is zero and both have blocks of 4B (min size is 4B).

F1: Paging all CS61C students (9 points)

Consider a byte-addressed machine with a 13-bit physical address space that can hold two pages in memory. Every process is given 16MiB of virtual memory and pages are evicted with an LRU replacement scheme.

Solve in order of blue circles

a) What are the sizes of the following fields in bits?

Virtual Addresses are $\log_2(16\text{MiB}) = \log_2(2^{24}) = 24\text{ bits}$

④ VPN = Addr Size - offset = $24 - 12$

Virtual Page Number: 12

③ Virtual Address Offset: 12

PPN bits = $\log_2(2) = 1$, "as we have 2 pages" pages are same size

① Physical Page Number: 1

Physical Address Offset: 12

b) Consider the following code snippet:

```
// a and b are both valid pointers to
// different arrays of length ARRAY_SIZE
void enumerate(int* a, int* b) {
    for (int i = 0; i < ARRAY_SIZE; i++) {
        a[i] = i;
        b[i] = ARRAY_SIZE - i;
    }
}
```

The compiled binary for the program containing this code snippet weighs in at 4096B. If this code was executed on the machine, what is the maximum value of `ARRAY_SIZE` that would allow this code to execute with 0 page faults in the best-case scenario? (Answer in IEC prefix: 8Gi, 32Ti, etc)

Each page is $2^{12} = 4096\text{B}$. Thus, the code fits in one page, so to have no faults, the data must be in the other. Each array is the same size, so $2 * \text{ARRAY_SIZE} \leq \text{PAGE_SIZE} \Rightarrow \text{max ARRAY_SIZE} = \underline{512}$

arr size is $\frac{4096\text{B}}{2} = 2048\text{B} \Rightarrow 512\text{ ints}$

c) How could we modify the above code snippet to allow a larger `ARRAY_SIZE` and execute with the fewest page faults in the best-case scenario? Write the new code below:

```
for (int i=0; i < ARRAY_SIZE; i++) {
    a[i] = i;
}
for (int i=0; i < ARRAY_SIZE; i++) {
    b[i] = ARRAY_SIZE - i;
}
```

We can make `ARRAY_SIZE = page size` because we no longer worry about a & b knocking each other, or the code, out of memory

F2: Why can't you use parallelism at a gas station? It might cause a spark. (10 points)

1. Optimize factorial() using SIMD intrinsic(AVX).

```
double factorial(int k) {
    int i;
    double f = 1.0;
    for (i = 1 ; i <= k ; i++) {
        f *= (double) i;
    }
    return f;
}
```

You might find the following intrinsics useful:

<code>__m256d __mm256_loadu_pd(double *s)</code>	returns vector(s[0], s[1], s[2], s[3])
<code>void __mm256_store_pd(double *s, __m256d v)</code>	stores p[i] = v _i where i = 0, 1, 2, 3
<code>__m256d __mm256_mul_pd(__m256d a, __m256d b)</code>	returns vector(a ₀ b ₀ , a ₁ b ₁ , a ₂ b ₂ , a ₃ b ₃)

```
double factorial(int k) {
    int i, j;
    double f_init[] = {1.0, 1.0, 1.0, 1.0};
    double f_res[4];
    double f = 1.0;
    // initialize f_vec
    __m256d f_vec = __mm256_loadu_pd(f)
    // vectorize factorial
    for (i = 1 ; i <= 4 * (k/4) ; i += 4) {
        double l[] = {
            (double) i, (double) i+1,
            (double) i+2, (double) i+3 };
        __m256d data = __mm256_loadu_pd(l);
        f_vec = __mm256_mul_pd(f_vec, data);
    }
    // reduce vector
    __mm256_store_pd(f_res, f_vec); → put result into f_res array
    for (j = 0 ; j < 4 ; j++) {
        f = f * f_res[j];
    }
    // handle tails
    for ( ; i <= k ; i++) {
        f *= (double) i;
    }
    return f;
}
```

Handwritten notes:

- initialize to all 1's (multiplicative identity)
- $4 * (k/4)$ is largest multiple of 4 $\leq k$
- reduce into f
- manually handle tail case

F2: (continued)

2. Cache Coherence:

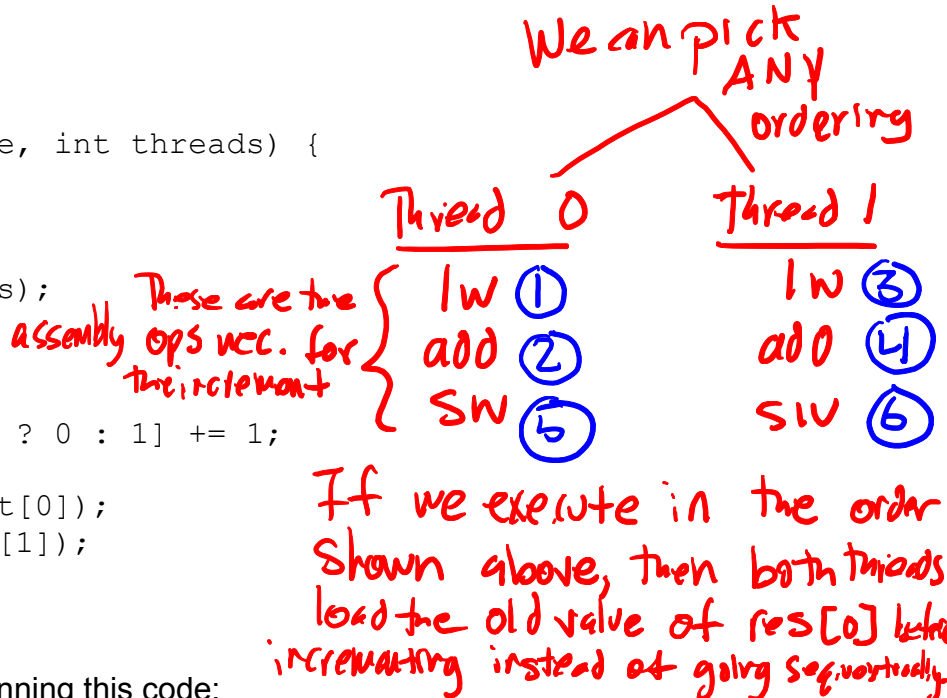
We are given the task of counting the number of even and odd numbers in an array, A, which only holds integers greater than 0. Using a single thread is too slow, so we have decided to parallelize it with the following code:

```
#include <stdio.h>
#include "omp.h"
void count_eo (int *A, int size, int threads) {
    int result[2] = {0, 0};
    int i, j;

    omp_set_num_threads(threads);

    #pragma omp parallel for
    for (j=0; j<size; j++)
        result[(A[j] % 2 == 0) ? 0 : 1] += 1;

    printf("Even: %d\n", result[0]);
    printf("Odd: %d\n", result[1]);
}
```



As we increase the number of threads running this code:

a) Will it print the correct values for Even and Odd? If not, explain the error.

Not always. Consider the case where we have 2 threads, and each has an index that corresponds to an even #. See the explanation above.

b) Can there be false sharing if the cache block size is 8 bytes?

Yes, the result array can be stored in a single block, so writing to any element will remove the block from all other caches.

c) What about 4 bytes?

No, false sharing is only applicable when the data being accessed by the threads are distinct. If the block is 4 bytes, then a single element of result is in a cache block, so we would invalidate blocks in other caches ONLY if they had the SAME element of the result array. Thus, false sharing is technically not possible.

F3: This isn't a bathroom. Why is there potpourri? (10 points)

1. T/F Questions (Circle one. If the circling is unclear, you will receive no credit.)

- 1) CPUs need separate instructions to access I/O devices. (True / **False**) *(consider memory mapped I/O)*
- 2) Segmentation (base + bound) has fragmentation problems. (**True** / False) *(contiguous, large blocks per process)*
- 3) Exceptions in early pipeline stages override exceptions in later stages for a given instruction. (**True** / False) *Exceptions in early stages can cause problems later on*
- 4) Exceptions are handled in the pipeline stage where they occur. (True / **False**) *Done by exception handler*

2. Polling, Interrupts, and DMA

1) Choose polling or interrupt for the following devices.

Device	Data Rate	Transfer Block Size	Polling/Interrupt?
A	80B/s	4B	Polling
B	400MB/s	4B	Interrupt
C	400MB/s	2KB	Interrupt

2) To support interrupts, the CPU should be able to save and restore the current state. Which of the following should be saved before handling interrupts to ensure correct execution?

- a.** Program Counter
 - b.** User Registers
 - c. TLB
 - d. Caches
- We need these 2 to know the state of our process. The others are just optimizations*

3) To which device in 1) is direct memory access (DMA) most beneficial? Explain briefly.

C. CPU can do work when transferring large blocks, and DMA enables this.

3. Warehouse Scale Computing and Amdahl's Law.

1) We are going to train convolutional neural networks on Amazon EC2. It turns out that 90% of training can be parallelized but the rest takes twice as long due to the communication overhead among the instances. What is the maximum speedup we can achieve?

10% is non parallelized, so this is doubled
90% is parallelized, sped up by 2

$$\Rightarrow \frac{1}{(0.1 \times 2) + \frac{0.9}{2}} = \frac{1}{0.2 + 0.45} \leq 5 \text{ (as } s \rightarrow \infty, \text{ expr} \rightarrow \frac{1}{0.2})$$

2) Which of the following can increase the maximum speedup in 1)?

- a. Use more instances
- b. Deploy the application across multiple arrays.
- c.** Reduce the number of reduce operations.
- d.** Increase network bandwidth.
- e. Increase the capacity of disks.

We need to improve the NON parallelizable part, which is the communication done. This communication is boosted by a better network OR if we do it less frequently. These options are d & c respectively.

F3: (continued)

4. Hamming Error-Correction Code (ECC)

1) Suppose we have a one-byte data value, 01101101_{two} . Fill in the encoded data in the following table.

Bit	1	2	3	4	5	6	7
Encoded Data							
P1	X		X		X		X
P2		X	X			X	X
P4				X	X	X	X

2) Assume that we have an encoded value, 1001110_{two} with a single-bit error. Indicate below each parity bit if it has an error:

Parity Bit	P1	P2	P4
OK/ERROR	OK	ERROR	ERROR

As both P2 & P4 have errors, there is an error in bit $2+4=6$ (should be 0)

Incorrect bit position: 6

Correct data: 0100

Fixed: 1001100 → 0100 is data

$P1: 1 \wedge 0 \wedge 1 \wedge 0 = 0$
 $P2: 0 \wedge 0 \wedge 1 \wedge 0 = 1$
 $P4: 1 \wedge 1 \wedge 1 \wedge 0 = 1$

} use table above

5. Dependability and RAID

1) Which of the following can increase the availability?

- a. Increasing MTTF
- b. Decreasing MTTF
- c. Increasing MTTR
- d. Decreasing MTTR
- e. Redundant data copies

We want to reduce failures OR decrease repair time. a & e follow the first reason, while d follows the second.

2) Explain very briefly why RAID1 is the most expensive form of RAID.

It requires a full copy of disks, so the overhead is 100%.

3) How many check disks are needed for RAID3?

One, RAID3 has a dedicated parity disk and uses byte-striping

4) Explain why RAID5 has a higher write throughput than RAID4.

The check information is distributed across disks, so not a single disk is queried for ALL checks. Thus, the load is balanced and the parity disk bottleneck is mitigated.