

# [Final] Past Exams - 2021 #1042



**Jero Wang** STAFF  
12 months ago in **Exam - Final**

1,112  
VIEWS



You can find the past exams here: <https://cs61c.org/fa23/resources/exams/>. Please check the linked past Piazza/Ed Q&A PDFs first before asking here. Many of the questions are already answered in those! Video walkthroughs (if available), are also linked on that page!

When posting questions, please reference the semester, exam, and question in this format so it's easier for students and staff to search for similar questions:

## Semester-Exam-Question Number

For example: **SP22-Final-Q1**, or **SU22-MT-Q3**



**Anonymous Armadillo** 12mth #1042ebc **Unresolved**

SP21-Final-Q6b

I understand the size of the struct foo is 8 bytes but why is the size of f 4 bytes



**Anonymous Armadillo** 12mth #1042eba **Unresolved**

SP21-Final-Q6a

if b and c swaps why does the size of the structure increase and what does it increase to



**Anonymous Whale** 12mth #1042ebd

Keeping in mind that we are following RISC-V alignment rules.

For part b, we swap b and c resulting in struct foo looking something like this:

```
struct foo{
```

```
char a;
```

```
char *c;
```

```
uint16_t b;
```

```
struct food *d;
```

```
}
```

char a is 1 byte so it would take up index 0. char \*c is a pointer so it is 4 bytes and according to alignment rules needs to start at multiple of 4 so it would take up indices 4, 5, 6, 7.

uint16\_t b is 2 bytes since it says 16 bits in variable definition and it would need to start at multiple of 2 so it takes up indices 8, 9. struct food \*d is a pointer so it is 4 bytes and it would

need to start at multiple of 4 so it would take up indices 12, 13, 14, 15 thus making size of struct 16 which is more than the 12 in the part a)i)

♡ ...



**Anonymous Reindeer** 12mth #1042eae

Unresolved

**SP21-Final-Q6b:**

Why is the offset for loadword in parentheses, isn't the register supposed to be in parens? I thought we do something like lw t0 4(t0) rather than lw t0 t0(4)

♡ ...



**Anonymous Boar** 12mth #1042eab

Resolved

**FA21-Final-Q1.8**

Where can we find the information about the number of parity blocks per data blocks like mentioned here (1/10 for RAID 5)? I don't remember the lecture or discussion ever getting this specific.

Update: do we just assume that 1 disk is always used for parity in RAID 4,5 and 2 disks for RAID 6?

♡ ...



**Anonymous Ant** 12mth #1042ead

same question! is this in scope?

♡ ...



**Jero Wang** STAFF 12mth #1042ecc

Yes, one parity block for RAID 4/5 and 2 for RAID 6. This is in scope.

♡ ...



**Anonymous Boar** 12mth #1042eee

Would it be 1 block or 1 disk, since each disk has multiple blocks?

♡ ...



**Anonymous Alligator** 12mth #1042dfe

Resolved

**FA21-Final-Q7.4**

Index	Tag 1	Valid 1	Tag 2	Valid 2
0b00	0b1011	1	0b1101	1
0b01	0b0011	1	0b0010	1
0b10	0b1110	1	0b0111	0
0b11	0b1111	0	0b0001	0

For 0b0011001101, I am confused that why this could be a compulsory miss. In index 0x00, both tag are valid.

♡ 1 ...



**Jero Wang** STAFF 12mth #1042ebf

Yes, both tags are valid. However, neither match the tag of the block we're trying to access, which is `0b0011`, so it has to be a miss.

♡ ...



Anonymous Meerkat 12mth #1042dfd

Unresolved

### SP21-MT-Q7a(iii)

I don't understand why the first instruction is `sll a0 s6 3` — why is it shifted by 3 specifically?

I also don't understand why the immediate is 4 for the `lw a0 4(a0)` instruction.

♡ ...



Ali Khani 12mth #1042dfc

Resolved

### SP21-Final-Q6Aii

I understand the byte alignment principle regarding the previous part of this question, but I'm not sure what the size of the new `struct` would be with `b` and `c` swapped. How many padding bytes would we need to add because the `uint16_t b` and `char *c` got swapped and each have a different size requiring separate offset intervals?

#### (a) (6.0 points) The Structure of Structures

- i. Assuming a 32-bit architecture with RISC-V alignment rules:

Consider the following structure definition and code:

```
struct foo {
    char a;
    uint16_t b;
    char *c;
    struct foo *d;
}
```

What is `sizeof(struct foo)` (Answer as an integer, with no units)?

12

- ii. If `b` and `c` are swapped, this increases the size of the structure:

True

♡ ...



Anonymous Tapir 12mth #1042dff

My understanding is that initially, we have `a` taking 1 byte, but since `b` has to be at an address of a multiple of 2, our layout is

0: `a` (1 byte)

2: `b` (2 bytes)

4: `c` (4 bytes)

8: `d` (4 bytes)

So we get 12 in total. If we swap `b` and `c`, we can't start `c` at 2 because it has to be at an address of a multiple of 4. We also can't start `d` at 10 for the same reason:

0: `a` (1 byte)

4: `c` (4 bytes)

8: b (2 bytes)

12: d (4 bytes)

So we get 16 bytes in total.

♡ ...



Anonymous Goldfish 12mth #1042ddc

✓ Resolved

### FA21-Final-Q9

Could Blank 2 also be  $n/4 * 4$ ?

General: How to know when to cast parameters when using SIMD instructions?

♡ ...



Jero Wang STAFF 12mth #1042ecd

Yes.

Pointers are implicitly converted for SIMD instructions, so you don't have to explicitly cast.

♡ ...



Anonymous Goldfish 12mth #1042ddb

✓ Resolved

### FA21-Final-Q8.1

What is the formula to find the number of virtual pages? What are the steps to get  $2^{32}/2^{12} = 2^{20}$ ?

Also what is the difference when asked for number of pages vs number of virtual pages?

♡ ...



Jero Wang STAFF 12mth #1042ece

$2^{32}$  because 4GiB of virtual memory (which equals virtual addresses of 32 bits), and  $2^{12}$  because page size is 4KiB (so offsets are 12 bits long).

If you're asking about the math itself, when dividing powers with the same base, you can subtract the exponents.

It should never ask for the number of "pages", it's always either virtual or physical. Is this referring to a specific question?

♡ ...



Anonymous Fox 12mth #1042dda

✓ Resolved

### Fa21-Final-Q6.3

Q6.3 (2.5 points) Par 3

An X is used to signify that either 1 or 0 can be outputted for the corresponding input.

W	Y	Z	Out
0	0	0	X
0	0	1	X
0	1	0	0
0	1	1	1
1	0	0	X
1	0	1	X
1	1	0	1
1	1	1	0

**Solution:** Staff solution:  $W \wedge Z$

Other answers may be possible (Notably, this question can yield a score well below par).

In this case, we note that there's never a time when  $Y$  distinguishes between 1 and 0; as such,  $Y$  never affects the result of our output, and we can look at the reduced truth table containing only  $W$  and  $Z$ .

Hi for this question, if there's an X in the output does that mean we simply ignore the row?

♡ ...

 **Nikhil Kandkur** STAFF 12mth #1042ddd

Yup! That means that the output for that specific combination of inputs does not matter.

♡ ...

 **Anonymous Goldfish** 12mth #1042dcf

✓ Resolved

**FA21-Final-Q7.4**

How do you identify given an address like in q7.4 whether there is a Hit, Compulsory Miss, Capacity Miss, or Conflict Miss? What happens if tag isn't in the index - is this compulsory or conflict? What happens when valid bit is off?

♡ ...

 **Jero Wang** STAFF 12mth #1042ecf

It could be either, since you don't know if its the first time that the cache has accessed this address.

If the valid bit is off, its the same as if the cache entry is empty.

♡ ...

 **Anonymous Manatee** 12mth #1042dce

Unresolved

**SP21-Final-Q1(b) ii**

Could someone explain how they got this answer? I'm not sure how to interpret the phrase "annualized failure rate" and how it relates to MTTF.

ii. Q2B

- A. Your company would like to restrict the annualized failure rate to be 1% for the individual machines in a large cluster. What does the Mean Time To Failure (MTTF) have to be to satisfy this annualized failure rate? Assume that the MTTF in this question is unrelated to that of part a. Write down your answer in years.


100

♡ 1 ...

 **Anonymous Tapir** 12mth #1042def

Here is my understanding: annualized failure rate (AFR) of 1% gives the probability that only 1 out of 100 devices in a cluster fails in 1 year. MTTF gives us the average number of years until a device fails. So then we would have  $AFR = 1/MTTF$ , so  $MTTF = 100$ . Intuitively, if devices fail once every 100 years, and we have 100 devices, 1 device will fail every year.

♡ ...

 **Anonymous Manatee** 12mth #1042dcd ✓ Resolved

SP21-Final-Q1(a)

Could someone explain how they got this answer?

(a) Q1

- i. You have a program that spends some percentage of its time waiting for requests and the rest of the time performing calculations. Suppose you have 8 threads which you can use to parallelize calculations, with no overhead or non-parallelizable calculations. What is the maximum fraction of time that your sequential program can spend on waiting for requests if we would like to achieve at least 4 times speedup? Leave your answer as a single simplified fraction.

$\frac{1}{7}$


♡ ...

 **Anonymous Dotterel** 12mth #1042ded 🔗 ENDORSED

Suppose that we spend a proportion of  $x$  time waiting. Then, the total time spent with the  $x8$  speedup would be  $x + \frac{1-x}{8}$ . Since we want a speedup of 4, we get

$$x + \frac{1-x}{8} = \frac{1}{4} \implies x = \frac{1}{7}$$

♡ ...

 **Anonymous Manatee** 12mth #1042dcc ✓ Resolved

FA21-Final-Q1.4

This slide explicitly says that we use polling when the data rate is low and interrupts when the data rate is high. However the solutions state the opposite. Which is correct?



# Polling, Interrupts and DMA

- **Low data rate (e.g. mouse, keyboard)**
  - Keyboard uses interrupts or polling (depending on the HW interface)
    - Overhead of interrupts ends up being low
  - Mouse input is typically polled
- **High data rate (e.g. network, disk)**
  - Start with interrupts...
    - If there is no data, you don't do anything!
  - Once data starts coming... Switch to Direct Memory Access (DMA)

Q1.4 (1 point) TRUE or FALSE: For high-performance network devices, polling tends to be used when there's a low data rate, while interrupts tend to be used when there's a high data rate.

TRUE

FALSE

**Solution:** False; generally, the opposite tends to happen.



N **Nikhil Kandkur** STAFF 12mth #1042dde

The slides list that either could be used for low data rate and that high data rate would start, but not solely use, interrupts. The solution provided in the exam is correct, and is the recommended approach: *interrupts for low data rate, and polling for high data rate.*



**Anonymous Manatee** 12mth #1042deb

So why does a mouse use polling then? Isn't a mouse an example of low data rate?



N **Nikhil Kandkur** STAFF 12mth #1042dec

↩ Replying to Anonymous Manatee

I'd say that's probably because one is more likely to move a mouse than type on a keyboard, so polling would be good then.






**Anonymous Armadillo** 12mth #1042dbf

Unresolved

sp21-final-q1D, why is option A included as a valid option? Due to data dependency won't parallel fibonacci lead to data race?






 **Anonymous Dotterel** 12mth #1042dee  
Notice that this is not an exact Fibonacci, and that we are actually incrementing by 3 each time. Therefore, there is no data race and splitting up iterations of the loop between threads will give a speedup.  
 ...

 **Anonymous Tiger** 12mth #1042dbc ✓ Resolved  
For fa21-final- 6.3 (and in general) - how do you approach identifying it's an xor.

I am using the approach where you sum up the products of all the true values, but like what result does it look like to know its an XOR. Thanks!!

 ...


 **Jero Wang** STAFF 12mth #1042eda  
You can expand an XOR -  $A \oplus B = (A \mid B) \ \& \ \sim(A \& B)$   
 ...

 **Anonymous Tiger** 12mth #1042dbb ✓ Resolved  
Fa21-final 1.9  
In the solution it implies that thread 1 reads  $y = 0$  and then sleeps.

therefore - wouldn't (after thread 2 runs setting  $y = 10$  and  $x = 0$ ) the while Loop for thread 2 return false when it reawakes. Since  $x = 0$ , it would skip any loops and return  $y = 0$ . Not  $y = 1$ .

A detailed answer is appreciated - thanks!

 2 ...

 **Andy Chen** STAFF 12mth #1042dfb  
The order of events for the worst case is as follows:

1. Thread 1 checks while loop condition --> true
2. Thread 1 reads  $y = 0$
3. Thread 2 starts and completes the loop ( $y$  will be 10 and  $x$  will be 0 at the end)
4. Thread 1 writes  $y = 1$
5. Thread 1 reads  $x = 0$ , then writes  $x = -1$
6. Thread 1 checks while loop condition --> false
7. Thread 1 exits the while loop

The while loop condition is only checked once every iteration, before starting to execute the loop code so 1 must be written to  $y$  when thread 1 wakes up. Hope this helps!

 ...

 **Anonymous Cassowary** 12mth #1042dba ✓ Resolved  
fa21-final-q4.5

could someone please explain the computation and how  $1/25\text{ps} = 40 \text{ GHz}$ ?



Q4.5 (2 points) What is the maximum allowable clock frequency for this circuit to function properly, in gigahertz?

GHz

**Solution:** 40 GHz

$$1 / (25 \text{ ps}) = 40 \text{ GHz}$$

♡ ...



N

Nikhil Kandkur STAFF 12mth #1042dea

$$\frac{1}{2.5 \times 10 \times 10^{-12} \text{seconds}} = \frac{1}{2.5} \times \frac{1}{10^{-11}} = \frac{4}{10} \times 10^{11} = 4 \times 10^{10} \text{Hz} = \frac{4 \times 10^{10} \text{Hz}}{1} \times \frac{1 \text{GHz}}{10^9 \text{Hz}} = 4 \times 10 \text{GHz} = 40 \text{GHz}$$

♡ 1 ...



Anonymous Narwhal 12mth #1042daf

Unresolved

SP21-Final-Q2I

Each PTE is 4 bytes, so the page table takes up  $220 \times 4 = 222$  bytes in total. When they say that each page table entry takes up 4 bytes how do they know that this is true? Also in the problem statement they say: After running one iteration of the inner loop for the code given in line, how many physical pages will our page table take up?

Would running one iteration of the inner loop change the number of physical pages our page table takes up in any other case because it doesn't seem to in this case.

♡ ...



Anonymous Koala 12mth #1042dae

Unresolved

Is there an alternative solution for 7a? I was thinking that we could modify the immediate generator and the ALU, so we can compare 0 with the register in the ALU. Nothing else needs to change really?

♡ ...



Anonymous Koala 12mth #1042daa

Unresolved

SP21-Final-Q1-c

What are parse trees?

♡ ...



Anonymous Lion 12mth #1042dca

+1

♡ ...



Anonymous Dogfish 12mth #1042cfa

Resolved

Can someone explain how partii is true? I thought that since pointers and uint16\_t types would both be 4 bytes so it shouldn't matter which one come first right?

## 6. C Structures

### (a) (6.0 points) The Structure of Structures

- i. Assuming a 32-bit architecture with RISC-V alignment rules:

Consider the following structure definition and code:

```
struct foo {
    char a;
    uint16_t b;
    char *c;
    struct foo *d;
}
```

What is `sizeof(struct foo)` (Answer as an integer, with no units)?

12

- ii. If `b` and `c` are swapped, this increases the size of the structure:

True

- True  
 False

♡ ...



**Anonymous Koala** 12mth #1042dad

uint16\_t is 2 bytes ( $8 \times 2 = 16$ ), hence with putting it next to char, it becomes  $1+2+(1 \text{ padding})$ . When we swap uint16\_t with char\*, now it becomes 4 (char now is 4 due to padding) + 4 + 4 (uint16\_t is also 4 due to padding) + 4 = 16, which is greater than 12.

♡ ...



**Anonymous Narwhal** 12mth #1042cef

Unresolved

### Sp21-Final-Q2g:

This question says that because entire array can fit in memory, and Because the access pattern is effectively the same, all accesses are hits.

How did they know that more than 4 blocks that we need access to would not have the same set index. If more than 4 blocks belonged to a set wouldn't that be a problem because we have a 4 way set associative cache?

♡ ...



**Anonymous Koala** 12mth #1042dab

Also, I thought compulsory hits exist?

♡ ...



**Anonymous Louse** 12mth #1042cee

Resolved

### FA21-Final-Q3.3

Why do we have to add 24 bytes to avoid overlapping the injected instructions? Wouldn't the `addi sp, sp, -24` from `verifypassword` already have been called before we inject these five instructions and run them (so the extra space was already allocated)?

Q3.3 (4 points) Regardless of your previous answer, assume that we can put up to 5 instructions in the buffer. Unfortunately, that's not really enough instructions to do much. Instead, we decide to inject code that lets us run longer programs, instead of only being limited to 5 instructions.

Complete the following 5-line code which does the above. You may use pseudoinstructions, as long as they resolve to exactly one instruction. Each blank is worth 1 point.

```
# Allocate a buffer of 256 bytes, which does not overlap with any data
# we already are using (such as the instructions injected in part 1)

1: addi _____

# Set the argument of Get20Chars to the start of the allocated buffer

2: mv _____ sp

# Set t1 so the next instruction jumps to Get20Chars

3: lui _____

4: jalr ra t1 -256 # Call Get20Chars

# Jump to the start of the buffer

5: _____
```

**Solution:** Line 1: `addi sp sp -280`. We allocate 256 bytes of buffer, plus 24 bytes to avoid overlap with the injected instructions and the saved ra register from part 1, for a total of 280 bytes. (This does mean that we rely on data after our stack pointer to stay constant, which is technically undefined behavior. In reality, the data will stay consistent for long enough that we move the stack pointer back down; since our sp starts in the middle of a page, the data will not be deallocated. The reasoning in the last sentence is not considered in scope, but is relevant in CS 162.)

Line 2: `mv a0 sp`. The `Get20Chars` function expects the address of a buffer as an argument in the a0 register.

Line 3: `lui t1 1`. This puts the value `0x1000` in `t1`, which then causes the next line to jump to `0x1000 - 256 = 0x0F00`, the address of `Get20Chars`. Note that we needed to use `lui` here, since `addi` can only increase by up to `0x7FF`. Line 4 could also not have been a `jal` operation, since our jump distance would have exceeded the size of a 20-bit immediate.


Line 5: `jr sp`. We want to jump to the address in the `sp` register, and we don't care about saving a return address anywhere.

```

1 verifypassword:
2     addi sp, sp, -24 # Space for:
3     sw ra 20(sp)    # ra
4     mv a0 sp        # 20-byte buffer
5     jal ra Get20chars
6     la t0 Password
7     mv t1 sp
8 Loop:
9     lb t2 0(t0)
10    lb t3 0(t1)
11    bne t2 t3 Fail
12    beq t2 x0 Pass
13    addi t0 t0 1
14    addi t1 t1 1
15    j Loop
16 Pass:
17    addi a0 x0 1
18    j End
19 Fail:
20    mv a0 x0
21 End:
22    lw ra 20(sp)
23    addi sp sp 24
24    jr ra

```

♡ 1 ...

 **Anonymous Fish** 12mth #1042cfe  
^same question

♡ 1 ...

 **Anonymous Camel** 12mth #1042dcb  
^ same

♡ ...

 **Anonymous Dunlin** 12mth #1042dfa

To get to the code that we inject into the buffer, we run verifypassword to completion, then load the injected ra into the ra register, and then add 24 to sp, then finally jump to ra, which is where the 20 bytes of injected code are. Thus, by the time we get to the injected code, the stack pointer has already been reset. To avoid overwriting the data we are already using (as stated in the comments for Q3.3) we must decrement by 24 for the injected code + ra + 256 for the buffer we want, for a total of 280.

♡ ...

 **Anonymous Louise** 12mth #1042ced ✓ Resolved

FA21-Final-Q5.3 in the first sentence of the solution, what is meant by "(since we don't change operation based off ID)"? Is it saying the instruction decode won't choose a different operation based on the values stored in the register?

Also, if we didn't initialize the registers, are we ok with using garbage values?

Q5.3 (2.5 points) Write a sequence of instructions that causes a hazard in a completely unoptimized 5-stage pipeline (no forwarding, no branch prediction, no synchronous read/writes, etc), but which would not cause a hazard if all mac instructions were changed to mul instructions. If no such sequence exists, write "Not Possible."

**Solution:** Note that regardless of the values stored in registers, we still need to stall (since we don't change operation based off ID), so we can write code without initializing registers.

The extra hazard that occurs as a result of this instruction is a data hazard on rd; with mul, we don't need to wait for the value of rd to get written back, but we do need to wait for rd for a mac.

Thus, a correct answer required a mac instruction up to 3 instructions after its rd got updated, and no other hazards involving other registers. Our staff solution was:

```
mac a0 a1 a1
```

```
mac a0 a1 a1
```



**Jero Wang** STAFF 12mth #1042edb

I think it's referring to the fact that the actual data we read from the regfile (in ID) doesn't affect whether we decide to stall or not.

I don't think you need to initialize the registers here.



**Anonymous Narwhal** 12mth #1042cde

Unresolved

### SP21-Final-Q1.2:

I am a little bit confused on the explanation that is given for this question. The explanation says:

Every 4000 hours, there is 1 failure. That failure takes 3 hours to repair. Therefore, the system is up for 3997 hours, for an availability of 3997/4000.

However in the question it says that the mean time between failure is 4000 which means no failure occurs in those 4000 hrs. However the answer says that the failure occurs within the 4000 hrs and that the system runs for 3997 minutes and is getting repaired for 3hrs within the 4000hrs. I might be understanding the problem incorrectly. Please let me know how to interpret this question.



**Anonymous Chimpanzee** 12mth #1042cdd

Resolved

### sp21 q2 b (ii)

How do we solve this question? There is no walkthrough or solution given..

Given the two-level page table we have in Figure 1, translate the following virtual addresses to physical addresses. Assume all page tables shown are active. If an address does not point to an active page table at any point, write "Invalid Address".

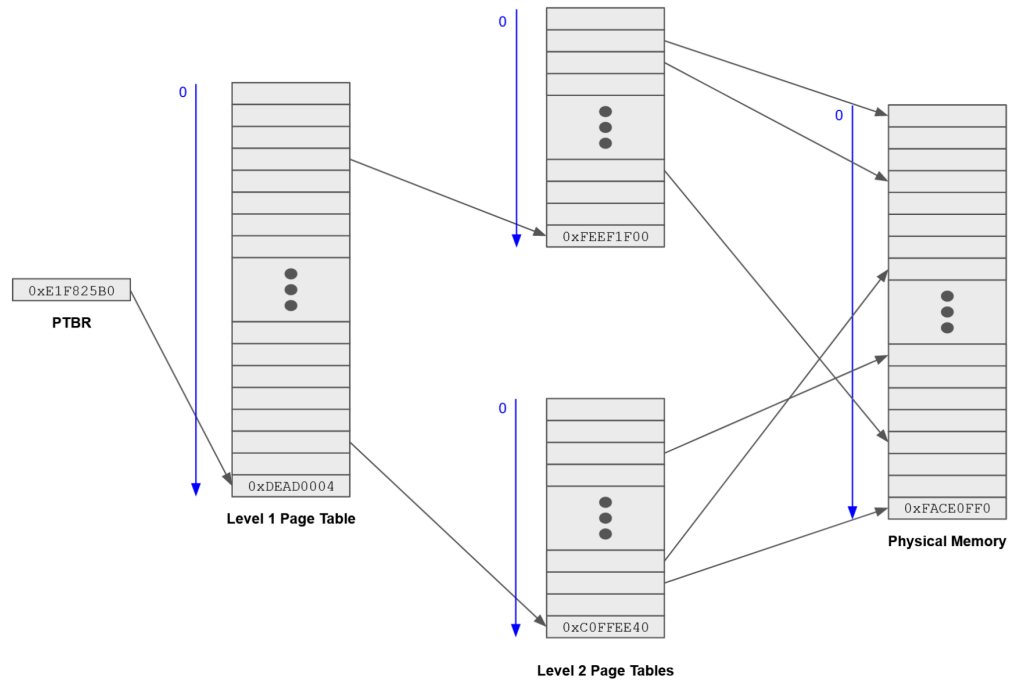


Figure 2

ii. A. Address: 0xB7C90AEA



Nikhil Kandkur STAFF 12mth #1042cfb

This is multilevel page tables, which I believe is out of scope for this semester.



Anonymous Cassowary 12mth #1042cdc

Unresolved

SP21-Final-Q10

iv. 1: endofstring:

Inputs: a0 is a pointer to a **nonempty** string

Output: a0 returns a pointer immediately after the end of the string.

Example: Let a0 be the pointer 0x10000000 to string s, which is the string "Hello". Then the expected output would be 0x10000006.

endofstring:

\_\_\_t0\_\_\_\_\_ #Code line 4. You must use t0 as the first register of this instruction.

\_\_\_\_\_ #Code line 5

\_\_\_\_\_ #Code line 6

ret

You may **not** use strlen or any other library function

Code line 4:

lb t0 0(a0)

v. Code line 5:

addi a0 a0 1

vi. Code line 6:

bneq t0 x0 endofstring OR bneq x0 t0 endofstring OR bnez t0 endofstring

I'm confused how this would be correct. if we use the 'hello' and we start at 0x10000000, we will get to t0 = 'o' and a0 = 0x10000006 but then bneq t0 x0 would take because t0 = 'o' and then we

go back to endofstring which will then make  $t_0=x_0$  and make  $a_0 = 0x10000007$  and then it would return the wrong thing in  $a_0$

♡ ...



Anonymous Crab 12mth #1042cdb

✓ Resolved

SP21-Final-Q2.a.i.B

**Assume all systems are 32-bit.**

- i. We're given a system with an 4-way set associative cache of size 512 KiB with 128 B blocks. How many bits are allocated to the tag, index, and offset bits respectively?

A. Tag:

15

Offset:  $32 - 10 - 7 = 15$  bits

B. Index:

10

Index:  $(2^{19}) / (4 * 128) = 10$  bits

C. Offset:

7

Blocks are  $2^7$  bits large.

Why are there 10 index bits? My understanding was that in N-way set associative caches, the index points to which set you search, so in this case we'd have 2 index bits to refer to the 4 sets in this cache?

♡ ...



Anonymous Ant 12mth #1042cea

Same question. Did they teach diff that sem?

♡ ...



Anonymous Ant 12mth #1042ceb **✳ ENDORSED**

Ahh I think i get it. Heres my understanding:

N-way set associative means that we group our blocks in sets of size N. In this case that means we group into 4 block sets.

We know that each block is  $2^7$  B. So we know each set is  $4 * 2^7$  B =  $2^9$  B (because each set will have 4 of our  $2^7$  B blocks).

We also know that our entire Cache is  $2^{19}$  B. So we can find the number of total sets:  $2^{19} / 2^9, = 2^{10}$ . This means that there are  $2^{10}$  or 1024 different sets, each with 4 blocks of size  $2^7$  B.

Now when we get some address. We look at the index to tell us which set the new block should belong to. If there is a block in that set with the same tag, it is a hit (given a 1 valid bit). Else, if it is a miss we use LRU on that set to kick-out a block if needed.

Hopefully that helps!

♡ 1 ...



Anonymous Ferret 12mth #1042ccd

✓ Resolved

was not able to get this as the solution, how did they get this?

I got  $\sim Z + (\sim Y * Z)$

Q6.2 (2.5 points) Par 2

W	Y	Z	Out
0	0	0	1
0	0	1	1
0	1	0	1
0	1	1	0
1	0	0	1
1	0	1	1
1	1	0	1
1	1	1	0

**Solution:** Staff solution:  $\sim (Y \& Z)$

Other answers may be possible.

Similar to the previous, we note that the answer does not depend on  $W$ .

♡ ...



Anonymous Fox 12mth #1042ccf

🔒 ENDORSED

Your answer can be simplified further. One of the laws of Boolean algebra says:

$$X + YZ = (X+Y)(X+Z), \text{ or in our notation, } X + Y\&Z = (X+Y)\&(X+Z)$$

Plugging in  $\sim Z$  for  $X$ ,  $\sim Y$  for  $Y$ , and  $Z$  for  $Z$ , and you get

$$\sim Z + \sim Y\&Z = (\sim Z + \sim Y)\&(\sim Z + Z) = \sim Z + \sim Y = \sim(Z\&Y)$$

It is just easier to note, though, that the only times the table is not 1 are when both  $Y$  and  $Z$  are 1.

♡ ...





Anonymous Reindeer 12mth #1042cbd

Unresolved

Sp21-final 2F

How do we do we know that nothing is evicted while filling up the cache?

When I was solving the problem I thought: we need to get the memory address, find their TIO, and check their index, and put the corresponding data in the cache at the corresponding index. If we do this, then can't there be a write through or write back if there are more than 4 elements in the index of the cache? How can we guarantee that there is no cache eviction?



Anonymous Walrus 12mth #1042bfd

Resolved

FA21-Final-Q1.10

Q1.10 (1 point) Justin purchased his HP Pavilion 15t-cs300 laptop 1,000 days ago. During this time, it has broken twice, and had to be repaired. Each repair took 10 days to complete, during which time the laptop was unusable. What is the mean time to failure (MTTF) of Justin's laptop, in days?

days

**Solution:** 490 days. 1,000 days minus 20 broken days = 980 days, divided by 2 failures.

Q1.11 (1 point) What is the availability of Justin's laptop?

**Solution:** 0.98. 980 days out of 1,000 days.

Author's note: Apparently I made a mistake when writing this question; I've only had my laptop for 2 years, not 3 years. Other than that, this is a fairly accurate representation of my laptop. Not that I'm bitter or anything...

If im not wrong, isn't availability =  $MTTF / (MTTF + MTTR) \%$ , then if availability is 980/1000, how is the MTTF in the q1.10, 490?



Anonymous Whale 12mth #1042bff


ENDORSED

I'm not sure if I completely understood your question but this was my thought process:


MTTF is (total time - unusable time) / #failures which is  $(1000 - (10 * 2)) / 2 = 490$

and availability would then be  $490 / (490 + 10)$  since 10 is the mean time to repair and 490 is mean time to failure, this simplifies to  $490/500 = .98 = 98\%$ .

♡ 1 ...

 **Anonymous Walrus** 12mth #1042caa  
oh yea makes sense thank you!

♡ ...

 **Anonymous Goldfinch** 12mth #1042adb ✓ Resolved

Q1.12 (1 point) We've devised an error-correcting code which is able to fix 1 bit errors. If 0x61C is a valid codeword, which of the following can NOT be a valid codeword, regardless of the error-correcting scheme we have? Select all that apply.

- |                                |  |
|--------------------------------|--|
| <input type="checkbox"/> 0x71C | <input type="checkbox"/> 0xC16             |
| <input type="checkbox"/> 0x51C | <input type="checkbox"/> 0x16C             |
| <input type="checkbox"/> 0x70D | <input type="checkbox"/> None of the above |

**Solution:** 0x71C and 0x51C cannot be valid codewords.

For 0x71C: If we received the data 0x71C, we would not be able to know if it was a correct codeword, or if we were supposed to receive 0x61C, but a bit got corrupted.

For 0x51C: If we received the data 0x71C, we would not be able to tell if the original word was 0x61C or 0x51C.

All others: The bit distance from 0x61C is far enough that we don't run into the problems above. We can construct an error-correcting code with those as codewords by defining our error-correcting code to have only those two as valid codewords; this is able to fix 1-bit errors, among other things.


## \*\*Fa21-Final -Q1.12\*\*

could anyone express that more clearly or in detail?

♡ ...

 **Justin Yokota** STAFF 12mth #1042bbd  
[#1042abc](#)

♡ ...

 **Anonymous Panther** 12mth #1042abf ✓ Resolved


If BrUn is 0 does it mean signed or unsigned?

♡ ...

 **Sonika Vuyyuru** STAFF 12mth #1042eaf

The signal BrUn = 0 indicates signed. So for example, blt instruction would have BrUn = 0, since it is a signed branch instruction.

♡ ...

 **Anonymous Stork** 12mth #1042aba ✓ Resolved

**FA21-Final-Q2.4**

Q2.4 (3.5 points)  $2 + 2 + 2 + 2 + \dots$

**Solution:** We don't need to worry about rounding until we get to the first nonrepresentable even number. This occurs at  $0b1.00001$ , with enough exponent that the right 1 becomes the 2s place. This is  $0b1000010 = 66$ . This will round down to 64, so we effectively get "stuck" there, and continuously get 64 from then. Thus, our answer is 64.

I'm also a bit confused about this question. How come the first nonrepresentable even number is at  $0b1.00001$ ? I thought that we had at most 4 mantissa bits (so I thought our answer was 32, not 64).

♡ ...

 **Justin Yokota** STAFF 12mth #1042bbe

Yes. If we have 4 mantissa bits, then  $1.0001$  just barely fits within our 4 bits (and thus is representable).

♡ ...

 **Anonymous Stork** 12mth #1042cbe

Oh sorry I think I worded my question incorrectly. I was wondering why we wouldn't consider  $1.00001 * 2^5 = 0b100001$  as our first nonrepresentable number?

♡ ...

 **Justin Yokota** STAFF 12mth #1042ccb

↩ Replying to Anonymous Stork

That number is nonrepresentable, but it's also the number 33; we never actually hit 33, since we're adding by 2s.

♡ 1 ...

 **Anonymous Panther** 12mth #1042cce

↩ Replying to Justin Yokota

How do we determine that  $1.00001 * 2^6$  is the highest?

♡ ...

 **Anonymous Boar** 12mth #1042cff

↩ Replying to Anonymous Panther

Isn't the largest exponent value 30, so the largest even value we should be able to represent is  $1.0000 * (2^{15})$ ?

♡ ...

 **Jero Wang** STAFF 12mth #1042edc

↩ Replying to Anonymous Boar

Yes, but that doesn't guarantee that every even number less than this number is representable. For this subpart, it gets stuck at 64 and never goes higher.

♡ ...

 **Jero Wang** STAFF 12mth #1042edd

↩ Replying to Anonymous Panther

It's the highest because this would get rounded down to  $1.0000 * 2^6$ , due to the fact that we only have 4 mantissa bits. If we try to add 2 again, we get  $1.00001 * 2^6$ , which gets rounded down to  $1.0000 * 2^6$  and so on....

♡ ...



Anonymous Panther 12mth #1042aaf

Unresolved

SP21 - Final Q7

Can someone explain the answers to this question for part B and C

**B.** What changes would you need to make in order for the instruction to be able to execute correctly? Assume all modifications and additions are done on top of the existing single cycle datapath. Select all that apply

- Modify the control signals to the ALU.
- Modify Branch Comparator logic.
- Modify the control logic for WBSel.
- Modify the control logic for parsing `instr[31:0]`.
- Modify control logic for ALU/ALUSel.
- Add an additional comparator.
- Add additional control signals for the writeback mux.
- Modify ALU buses.
- Modify the control logic for the Branch Comparator.
- None of the above.

**C.** `is_null rd, rs1` is not in a standard RISC-V instruction format; as we're attempting to reduce the number of hardware changes in our datapath. We instead choose to implement our instruction as a pseudoinstruction in the following format. Which of the following statements is true? Assume earlier changes propagate. Select all that apply.

**Format: R-Type Instruction**

- We need to wire `x0` as `rs2` and modify the control signals.
- We need to provide a second argument `x0` when calling the instruction and modify the control signals.
- We need to provide a second argument `x0` as a comparator for all branch comparisons.
- We need to wire `x0` as a comparator for all branch comparisons.
- It is impossible to represent as an R-Type instruction.

♡ 1 ...



Anonymous Pigeon 12mth #1042abe

I am also confused on part c. Can someone explain this choice?

♡ ...



Anonymous Pigeon 12mth #1042aca

specifically, what is the difference between the first two options:

- wiring `x0` as `rs2` and modifying the control signals
- providing a second argument `x0` when calling the instruction and modifying the control signals.

♡ ...



Anonymous Dogfish 12mth #1042cfd

↩ Replying to Anonymous Pigeon

They're asking what needs to be true for `is_null rd rs1` to be a valid pseudo-instruction (it is worded pretty weirdly). As there isn't a second argument passed in you can't say

that you need to since the question is asking what needs to be true for `is_null` to be valid. And the only answer in that case is 1.

I have no idea about part b though so if someone could explain that I'd appreciate it.

♡ ...



Anonymous Stork 12mth #1042aae ✓ Resolved

### FA21-Final-Q2.3

Could someone walk me through how to solve this? I'm not entirely sure how they got the mantissa bits here and concluded that the sum converges to 2.

Q2.3 (3.5 points)  $1 + (1/2) + (1/4) + \dots$

**Solution:** We can look at how our mantissa changes:

0b1.0000

0b1.1000

0b1.1110

0b1.1111

0b1.11111 -> 0b10.000

0b10.000001 -> 0b10.000

Our answer is thus 2, which is mathematically correct.

♡ ...



Anonymous Armadillo 12mth #1042afb

I have the same question. Specifically, how we go from

0b1.11111 -> 0b10.000... Isn't 0b1.11111 halfway between 0b1.11110 and 10.000 so we round down to get 0b1.1111?

♡ ...



Justin Yokota STAFF 12mth #1042bbf

Per the question, floating point numbers round to the "most even" number (formally, the number divisible by the most factors of two). 0b1.1111 is less even than 0b10.000, so we prefer rounding to 10.000

♡ 1 ...



Anonymous Stork 12mth #1042cbf

↩ Replying to Justin Yokota

Sorry, I'm still a little confused as to how they calculated the mantissa changing. Shouldn't the third line be 1.1100?

♡ ...



Justin Yokota STAFF 12mth #1042cca

↩ Replying to Anonymous Stork

Yes, there are two lines omitted between the second and third line.

♡ 1 ...



Anonymous Stork 12mth #1042aad ✓ Resolved

## FA21-Final-Q9

For the last line of code for this question, I'm still confused as to why `output[0] + output[1] + output[2] + output[3] == 0` would be incorrect. The explanation from the solutions didn't make a whole lot of sense to me:

`output[0]+output[1]+output[2]+output[3] == 0` is incorrect, since we could have received outputs that happened to add to 0, even if they aren't all 0. As an example consider the inputs  $A = [0, 0, 0, 0]$ ,  $B = [0, 0, 0, 0]$ ,  $C = [1 \ll 30, 1 \ll 30, 1 \ll 30, 1 \ll 30]$ .

How does the example result in nonzero outputs that happen to add to 0?

♡ 1 ...



**Justin Yokota** STAFF 12mth #1042abb

Apologies; those should be  $1 \ll 62$ s instead of  $1 \ll 30$ s. Keep in mind that all numbers are 64-bit integers. After running the loop, we end up with `output[0]==output[1]==output[2]==output[3]== $2^{62}$` .  $2^{62} * 4 = 2^{64}$ , which overflows to 0.

♡ 1 ...



**Anonymous Moose** 12mth #1042aac

✓ Resolved

## FA21-Final-Q3.6

The explanation and the answer don't match up here. What is the case for the first option `jalr ra, t2, -256`? Which are valid?

Q3.6 (2 points) Which of the following jump instructions can we use in our injected code? Don't worry about these lines not properly calling `Get20Chars`; we just want a valid RISC-V jump without running into the problem identified in part 5 (Hint: use the conversion you already did in part 4).

Note that `+3840 == 0x0000 0F00`.

- `jalr ra, t2, -256`
- `jalr ra, t0, 16`
- `ret`
- `jalr s0, x0, 3840`
- `jalr x0, t1, -256`
- `jalr s0, t1, -256`
- None of the above

### Solution:

`ret = jr ra = jalr x0, ra, 0` and `jalr x0, t1, -256` contain a null byte when translated to machine code: `0x00008067` and `0xF0030067`. Note that we can reuse our translation from part 4 and note only the places that get changed.

`jalr s0, x0, 3840` is an invalid instruction because the immediate 3840 is greater than 2047. I-type immediates must be between -2048 and +2047.

`jalr s0, t1, -256` and `jalr ra, t0, 16` are valid, and don't have a null byte; note that `jalr s0, t1, -256 = 0xF0030467` doesn't have a null byte, because it gets split up into bytes as `0x67 0x04 0x03 0xF0`.

♡ 3 ...

Jero Wang STAFF 12mth #1042ebe

The first answer choice translates into `0xf00380e7`, which is valid as it's a valid instruction and doesn't have a null byte. We've noted the error in this solution on our end as well.

♡ ...


Anonymous Moose 12mth #1042fe

✓ Resolved

### FA21-Final-Q4

In the following question, they don't consider clk-to-q and setup time when calculating the combinatorial logic path that takes the longest time.

Do combinatorial path delays only consider the time for AND, OR, NOT, etc gates?



Q4.1 (2 points) What is the largest combinational delay of all paths in this circuit, in picoseconds?

ps

**Solution:** 16 ps

The longest path between two registers goes through the AND gate, then the OR gate, for a total delay of  $8+8=16$  ps.

♡ ...

Jero Wang STAFF 12mth #1042ede

CL delay does not include state elements, such as registers.

♡ ...



Anonymous Peafowl 12mth #1042fd

✓ Resolved

### SP21-Final-Q2I

How did they get the answer 1024?

**H. NOTE: for all parts, assume changes propagate unless otherwise stated.**

As we're working on running the code snippet, we realise we want to run different instances of the same code. We choose to employ virtual memory on our memory space. We have 4 GiB of virtual memory and 16 MiB of physical memory mapped with a single level page table with a page size of 4 KiB. We choose to store 8 bits of metadata with each page table entry.

**I.** After running one iteration of the inner loop for the code given in line, how many physical pages will our page table take up?

1024

NOTE: because we did not specify alignment for the system, if you solved for a byte-aligned system, you will get the points for 768 pages.

♡ 5 ...



Anonymous Reindeer 12mth #1042eac

I don't understand this either

♡ ...



Anonymous Cat 12mth #1042ebb

me neither

♡ ...



Andrew Liu STAFF 12mth #1042eed

A PTE consists of  $VPN \mid PPN \mid \text{Metadata}$ , where  $\mid$  denotes concatenation. You use this to calculate the size of all of PTE needed (one per VPN), and then from that size, you calculate the number of pages needed.

♡ ...



Anonymous Butterfly 12mth #1042fc

✓ Resolved

### Final Fall 2021 Question 6.3:



Q6.3 (2.5 points) Par 3

An X is used to signify that either 1 or 0 can be outputted for the corresponding input.

W	Y	Z	Out
0	0	0	X
0	0	1	X
0	1	0	0
0	1	1	1
1	0	0	X
1	0	1	X
1	1	0	1
1	1	1	0

**Solution:** Staff solution:  $W \wedge Z$

Other answers may be possible (Notably, this question can yield a score well below par).

In this case, we note that there's never a time when  $Y$  distinguishes between 1 and 0; as such,  $Y$  never affects the result of our output, and we can look at the reduced truth table containing only  $W$  and  $Z$ .

Would it be okay to say  $Y(W \wedge Z)$  as the answer or should the  $Y$  not be in the answer at all. I thought  $Y$  has to be there because the only cases where the output is 1 is when  $Y$  is 1 as well.

♡ ...



**Justin Yokota** STAFF 12mth #1042acc

That would be a valid, but less-efficient answer. You can choose whatever value you want for the Xs, and as it turns out, the most efficient solution requires assigning some Xs 1s (when  $Y=0$ )

♡ ...



**Anonymous Chimpanzee** 12mth #1042fb

Unresolved

**Sp 21 MT Q5 c iv**

Could someone explain why the least number of NO OPs required would be 10.

(c) (1.5 points) **Pipelining**

Assume we've added pipeline registers to create a 6-stage pipeline with our updated datapath, as seen below (Figure 3). This pipelined datapath acts similarly to our standard RISC-V 5-stage pipeline with the additional change of the memory section being separated into two sections. Assume no forwarding logic has been implemented, no branch prediction, and one register operation per cycle.

For the given code, what hazards might exist and due to which lines? Assume all registers have been initialised and that all labels are defined and that all branches are taken. **HINT: having a table open will be useful for scratch work.**

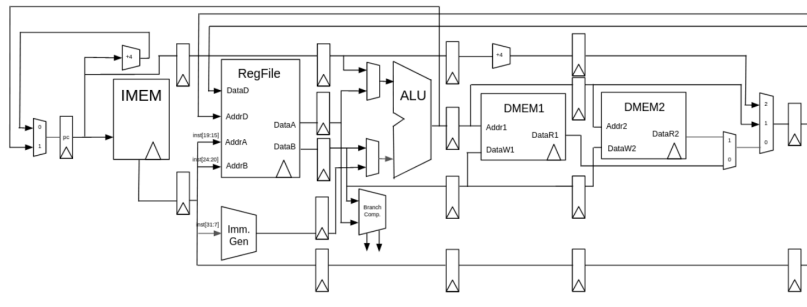


Figure 3

```

1 bne x0, t0, next
2 addi t1, t0, 1
3 lb s0, 0(t1)
4 shw s0, 4(t0)
    
```

Here's my logic:

St stands for stall

1. IF ID EX MEM WB
2. St IF ID. EX. MEM WB
3. St St. IF ID EX. MEM. WB
4. St. St. St. St. IF. ID. EX. MEM. WB

Correct me if I am wrong: there would be 1 stall in line 2 since the execute stage determines what gets added to the ALU (PC counter).

There would be 2 stalls in line 3 since we get the value of t1 in the execute stage of line 2.

there would be 4 stalls in line 4 since we load the value into s0 in the mem stage of line 3.

This only adds up to 7 stalls. How is the answer 10

CommentEditDelete

♡ 1 ...

Anonymous Panther 12mth #1042ada

Im confused about this question too. Is it because there's 2 mem stages now?

♡ ...

Anonymous Butterfly 12mth #1042eb

✓ Resolved

Q1.9 (1 point) We run the following code on two threads.

```
1 int y = 0;
2 int x = 10;
3 #pragma omp parallel
4 {
5     while(x > 0)
6     {
7         y = y + 1;
8         x = x - 1;
9     }
10 }
```

What is the smallest possible value y can contain after this runs?

**Solution:** This one's a bit tricky. The optimal sequence is:

Thread 1 reads  $y=0$  and goes to sleep.

Thread 2 runs to completion.

Thread 1 wakes up and writes  $y=1$ , reads  $x=0$ , sets  $x=-1$ , then sees  $x == -1$  and stops the loop.

**Fall 2021 Question 1.9:** I am still a little confused on how to get the answer for this one. Could someone explain the logic? If y is a shared variable, shouldn't the y value from Thread 2 be the same as Thread 1?

♡ ...

J Jero Wang STAFF 12mth #1042ec  
Please label your questions!

♡ ...

Anonymous Butterfly 12mth #1042fa  
Sorry about that! Just edited it.

♡ ...

J Justin Yokota STAFF 12mth #1042ff  
↩ Replying to Anonymous Butterfly

Yes, y should be. But in order to do math on y, you need to save it to a register. By interrupting the calculation, we can end up with the thread thinking y is some value when it's been updated to a different value already.

♡ ...

Anonymous Yak 12mth #1042ea ✓ Resolved

Fa21-Final-Q8(and VM as a whole):

I'm confused what happens once we reach page number 256? Since we have  $2^{20}$  pages, but only two hex numbers to use to assign the pages to in physical memory.

♡ ...

J Jero Wang STAFF 12mth #1042ed  
Please label your questions!

♡ ...

J Jero Wang STAFF 12mth #1042edf

You only have 256 physical pages,  $2^{20}$  pages is referring to virtual pages.

When you reach 256, you're out of physical pages.



T **Theo Putterman** 12mth #1042df **Unresolved**

Fa 2021 Final 7.4.2.


Couldn't this be a capacity miss? I thought a capacity miss was a miss that occurred in the fully associative cache but not the infinite cache. If we accessed 2 items at index 0, 1 item at index 2, 1 item at index 3, and then 1 billion items at index 01, a fully associative model would still have a miss, but an infinite model would not. As such, couldn't this be a capacity miss?



J **Jero Wang** STAFF 12mth #1042ee

Please label your questions!




 **Anonymous Pigeon** 12mth #1042de **Unresolved**

**SP21-Final-Q2E**

Does the second set of nested for-loops assume we leave off with the cache from the first set or is the question asking to independently assess this part and start with a new cache?




 **Anonymous Seahorse** 12mth #1042dc **Resolved**

Sp 21 Q 4

When calculating minimum clock periods in general, do we have to also take in consideration time from input to register? I always thought it was just register to register

#### 4. SDS



#### State Machine

Assume input A and input B come from registers. Please include ns in your answer.

(a) Assume all 2-input logical gates have a 10 ns propagation delay. The NOT gate has a 5 ns delay. All registers have a clk-to-q of 15 ns and setup time of 20 ns.

i. Find the minimum clock period to ensure the validity of the circuit.

75 ns

We have the following paths:

- Input A (clock-to-q) -> NOT -> Register (setup) = 15 ns + 5 ns + 20 ns = 40 ns
- Input A (clock-to-q) -> NOT -> AND -> NOR -> AND -> Register (setup) = 15 ns + 5 ns + 10 ns + 10 ns + 10 ns + 20 ns = 70 ns
- Input B (clock-to-q) -> NAND -> AND -> NOR -> AND -> Register (setup) = 15 ns + 10 ns + 10 ns + 10 ns + 10 ns + 20 ns = 75 ns
- Register (clock-to-q) -> NOT -> NOR -> AND -> Register (setup) = 15 ns + 5 ns + 10 ns + 10 ns + 20 ns = 60 ns

So we need the max of them which would be 75 ns.

♡ ...

⌋ C Catherine Van Keuren STAFF 12mth #1042dd

I believe since in the question it indicates to "Assume input A and input B come from registers," you would take them into account in this case.

♡ ...

⌋ Anonymous Opossum 12mth #1042db ✓ Resolved

#### FA21-Final-Q1.12

Can anyone explain the correcting logic? I still don't understand why 0x71C and 0x51C cannot be valid codewords after reading the answer.

♡ 1 ...

⌋ Justin Yokota STAFF 12mth #1042abc

Both 0x71C and 0x51C are "too close" to 0x61C; there's a one-bit difference between 0x71C and 0x61C, and a two-bit difference between 0x51C and 0x61C. The following paragraphs describe a specific input that becomes ambiguous as a result of two valid codewords being too close to each other.

♡ ...

⌋ Anonymous Goldfinch 12mth #1042adc

sorry can u express the answer more in detail?

or what kind of situation is "far enough"

♡ ...

J **Justin Yokota** STAFF 12mth #1042ade

↩ Replying to Anonymous Goldfinch

To use an analogy, I'm going to think of some English words, then change up to one letter in each word and write it (so if I think about the word "bat", I may write down "bat", "aat", "cat", "bot", "ban", "bqt", etc.). For which of the following words could you detect that I had changed a letter, and for which of the following words can you deduce what my original word was (knowing that the original word was a valid English word)? Why?

bat

sthdy

cail

weight

creste

scientific

♡ ...

 **Anonymous Goldfinch** 12mth #1042adf

↩ Replying to Justin Yokota

icic ,but if there is an exact number to describe "how far"

♡ ...

J **Justin Yokota** STAFF 12mth #1042aeb

↩ Replying to Anonymous Goldfinch

Well, how far apart do two words in English need to be before it's unambiguous?

♡ ...

 **Anonymous Goldfinch** 12mth #1042aec

↩ Replying to Justin Yokota

idk haha,for me maybe 1 to 2 letters?

♡ ...

J **Justin Yokota** STAFF 12mth #1042aed

↩ Replying to Anonymous Goldfinch

What were the possible words for each line? How far apart were they?

♡ ...

 **Anonymous Goldfinch** 12mth #1042aee

↩ Replying to Justin Yokota

bat 0

study 1

call 1

weight 0

create 1

scientific 0



J **Justin Yokota** STAFF 12mth #1042aef

↩ Replying to Anonymous Goldfinch

Were there any other possible words? Some lines could have started from multiple different words (and thus were ambiguous).



 **Anonymous Goldfinch** 12mth #1042afa

↩ Replying to Justin Yokota

ohoh so we must keep them "far" enough ,but there is no exact distance



J **Justin Yokota** STAFF 12mth #1042afc

↩ Replying to Anonymous Goldfinch

Well, there is an exact distance. But you should be able to deduce that distance from first principles.



 **Anonymous Goldfinch** 12mth #1042afd

↩ Replying to Justin Yokota

sorry i am little confused



J **Justin Yokota** STAFF 12mth #1042afe

↩ Replying to Anonymous Goldfinch

For each of the lines in my example, find ALL possible words that I could have started with; most of these lines will have more than one possible word. In order to error-correct, there must be only one possible word in a line (since otherwise you don't know which word I started with).



 **Anonymous Goldfinch** 12mth #1042aff

↩ Replying to Justin Yokota

bat bot but ...



J **Justin Yokota** STAFF 12mth #1042bab

↩ Replying to Anonymous Goldfinch

Yes, but there are more, like "cat" and "ban". What about for the other lines?



 **Anonymous Goldfinch** 12mth #1042bac

↩ Replying to Justin Yokota

haha yeah

study

call calt

weight



**J Justin Yokota** STAFF 12mth #1042bad  
↩ Replying to Anonymous Goldfinch  
"calt" is two letters away from "cail"; I could not have started with that. There's also another word close to "weight".  
♡ ...

**Anonymous Goldfinch** 12mth #1042bae  
↩ Replying to Justin Yokota  
what word lol  
♡ ...

**J Justin Yokota** STAFF 12mth #1042baf  
↩ Replying to Anonymous Goldfinch  
"height", "weighs", "wright"  
♡ ...

**Anonymous Goldfinch** 12mth #1042bba  
↩ Replying to Justin Yokota  
oh yes  
♡ ...

**J Justin Yokota** STAFF 12mth #1042bbb  
↩ Replying to Anonymous Goldfinch  
Now, consider a universe where the only words in the English language were "height" and "wright". If I wrote "weight", would that word be ambiguous (you can't tell what word I started with)?  
♡ ...

**Anonymous Goldfinch** 12mth #1042bbc  
↩ Replying to Justin Yokota  
yes i will be kind of ambiguous about whether it is "height "or "wright"  
♡ ...

**J Justin Yokota** STAFF 12mth #1042bca  
↩ Replying to Anonymous Goldfinch  
What if the only two words in the English language were "carrot" and "marmot"? Is there anything I could possibly write that is ambiguous?  
♡ ...

**Anonymous Goldfinch** 12mth #1042bcb  
↩ Replying to Justin Yokota  
carmot or marrot  
♡ ...

**J Justin Yokota** STAFF 12mth #1042bcc  
↩ Replying to Anonymous Goldfinch  
How did you get those words? In general, if I asked you to find an ambiguous word if the English language had only the two words X and Y, when would you be able to find an ambiguous word?  
♡ ...

**Anonymous Goldfinch** 12mth #1042bcd  
↩ Replying to Justin Yokota



sorry i dont make sense



J **Justin Yokota** STAFF 12mth #1042bce

↩ Replying to Anonymous Goldfinch

How did you determine that carmot and marrot were ambiguous? What property of the words "carrot" and "marmot" did you rely on to create those words (i.e. if another pair of words had this property, you'd be able to find ambiguous words in the same way)?



 **Anonymous Goldfinch** 12mth #1042bcf

↩ Replying to Justin Yokota

they have same difference



J **Justin Yokota** STAFF 12mth #1042bda

↩ Replying to Anonymous Goldfinch

Well, "abacus" also is the same difference from "carrot" and "marmot" (all 6 characters are different). Why wasn't this an ambiguous word?



 **Anonymous Goldfinch** 12mth #1042bdb

↩ Replying to Justin Yokota

so far?



J **Justin Yokota** STAFF 12mth #1042bdc

↩ Replying to Anonymous Goldfinch

I'm not quite sure I understand? Can you elaborate on what you mean by "so far", and what you think has changed?



 **Anonymous Goldfinch** 12mth #1042bdd

↩ Replying to Justin Yokota

i mean the "abacus" is so different from the raw two words

even though it has the same difference, but it is totally different with "carrot" or "marmot"



J **Justin Yokota** STAFF 12mth #1042bde

↩ Replying to Anonymous Goldfinch

Ah, yes. "abacus" is far enough from the target words that it would have been impossible for me to even write that. So it's not just that "carmot" was the same distance away from "carrot" and "marmot". What specifically makes "carmot" ambiguous?



 **Anonymous Goldfinch** 12mth #1042bdf

↩ Replying to Justin Yokota

not so far away and has same difference with raw words



J **Justin Yokota** STAFF 12mth #1042bea  
↩ Replying to Anonymous Goldfinch  
What specific distance? Also, if the only two words in the English language were "weight" and "height", would there be an ambiguous word I could write?  
♡ ...

**Anonymous Goldfinch** 12mth #1042beb  
↩ Replying to Justin Yokota  
no?  
♡ ...

J **Justin Yokota** STAFF 12mth #1042bec  
↩ Replying to Anonymous Goldfinch  
There is an ambiguity. If I write "weight", do you know for a fact that I wrote "weight", or is it possible that I meant to write "height" and the first letter was changed?  
♡ ...

**Anonymous Goldfinch** 12mth #1042bed  
↩ Replying to Justin Yokota  
i think i dont know, the two words are so similiar  
♡ ...

J **Justin Yokota** STAFF 12mth #1042bee  
↩ Replying to Anonymous Goldfinch  
Yup. If two words are too similar, you can't disambiguate. Exactly how similar do two words need to be?  
♡ ...

**Anonymous Goldfinch** 12mth #1042bef  
↩ Replying to Justin Yokota  
emmm maybe "not so similiar" haha  
♡ ...

J **Justin Yokota** STAFF 12mth #1042bfa  
↩ Replying to Anonymous Goldfinch  
What is that threshold? "height" and "weight" led to an ambiguity, as did "carrot" and "marmot". At what point do words become "not so similar"?  
♡ ...

**Anonymous Goldfinch** 12mth #1042bfb  
↩ Replying to Justin Yokota  
sorry idk  
♡ ...

J **Justin Yokota** STAFF 12mth #1042bfc  
↩ Replying to Anonymous Goldfinch  
Why was "carmot" ambiguous? What is the quantifiable property that "carmot" had that made it ambiguous (and words like "abacus" not ambiguous)?  
♡ ...

**Anonymous Goldfinch** 12mth #1042cad  
↩ Replying to Justin Yokota

it is close to the two words and have same difference

♡ ...

J **Justin Yokota** STAFF 12mth #1042cae

↩ Replying to Anonymous Goldfinch

How close is close? Did "weight" have the same difference from "height" and "weight"?

♡ ...

 **Anonymous Goldfinch** 12mth #1042caf

↩ Replying to Justin Yokota

$\text{diff}(\text{weight}, \text{height}) = 1$

$\text{diff}(\text{weight}, \text{weight}) = 0$

♡ ...

J **Justin Yokota** STAFF 12mth #1042cba

↩ Replying to Anonymous Goldfinch

So what is required for a word to be ambiguous? What is the reason why "carmot" and "weight" were ambiguous, but "abacus" wasn't?

♡ ...

 **Anonymous Goldfinch** 12mth #1042cbb

↩ Replying to Justin Yokota

the "abacus" is so different from the raw two words

even though it has the same difference, but it is totally different with "carrot" or "marmot"

♡ ...

J **Justin Yokota** STAFF 12mth #1042cbc

↩ Replying to Anonymous Goldfinch

Yes, but what is the quantitative difference? A computer can't say "This is very different"; what specifically is needed? As another example, is "carmat" ambiguous in the same way "carmot" is ambiguous?

♡ ...

 **Anonymous Opossum** 12mth #1042add

I have the same questions Goldfinch asked. Why "too close" -> ambiguous result?

♡ ...

 **Anonymous Fox** 12mth #1042ccc

↩ Replying to Anonymous Opossum

I followed along to try to figure this out, too, because I also initially didn't follow until I really combed through the examples, but I think I've got it now. To detect a single letter error in words, you need a minimum of three letters difference between two code words. Carrot and marmot have two letters difference in **carrot** and **marmot** so you can formulate a word that is one away from each, like **carmot**. With three letters difference, this no longer happens.

I'll use **harper** and **marked** as my example words. All possible one letter mistakes either mess up the first, fourth, or sixth letter. Since only one of those three are affected, the other two must be fine, and you can reconstruct the word from those two unaffected letters.

♡ ...

J **Justin Yokota** STAFF 12mth #1042cec

↩ Replying to Anonymous Fox

Yup!

♡ ...

 **Anonymous Fish** 12mth #1042cfc

↩ Replying to Justin Yokota

is it a general rule that when you want d-bit error correction that you need a minimum hamming distance of  $2d+1$ ? If the hamming distance is less than this, then you can't reconstruct your bitstring / word?

♡ ...

J **Justin Yokota** STAFF 12mth #1042dac

↩ Replying to Anonymous Fish

Yup. Can you prove that?

♡ ...

 **Anonymous Seahorse** 12mth #1042da

✓ Resolved

Sp21 final q 1d

#### (d) Q4

Which of the following code blocks would see a performance improvement if we placed a `#pragma omp for` over the outer for loop? Select all that apply.

i. A

```
for (int i = 0; i < 5000 - 3; i += 3) {  
    a[i+2] = a[i] + a[i+1];  
}
```

B

```
for (int i = 0; i < 5000 - 2; i += 2) {  
    a[i+2] *= a[i];  
}
```

C

```
for (int i = 0; i < 5000 - 3; i++) {  
    a[i+2] = a[i] + a[i+1];  
}
```

D

```
for (int i = 0; i < 5000; i++) {  
    a[i] = 100;  
}
```

A

B

C

D

E. None of the above

Isn't there a chance that for D, we lose spatial locality in caches, and it would be slower, even if we used multithreading? I remember one of the hw questions was similar to this

♡ ...

 J **Justin Yokota** STAFF 12mth #1042ace

Not likely, since we still need to hit the same set of blocks anyway; we don't lose spatial locality because we needed to go that distance in the naive code.

♡ ...

 **Anonymous Ferret** 12mth #1042ce ✓ Resolved

for fa21 final q7.4, what would a capacity miss look like for the problem?

Also for the second question in 7.4, why couldn't it be a capacity miss? The solution says the cache isn't full yet, but for that index, the cache is full.

♡ 1 ...

 **Jero Wang** STAFF 12mth #1042eca

Capacity miss occurs only when the entire cache is full (when decreasing associativity won't make this a hit).

♡ ...

 **Anonymous Ferret** 12mth #1042cd ✓ Resolved


for fa21 final q6.3, is there an algorithm or way to solve the problem besides just inspection?

♡ ...

 **Justin Yokota** STAFF 12mth #1042acf

The problem is NP-hard in general, so no polynomial-time algorithm is known to exist.

♡ ...

 **Anonymous Ferret** 12mth #1042cc ✓ Resolved

For fa21 final q5.3, would the following be correct?

li s3 8

li s2 7

li s1 6


mac s1 s2 s3

♡ ...

 **Jero Wang** STAFF 12mth #1042eea

Yeah, that should be correct.

♡ ...

 **Anonymous Jaguar** 12mth #1042cb ✓ Resolved

**sp21-Final-q2K**


Is the 2-way page table in scope? I don't think it's covered in this semester's lecture

♡ 1 ...

 **Jero Wang** STAFF 12mth #1042ecb

No, 2 level PTs are not in scope.

♡ ...

 **Anonymous Cassowary** 12mth #1042be ✓ Resolved

**fa21-Final-q2.2**

Can someone explain how we get the answer for this question, specifically, I am confused how we made this step "0b1.0101 \* 2<sup>-2</sup> , or 0b10101 \* 2<sup>-6</sup>" Thanks.

Q2.2 (3.5 points) 1/3 (whose binary representation is 0b0.0101 0101...)

**Solution:** Answer:  $21 * 2^{-6} = 0.328125$ . We can move the binary point to get our floating point representation  $0b1.010101... * 2^{-2}$ . The mantissa rounds down, so our float would be  $0b1.0101 * 2^{-2}$ , or  $0b10101 * 2^{-6}$  or  $21 * 2^{-6}$

♡ 1 ...

 **Anonymous Moose** 12mth #1042aaa

Our goal is to put it into the format of a floating point number first. ie.  $1.mantissa * 2^{(exp+bias)}$

Step 1. Take the existing binary representation and shift the decimal over twice to get:  $1.01010101... * 2^{-2}$

Step 2: We have a repeating mantissa without end, however because of our limited bits, we have to cut off and round. The closest rounded value is  $1.0101 * 2^{-2}$  given that the problem limits us to 4 mantissa bits.


♡ ...

 **Anonymous Opossum** 12mth #1042bd Unresolved

**sp21-MT-5c(iv)**

Why the answer is 10? I calculated 8 NOPs(2 for 1/2, 3 for 2/3, 3 for 2/3). During Branch instruction, at which stage we can know whether the branching is successful?

♡ ...

 **Anonymous Pigeon** 12mth #1042bc Unresolved

**SP21-Final-Q2**

Q3.2 (2 points) The Mean Time Between Failures (MTBF) for this particular data center is 4000 hours. The Mean Time To Repair (MTTR) is 3 hours. What is the availability for this datacenter? Express your answer as a single simplified fraction.

**Solution:**  $\frac{3997}{4000}$

MTBF only tells you how much time passes between failures; it doesn't account for repair time.

Every 4000 hours, there is 1 failure. That failure takes 3 hours to repair. Therefore, the system is up for 3997 hours, for an availability of  $\frac{3997}{4000}$ .

The solution here says MTBF doesn't account for repair time, but from my understanding (and the lecture slides)  $MTBF = MTTF + MTTR$ . Doesn't this include repair time?

$MTBF = 4000$ , so availability is  $\frac{4000-3}{4000}$ .

♡ ...

 **Anonymous Chimpanzee** 12mth #1042ba Unresolved

**sp21 midterm 5c iii**

Why is the answer Data and not Structural. Isn't shw an invalid risc v function? Wouldn't that make this a structural hazard?

♡ ...



Anonymous Chimpanzee 12mth #1042af

Unresolved

### sp21 Final Q7 Part b Cached Datapath

The solution says 'See Note'. Where is this note?

Also, how can one replace IMEM and DMEM with Cache? Can someone pls break this down for me.

♡ ...



Anonymous Seahorse 12mth #1042ae

Resolved

fa23 q3.1-

Do we write "0xEF 0xBE 0xAD 0xDE" because when we are writing, we are moving up the stack pointer so we want less significant values to be lower on the stack pointer position?

♡ ...



Anonymous Armadillo 12mth #1042baa

For this question, if we added a null character (0x00) at the end of the string would we get credit?

♡ ...



Jero Wang STAFF 12mth #1042eec

I don't remember how this question was graded, but I assume a small point penalty since it might have messed up other values saved on the stack.

♡ ...



Jero Wang STAFF 12mth #1042eeb

This is because we're in a little endian system, and we write from lower to higher addresses. The least significant byte is at the lowest address.

♡ ...



Anonymous Seahorse 12mth #1042ad

Resolved

fa21 q1.8-

... .., we don't have any redundancy, so we get a capacity of 9 TiB

Q1.8 (0.5 points) RAID 5?

TiB

**Solution:** In RAID 5, we get one parity block for every 9 data blocks, so we get a total capacity of 9 TiB

Could someone explain how we get 9 TiB? Im confused of the breakdown

♡ ...



Anonymous Moose 12mth #1042aab

We have 10 disks (aka 10 cylinders), and in RAID 5, for each row of cylinders, you need one disk that's dedicated to storing the parity.

10 - 1 parity disk = 9 data disks

♡ ...



**Anonymous Boar** 12mth #1042eaa

Is this somewhere on the slides, I can't seem to find it. All I can find is "block level striping, but distribute the parity across all disks"

♡ ...



**Anonymous Ferret** 12mth #1042d

✓ Resolved

Is q1.4 in the fa21 final in scope?

♡ 1 ...



**Nikhil Kandkur** STAFF 12mth #1042e

Yup! Check out the I/O lecture for more information on this.

♡ ...



**Anonymous Lapwing** 12mth #1042bf

FA21-Final-1.4:

I am wondering why the opposite is true. In I/O lecture slide, it's said that we use polling(or interrupt) for low data rates and interrupt & DMA for high data rates.

♡ 1 ...



**Anonymous Walrus** 12mth #1042aea

↩ Replying to Anonymous Lapwing

^^ same

♡ ...



**Anonymous Whale** 12mth #1042bfe

↩ Replying to Anonymous Lapwing

I think the idea behind it is that polling is not efficient for active devices/devices that transfer a lot of data since it's effectively similar to a loop like `while (read_status(device) == 1) { // keep waiting }`.

An analogy would be like when we were little and our parents told us to wait by the door to open it for any guests to arrive. If the guests/data arrive at a fast rate(high data rate) then it is pretty efficient for us to just wait by the door otherwise if we choose to keep playing video games with our friends then we keep having to go open the door continuously which is annoying and inefficient. Interrupts are like the door bell so if guests are arriving at a slow rate like 2 guests an hour(low data rate) then it would be better for us to just hang around playing games and wait for the doorbell to then go open it.

♡ ...



**Anonymous Lapwing** 12mth #1042cab

↩ Replying to Anonymous Whale

Thank you! But I thought data rate means the amount of data transferred instead of frequency.

♡ ...



**Anonymous Llama** 12mth #1042c

✓ Resolved

SP21-Final-2H:



J. If our caching system remains as seen in question 1, how many caches would be needed to fully fit our page table? Give your answer as a decimal to two decimal places.

8

Alternate solution: 6 caches

How do I setup the problem?

I was thinking (page entry size) \* (total #pages) / (cache size) but that didn't yield 8 or 6.

♡ ...



N

Nikhil Kandkur STAFF 12mth #1042f

Your setup is correct, but there might have been some confusion over page entry size:  
page entry size = #metadata bits +  $\log_2(\# \text{ physical pages}) = 8 + \log_2(16 \text{ MiB}/4 \text{ KiB}) = 8 + 12$   
bits = 20 bits.

Unfortunately alignment was not specified, so students either rounded it up to the nearest byte (24 bits = 3 bytes), or the nearest word (32 bits = 4 bytes). You would proceed with your calculations from there.

♡ ...



Anonymous Llama 12mth #1042a

✓ Resolved

FA21-Final-Q8.9:

Regardless of your answer to Q8.2, assume that physical addresses are 20 bits long. We run two programs (with no shared memory), which access the following virtual memory addresses in order. For each memory access, determine the physical address that gets accessed, writing your answer in hexadecimal.

Assume that no physical pages are in use prior to the first memory access, and that physical pages get assigned in order of physical page number (so page 0 is assigned first, then page 1, and so on).

Q8.9 (1 point) Program 1: 0x12345664

**Solution:** This is a hit on program 1's version of this page, so we reuse the same page number from before. We get the physical address 0x01664.

Why are we able to get a hit on program 1's version of this page? "0x01678" should be stored at index 01 at the current time and  $664 \neq 678$ .

♡ ...



N

Nikhil Kandkur STAFF 12mth #1042aa

By the time we get to 8.9, the valid entries in each of the program's page tables would look like this:

Program 1:

VPN: 0xABCDE, PPN: 0x00

VPN: 0x12345, PPN: 0x01

VPN: 0xABCDD, PPN: 0x02

Program 2:

VPN: 0xABCDE, PPN: 0x03

VPN: 0x12345, PPN: 0x04

When we get to 8.9, we see that our VPN is 0x12345, and we already have a PPN set for that VPN in the Program 1 page table, so we do not need to get a new free physical page. With that in mind, we just get the PPN that is mapped to VPN 0x12345, which is 0x01, and we append our offset to the PPN, which is 0x664, so we get a physical address of 0x01664.

♡ 1 ...



*This comment was deleted*



N

**Nikhil Kandkur** STAFF 12mth #1042ac

[#1044bb](#)

♡ ...