

PRINT your name: \_\_\_\_\_,  
(last) (first)

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

---

You have 110 minutes. There are 8 questions of varying credit (100 points total).

Question:	1	2	3	4	5	6	7	8	Total
Points:	7	16	20	6	15	7	16	13	100

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

- Unselected option (completely unfilled)
- Only one selected option (completely filled)

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

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

Anything you write that you ~~cross-out~~ will not be graded. Anything you write outside the answer boxes will not be graded.

If an answer requires hex input, make sure you only use capitalized letters! For example, `0xDEADBEEF` instead of `0xdeadbeef`. Please include hex (0x) or binary (0b) prefixes in your answers unless otherwise specified. For all other bases, do not add any prefixes or suffixes.

---

**Read the following honor code and sign your name.**

I understand that I may not collaborate with anyone else on this exam, or cheat in any way. I am aware of the Berkeley Campus Code of Student Conduct and acknowledge that academic misconduct will be reported to the Center for Student Conduct and may further result in, at minimum, negative points on the exam.

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

Questions begin on the next page.

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

Q1.1 (4 points) Select the 2 Boolean expressions from below that are equivalent to each other:

- $(A+!B)(A+!C)$
- $!B!C + A$
- $A+!(B + C) + (!B+!C)(!B + C)$
- $(A + B)(A + C)(!BC)$

Convert the following 2-byte hex numbers to decimal using signed 2's complement.

Q1.2 (1 point) 0xFF40

Q1.3 (1 point) 0x009C

Q1.4 (1 point) Consider a 12-bit biased number representation scheme with a bias of  $-2047$ . Which of the following is not a representable number in this scheme?

- 0
- 2048
- $-2048$
- $-1$

**Q2 C: Filed Away****(16 points)**

You are bored over summer break so you decide to write up a file system in C!

The struct `file_item` represents a file or a folder. The data union holds either the contents of the file (a string), or an array of pointers to children `file_items`.

In this question, assume that pointers are 4 bytes long.

```
1 typedef struct file_item {
2     char *name;
3     bool is_folder;
4     file_item_data data;
5 } file_item;
6
7 typedef union file_item_data {
8     char contents[X];
9     struct file_item* children[16];
10 } file_item_data;
11
12 // Copies all characters from src to dest including the NULL terminator
13 char *strcpy(char *dest, const char *src);
```

We set `X` to be the largest possible value that doesn't increase the union size. What is the `strlen` of the longest string we can store in a single file?

Fill in the code on the next page to create files and folders. Your code must still work even if the input strings are freed later on. Assume that the input strings will fit inside of a `file_item_data` union.

You may use fewer lines than provided, but you may not add more lines.

(Question 2 continued...)

```
1 /* Creates a file with the given name and contents,  
2    and returns a pointer to that file. */  
3 file_item* create_file(char* contents, const char *name) {  
  
4     _____  
  
5     _____  
  
6     _____  
  
7     _____  
  
8     _____  
  
9     _____  
  
10    _____  
  
11    _____  
  
12    _____  
  
13    _____  
14 }  
15 /* Creates a folder with the given name and no children,  
16    and returns a pointer to that folder. */  
17 file_item* create_folder(const char *name) {  
  
18    _____  
  
19    _____  
  
20    _____  
  
21    _____  
  
22    _____  
  
23    _____  
  
24    _____  
  
25    _____  
  
26    _____  
  
27    _____  
28 }
```

(Question 2 continued...)

**Q3 Error-Introducing Code****(20 points)**

In machine learning, some data scientists add random noise to a training dataset in order to improve their models. Here, we will take a dataset and transform it into usable data in RISC-V!

The function `jitter` is defined as follows:

- Inputs:
  - `a0` holds a pointer to an integer array.
  - `a1` holds a buffer large enough for `n` integers (which does not overlap with the array in `a0`).
  - `a2` holds `n`, the length of the arrays.
- Output:
  - `a0` holds a pointer to the buffer originally in `a1`. The buffer is filled with the result of calling `noisen` on each element in the `a0` array.

The function `noisen` is defined as follows:

- Input: `a0` holds an integer.
- Output: `a0` returns the integer with noise added.

Eric has provided the correct implementation of `jitter` below, following calling convention:

```
1 jitter:
2 # BEGIN PROLOGUE
3 addi sp sp <BLANK 1>
4 # (multiple lines omitted)
5 # END PROLOGUE
6 mv s0 a0
7 mv s1 a1 # Hold beginning of output arr
8 mv s2 a1
9 mv s3 a2 # Hold counter
10 loop:
11 beq s3 x0 end
12 lw a0 0(s0)
13 jal ra noisen
14 sw a0 0(s1)
15 addi s0 s0 4
16 addi s1 s1 4
17 addi s3 s3 -1
18 j loop
19 end:
20 mv a0 s2
21 # BEGIN EPILOGUE
22 # (multiple lines omitted)
23 addi sp sp <BLANK 2>
24 # END EPILOGUE
25 ret
```

(Question 3 continued...)

Q3.1 (1 point) To follow calling convention, what numbers should go in the blanks?

<BLANK 1>

<BLANK 2>

Q3.2 (1 point) List all registers that Eric needs to save on the stack in order to follow calling convention.

Q3.3 (6 points) Write a sequence of at most two instructions or pseudoinstructions that are equivalent to the `j loop` instruction.

You must use a `jalr` instruction or `jalr` pseudoinstruction in at least one of the blanks. You may not use a `jal` instruction, branch instruction, or `jal` pseudoinstruction in any of the blanks.

1	_____
2	_____



(Question 3 continued...)

Now, Eric wants to implement `noisen` to add some random offset to an integer `a0`. Unfortunately, Eric only has access to the `randomBool` function, which takes in no inputs and randomly returns either 1 or 0 in `a0`.

If Eric implemented `noisen` to return `a0+randomBool()`, the integer would always get shifted in the positive direction. Instead, Eric suggests implementing `noisen` so that `noisen` alternates between returning `a0+randomBool()` and returning `a0-randomBool()`.

Q3.4 (12 points) Fill in the blanks to complete Eric's suggested implementation of `noisen`.

Assume that you can read from and write to any memory addresses, in any section of memory.

```
1 noisen:
2  addi sp sp -8 # Prologue
3  sw ra 0(sp)
4  sw s0 4(sp)
5  mv s0 a0

6  jal ra randomBool
7  add s0 s0 a0

8  _____ # one RISC-V instruction

9  xori t0 t0 _____ # some immediate

10 _____ # one RISC-V instruction

11 mv a0 s0 # Epilogue
12 lw ra 0(sp)
13 lw s0 4(sp)
14 addi sp sp 8
15 ret
```

**Q4 Floating Loading****(6 points)**

RISC-V has a floating-point extension to support IEEE-754 single-precision floats (1 sign bit, 8 exponent bits, 23 mantissa bits). The extension adds the instructions in the following table, as well as 32 floating-point registers (f0 to f31) that each hold 32 bits.

Instruction	Name	Description
<code>fadd.s rd rs1 rs2</code>	Float ADD	$rd = rs1 + rs2$
<code>fsub.s rd rs1 rs2</code>	Float SUBtract	$rd = rs1 - rs2$
<code>fmul.s rd rs1 rs2</code>	Float MULTiply	$rd = rs1 * rs2$
<code>fdiv.s rd rs1 rs2</code>	Float DIVide	$rd = rs1 / rs2$
<code>fmv.w.x rd rs1</code>	MoVe from Integer	$rd = ((float) rs1)$

Note that the 4 floating-point arithmetic instructions can only operate on floating-point registers, and floating-point registers cannot be used with base instructions. For example, `fadd.s f0 t1 t2` and `addi f0 t1 5` are not valid instructions.

`fmv.w.x` takes the bitstring in an integer register `rs1`, and transfers the bitstring over to a floating-point register `rd`.

For example, the value 4 is represented in single-precision floating point numbers as `0x4080 0000`. If `t0` contained the hexadecimal value `0x4080 0000` and we ran `fmv.w.x f0 t0`, then `f0` would be set to the floating point value 4.0.

Q4.1 (2 points) Translate the value 3 into its single-precision floating point representation, in hexadecimal.

0x

Q4.2 (4 points) Write instructions to put the float (as close as possible to) 1.33333... into a register `f1`. You may not use any floating-point instructions outside the five listed above.

Only one RISC-V instruction or pseudoinstruction is allowed per line. You may use fewer lines than provided, but you may not add more lines.

1 \_\_\_\_\_

2 \_\_\_\_\_

3 \_\_\_\_\_

4 \_\_\_\_\_

5 \_\_\_\_\_

**Q5 Error-Correcting Code****(15 points)**

Recall that a Hamming Error Correcting Code can be used to fix single-bit errors. In situations where data is error-prone (such as on satellites, or in L1 caches), it is often useful to store blocks of data in Hamming codes, and internally fix errors during memory retrieval.

For this question, we will consider a (7,4) Hamming Code with even parity, which uses 7 total bits to store 4 data bits. Recall the following bit pattern for the (7,4) Hamming Code:

Bit	1	2	3	4	5	6	7
Transmitted Bit	$p_1$	$p_2$	$d_1$	$p_4$	$d_2$	$d_3$	$d_4$
$p_1$	✓		✓		✓		✓
$p_2$		✓	✓			✓	✓
$p_4$				✓	✓	✓	✓

We adopt the convention that bit 1 is the most significant bit of the data, and bit 7 is the least significant bit.

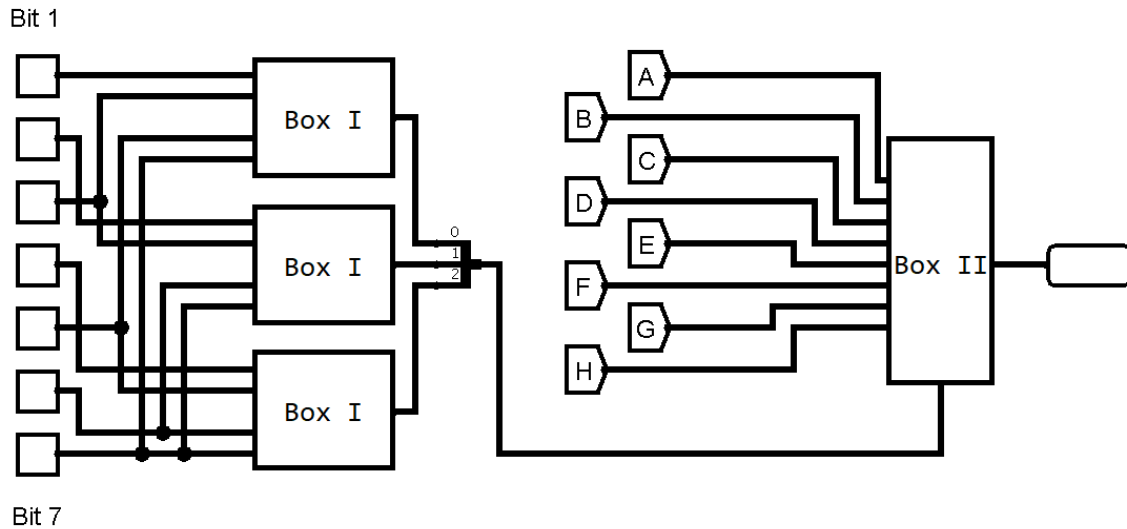
Q5.1 (3 points) In order to store a full byte, we concatenate two (7,4) Hamming codes, with the first 7 bits storing the most significant nibble of data, and the last 7 bits storing the least significant nibble. After storing a byte, we retrieve the following raw data (spacing has been added for readability):

0b 1001110 100011

After error correction, what byte gets returned? Express your answer in hexadecimal.

(Question 5 continued...)

We construct a circuit that, given a (7,4) Hamming Code, outputs a corrected nibble of data. What components should be placed in the labeled boxes? Assume that multiple input gates are made by chaining two-input gates.



Q5.2 (2 points) What goes in Box I?

- AND gate
- OR gate
- XOR gate
- Adder gate
- Multiplier gate
- Multiplexer (with 0 input on top)
- Demultiplexer (with 0 input on top)
- Priority Encoder (with 0 input on top)

Q5.3 (2 points) What goes in Box II?

- AND gate
- OR gate
- XOR gate
- Adder gate
- Multiplier gate
- Multiplexer (with 0 input on top)
- Demultiplexer (with 0 input on top)
- Priority Encoder (with 0 input on top)

Q5.4 (2 points) Which four labels among A-H should have the same value?

- |                            |                            |                            |                            |
|----------------------------|----------------------------|----------------------------|----------------------------|
| <input type="checkbox"/> A | <input type="checkbox"/> C | <input type="checkbox"/> E | <input type="checkbox"/> G |
| <input type="checkbox"/> B | <input type="checkbox"/> D | <input type="checkbox"/> F | <input type="checkbox"/> H |

(Question 5 continued...)

Q5.5 (4 points) Assume that the component in Box I has a delay of 3 ns, the component in Box II has a delay of 5 ns, and that computing the values of labels A-H take 1 ns. If inputs arrive at time 0, what is the earliest and latest time the output can change in response?

Earliest:


Latest:

Q5.6 (2 points) Is it more important to add this component around the IMEM or the DMEM? Explain your reasoning in 20 words or fewer.

**Q6 00100**

**(7 points)**

Consider the following FSM:



On each transition, we receive a 1-bit input, and output a character.

For the following inputs, what string would the FSM output?

Q6.1 (1 point) 000100

Q6.2 (1 point) 0110

Q6.3 (1 point) 010000

Q6.4 (4 points) If the FSM ends in a state labeled *T* after processing the entire input, then the input is accepted. Otherwise, the input is rejected.

Write a rule for determining whether an input will be accepted by this FSM.

**Q7 Virtual Memory**

**(16 points)**

Assume the virtual memory address 0xABCD123 maps to the physical address 0x123123.

Q7.1 (1.5 points) What is the maximum possible page size? Include units in your answer.

Q7.2 (1.5 points) What is the minimum possible size of virtual memory? Include units in your answer.

Q7.3 (2.5 points) This subpart is independent of the previous subparts.

We have 1 GiB of virtual memory, 24-bit physical addresses, a 4 KiB page size, and a single-level page table.

If a page table entry has 4 bits of metadata, how many bits is a page table entry, with no padding?

(Question 7 continued...)

This subpart is independent of the previous subparts.

We have 4 GiB of virtual memory, 16 MiB of physical memory, a 4 KiB page size, and a single-level page table. We also have a 2-entry, fully-associative TLB with LRU replacement policy.

Assume that the TLB starts empty, and that physical pages are assigned in order starting with page 0 (e.g. page 0, page 1, etc.). Also assume that we are working with a single-core system that will context-switch between processes.

Q7.4 (10.5 points) Each row in the table represents a memory access. For each row, fill out the corresponding physical address, and whether the access causes a TLB hit, a TLB miss but page table hit, or a page fault.

Virtual Address	Process ID	Physical Address	Access Type
0xDEADBEEF	1	0x000EEF	<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault
0xDEADBEEF	2		<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault
0x0000061C	2		<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault
0xDEADB61C	2		<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault
0xDEADB61B	1		<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault
0xDEADB61C	1		<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault
0x0000061A	2		<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault
0x0000061A	1		<input type="radio"/> TLB Hit <input type="radio"/> TLB miss, page table hit <input type="radio"/> Page Fault



**Q8 DLP Meets TLP****(13 points)**

Grace wants to write a function called `vector_mul_positive`, defined as follows:

Inputs:

- `a`, a pointer to an integer array
- `b`, a pointer to an integer array
- `N`, the length of each array. You may assume that `N` is a multiple of 4.

Outputs:

- Compute the element-wise products, and output the sum of only the positive products.

Example input: `a=[-1, 2, 3, 4]`, `b=[5, 6, -7, 8]`, and `N = 4`.

Example output:  $2 \times 6 + 4 \times 8 = 44$ . Note we skip  $-1 \times 5$  and  $3 \times -7$  because these products are negative.

For this question, you may use these SIMD functions (`__m128i` represents packed signed 32-bit integers):

`__m128i vmul(__m128i a, __m128i b)`: Return the result of multiply values in `a` and `b`

`__m128i vset (int32_t x)`: Set the 4 signed 32-bit integers to `x`

`__m128i vload (__m128i* a)`: Return 128-bits of integer data stored at pointer `a`

`void vstore (__m128i* b, __m128i a)`: Store 128-bits of integer data from `a` into `b`

`__m128i vcmpgt (__m128i a, __m128i b)`: Compare `a` and `b` for greater-than and return result

`__m128i vadd (__m128i a, __m128i b)`: Return the addition if `a` and `b`

`__m128i vsub (__m128i a, __m128i b)`: Return the subtraction of `b` from `a`

`__m128i vxor (__m128i a, __m128i b)`: Return the bitwise XOR of `a` and `b`

`__m128i vor (__m128i a, __m128i b)`: Return the bitwise OR of `a` and `b`

`__m128i vand (__m128i a, __m128i b)`: Return the bitwise AND of `a` and `b`

(Question 8 continued...)

Implement a parallelized version of `vector_mul_positive`. You may use fewer lines than provided, but you may not add more lines.

```
1 int32_t vector_mul_positive(int32_t *a, int32_t *b, int32_t N) {
2     int32_t result[4];

3     __m128i sum_v = _____;
4     __m128i cond_v = _____;
5     #pragma _____
6     for (_____ ) {
7         __m128i curr_v1 = vload(_____);
8         __m128i curr_v2 = vload(_____);
9         __m128i mul = _____;
10        __m128i tmp = _____;
11        _____;
12        #pragma _____
13        _____;
14    }
15    _____;
16    _____;
17 }
```

Nothing on this page will be graded.

Is there anything you want us to know?