

Final Exam

CSC 252

12 May 2021

Computer Science Department
University of Rochester

Instructor: Yuhao Zhu

TAs: Rongcui Dong, Elana Elman, Kalen Frieberg, Sudhanshu Gupta, Yiyao (Jack) Yu, Vladimir Maksimovski, Nathan Reed, Raffaello Sanna

Name: _____

Problem 0 (3 points):

Problem 1 (14 points):

Problem 2 (15 points):

Problem 3 (16 points):

Problem 4 (22 points):

Problem 5 (30 points)

Total (100 points):

Extra Credit (15 points)

Remember “**I don’t know**” is given 15% partial credit, but you must erase everything else. This does not apply to extra credit questions.

Your answers to all questions must be contained in the given boxes. Use spare space to show all supporting work to earn partial credit.

You have 3 hours to work.

Please sign the following. I have not given nor received any unauthorized help on this exam.

Signature: _____

GOOD LUCK!!!

Problem 0: Warm-up (3 Points)

What's the most surprising thing about computers you learned from 252?

Problem 1: Miscellaneous (14 points + 12 points extra credit)

Part a) (3 points) Convert the decimal number 334 to hexadecimal.

Part b) (3 points) Why wouldn't a user program that contains an infinite loop freeze the entire computer even if the computer has only one processor?

Part c) (8 points) Consider a shell in Linux, `sh`, that supports basic foreground/background job control, just like the one you implemented in Assignment 4. Initially, `myprogA` is running in the foreground, `myprogB` is running in the background as job 1, and `myprogC` is stopped in the background as job 2. **Assume this initial state for each question below.**

(4 points) The user presses `Ctrl-C` and kills the foreground process. List all signals that are generated and the receivers of these signals.

(4 points) `myprogA` finishes, then the user presses `Ctrl-Z` at the shell's command prompt. List all the signals that are generated and the receivers of these signals.

--

Part d) (this entire part is extra credit; 12 points) You are designing a new ISA called URISA, which supports the following instructions.

Instruction	Semantics
<code>ADDI rd, rs1, imm</code>	Sum immediate <code>imm</code> and register <code>rs1</code> , store the sum in register <code>rd</code>
<code>ADD rd, rs1, rs2</code>	Sum two registers <code>rs1</code> and <code>rs2</code> , store the result in register <code>rd</code>
<code>LW rd, delta(rs1)</code>	Load 4 bytes to register <code>rd</code> from memory location specified by <code>rs1 + delta</code>
<code>SW rd, delta(rt)</code>	Store 4 bytes from register <code>rd</code> to memory location specified by <code>rt + delta</code>
<code>BLTU rs, rt, label</code>	Jump to <code>label</code> if register <code>rs < rt</code> ; both <code>rs</code> and <code>rt</code> are interpreted as unsigned numbers.

In addition, these registers are supported in URISA:

Register(s)	Description	X86 Equivalent
<code>zero</code>	Always reads zero; write ignored	-
<code>ra</code>	Return address of a function	On stack
<code>a1</code>	Return value of a function	<code>%eax</code>
<code>sp</code>	Stack pointer	<code>%esp</code>
<code>t0-6</code>	Temporary registers	Any free register, e.g., <code>%ecx</code> , <code>%edx</code>

Translate each of the following x86 instructions to the equivalent URISA instructions using only the instructions listed above. You may use any register above, but no new memory location other

than what is used in the original x86 instruction. Note that the x86 instructions use the AT&T/GAS syntax, i.e., `opcode src, dst`.

For example, the x86 instruction `pop (%eax)` would be equivalent to:

```
LW t0, 0(sp)
SW t0, 0(a1)
ADDI sp, sp, 4
```

(3 points extra credit) `add $0x8, %esp`

(3 points extra credit) `add %eax, 8(%esp)`

(3 points extra credit) `push $0x252`

(3 points extra credit) This processor does not have an Overflow flag. Fill in the following assembly so that if the ADD instruction causes an unsigned overflow, the program jumps to the overflow label.

```
ADD t3, t1, t2
_____ // Your answer here
;... code omitted
overflow:
;... code omitted
```

Problem 2: Floating-Point Arithmetics (15 points)

Part a) (3 points) Put $6\frac{5}{16}$ into the binary normalized form.

Part b) (12 points) The IEEE has decided to introduce a floating-point standard of some unknown size, whose main characteristics are consistent with existing floating-point number representations that we discussed in the class.

The following bit sequence contains the exact encoding of $6\frac{5}{16}$ in this new standard, but is padded with arbitrary bits at the beginning and in the end: 10010100110010110.

(3 points) How many bits are used for the fraction in this new standard?

(3 points) How many bits are used for the exponent in this new standard?

(3 points) What is the bias in this new standard?

(3 points) What is the smallest positive value that can be precisely expressed using this new floating-point format? You can write your answer either in the binary-normalized form or as a bit sequence.

Problem 3: Assembly Programming (16 points)

Conventions:

1. For this section, the assembly shown uses the AT&T/GAS syntax `opcode src, dst` for instructions with two arguments where `src` is the source argument and `dst` is the destination argument. For example, this means that `mov a, b` moves the value `a` into `b` and `cmp a, b` then `jge c` would compare `b` to `a` then jump to `c` if `b ≥ a`.
2. All C code is compiled on a 64-bit machine, where arrays grow toward higher addresses.
3. For functions that take two arguments, the first argument is stored in `%edi` and the second is stored in `%esi` at the time the function is called. The return value of this function is stored in `%eax` at the time the function returns.

Part a) (9 points) The following is assembly code for a C function `find_next()`:

```
00000000000001169 <find_next>:
    1169:    mov    %edi,%eax
    116b:    inc   %eax
    116d:    mov    $0x2,%edx
    1172:    cmp   %eax,%edx
    1174:    jge   1184 <find_next+0x1b>
    1176:    mov   %eax,%ecx
    1178:    sub   %edx,%ecx
    117a:    test  %ecx,%ecx
    117c:    jg    1178 <find_next+0xf>
    117e:    je    116b <find_next+0x2>
    1180:    inc   %edx
    1182:    jmp   1172 <find_next+0x9>
    1184:    ret
```

(3 points) What is the return value of `find_next(2)`?

(3 points) What is returned when a non-positive `x` is provided to `find_next(x)`? Show your answer in terms of `x`.

(3 points) For positive values of x as input, what does `find_next(x)` do?

Part b) (7 points) Some students wrote a CSC252 project in assembly. They discover right before submission that the machine used for grading has a faulty `call` instruction. They decide to rewrite the project without using `call`. The behavior of a `call` can be reproduced with **exactly two** other instructions.

(3 points) On a **correct 64-bit machine**, which register(s) does `call 1169` update? Explain how the register value(s) change.

(4 points) Which instructions would they use to replace the faulty `call`? Provide **just the names** of the instructions in the order they are used.

Problem 4: Cache (22 points)

You are designing a cache for an 8-bit machine that has a byte-addressable memory. Remember that an 8-bit machine means that the length of a memory address is 8 bits. The cache line size is 2 bytes. As a first attempt, you design a 4-way set associative cache with 2 sets and a random replacement policy. A random replacement policy chooses a line to replace at random within the set. To reduce memory traffic, you design it to be a write-back cache.

Assuming that the machine runs only one program, which will generate the following memory access sequence.

Access	Address
1	00110000
2	10100010
3	10011001
4	01011101
5	00110111
6	10011010
7	01011110
8	01101100
9	10100010

(6 points) How many bits do you need for the tag, the set index, and the cache line offset?

(3 points) How many overhead bits does this cache have? Recall that overhead includes the tag bits, the dirty bits (if needed), and the replacement bits (if needed). Show your work to earn full credit.

(3 points) Assuming the cache is initially empty. How many cache misses will the program generate when it's run for the first time?

(3 points) How many cache misses will the program generate when it's run for the second time?

(4 points) Looking carefully at the access pattern, you realize that it is possible to design a cache with a smaller overhead that achieves a 100% hit rate for the program after the initial run.

Without changing the cache line size, explain the new design of your cache, i.e., its associativity and the number of sets.

(3 points) What is the overhead of your new cache design?

Problem 5: Virtual Memory (30 points + 3 points extra credit)

You are building a machine with virtual memory support. The virtual address space is 256 KB (2^{18} bytes). The physical memory size is 32KB. The system uses a one-level page table.

Part a) Basic organization (13 points + 3 points extra credit)

(4 points) Assume the page size is 1KB. How many bits are used for the physical page number (PPN), and how many for the virtual page number (VPN)?

(3 points) What is the size of one PTE, assuming that the only overhead in the PTE is one valid bit and that the disk address is 2 bits longer than the PPN.

(3 points) You decide to experiment with changing the page size to 4KB. What is the size of one PTE now? Continue to assume that the only overhead in the PTE is one valid bit and that the disk address is 2 bits longer than the PPN.

(3 points) The system has a TLB. What is likely to happen to the TLB hit rate when we increase the page size to 4KB? Will it increase or decrease? Explain your answer to receive ANY credit.

(3 points extra credit) If the machine were to use a two-level page table with a 4KB page size, how many PTE entries are there in the level 1 page table?

Part b) Protection (8 points)

The major customer for your system is the US Government. They want a secure system that is safer from buffer overflow attacks. The memory protection system should help protect against arbitrary code execution in the event that someone manages to exploit a buffer overflow.

(4 points) Briefly explain the mechanism by which an attacker exploits a buffer overflow. Your answer must cover 1) where in memory the attacker can place the code and 2) how the injected code can be executed.

(4 points) Now you want to augment the PTE structure with permission information to defend against buffer overflow attacks. In order to keep the overhead low, you decide that you will use a maximum of 2 permission bits, which will indicate whether a page can be written to, whether a page can be read from, and whether a page's contents can be executed as code. Assuming defending buffer overflow attacks is your sole goal in this design process, how will you use the 2 permission bits? Explain the semantics of each bit to receive full credit.

Part c) Performance (9 points)

Management decided to go with the 1KB page size for this machine and removed the TLB to cut costs. Consider the C code below. Assume:

1. A 1-level page table translation scheme with the page table starting empty, and
2. `arr` is aligned such that it starts 8 bytes before a page boundary.

```
void func() {
    int arr[2048];

    for(int i = 0; i < 8; i++) {

        if(rand_0_1() == 1) {
```

```
        printf("%d\n", arr[(i << 8) + 1]);
        printf("%d\n", arr[(i << 8) + 2]);
    } else {
        printf("%d\n", arr[i]);
        printf("%d\n", arr[i + 2]);
    }
}
}
```

`rand_0_1()` is a function that returns either 0 or 1, at random.

For the following 3 questions, only count page faults generated by accessing elements of `arr`.

(3 points) What is the maximum number of page faults that this code could possibly generate?

(3 points) What is the minimum number of page faults that this code could possibly generate?

(3 points) Suppose that the `rand_0_1()` function is faulty and it returns 0 40% of the time and 1 60% of the time. How many page faults, on average, will this code generate?