

Midterm Exam

601.229 Computer Systems Fundamentals

October 1, 2021

Complete all questions.

Use additional paper if needed.

Time: 50 minutes.

I affirm that I have completed this exam without unauthorized assistance from any person, materials, or device.

Signed: Solution

Print name: _____

Date: _____

Reference

Powers of 2 ($y = 2^x$):

<i>x</i>	0	1	2	3	4	5	6	7	8	9	10	11	12
<i>y</i>	1	2	4	8	16	32	64	128	256	512	1,024	2,048	4,096

<i>x</i>	13	14	15	16
<i>y</i>	8,192	16,384	32,768	65,536

Note that in all questions concerning C:

- `uint8_t` is an 8-bit unsigned integer type
- `uint16_t` is a 16-bit unsigned integer type
- `uint32_t` is a 32-bit unsigned integer type
- `int8_t` is an 8-bit signed two's complement integer type
- `int16_t` is a 16-bit signed two's complement integer type
- `int32_t` is a 32-bit signed two's complement integer type

x86-64 registers:

Callee-saved: `%rbx, %rbp, %r12, %r13, %r14, %r15`

Caller-saved: `%r10, %r11`

Return value: `%rax`

Arguments: `%rdi, %rsi, %rdx, %rcx, %r8, %r9`

Note that argument registers and return value register are effectively caller-saved.

Registers and sub-registers:

Register	Low 32 bits	Low 16 bits	Low 8 bits
<code>%rax</code>	<code>%eax</code>	<code>%ax</code>	<code>%al</code>
<code>%rbx</code>	<code>%ebx</code>	<code>%bx</code>	<code>%bl</code>
<code>%rcx</code>	<code>%ecx</code>	<code>%cx</code>	<code>%cl</code>
<code>%rdx</code>	<code>%edx</code>	<code>%dx</code>	<code>%dl</code>
<code>%rbp</code>	<code>%ebp</code>	<code>%bp</code>	<code>%bpl</code>
<code>%rsi</code>	<code>%esi</code>	<code>%si</code>	<code>%sil</code>
<code>%rdi</code>	<code>%edi</code>	<code>%di</code>	<code>%dil</code>
<code>%r8</code>	<code>%r8d</code>	<code>%r8w</code>	<code>%r8b</code>
<code>%r9</code>	<code>%r9d</code>	<code>%r9w</code>	<code>%r9b</code>
<code>%r10</code>	<code>%r10d</code>	<code>%r10w</code>	<code>%r10b</code>
<code>%r11</code>	<code>%r11d</code>	<code>%r11w</code>	<code>%r11b</code>
<code>%r12</code>	<code>%r12d</code>	<code>%r12w</code>	<code>%r12b</code>
<code>%r13</code>	<code>%r13d</code>	<code>%r13w</code>	<code>%r13b</code>
<code>%r14</code>	<code>%r14d</code>	<code>%r14w</code>	<code>%r14b</code>
<code>%r15</code>	<code>%r15d</code>	<code>%r15w</code>	<code>%r15b</code>

Stack alignment: `%rsp` must contain an address that is a multiple of 16 when any `call` instruction is executed.

Operand size suffixes: **b** = 1 byte, **w** = 2 bytes, **l** = 4 bytes, **q** = 8 bytes (Examples: `movb`, `movw`, `movl`, `movq`)

Question 1. [10 points] Show the binary (base 2) representation of the following integer values:

- $15 = 8 + 4 + 2 + 1 = \boxed{1111}$
- $225 = 128 + 64 + 32 + 1 = \boxed{11100001}$

Question 2. [10 points] What output is printed by the following C code? Explain briefly.

$$\begin{array}{r} 128 \\ 107 \\ \hline 21 \end{array}$$

```
uint8_t a = 197;
uint8_t b = 65;
uint8_t sum = a + b;
printf("%u\n", (unsigned) sum);
```

$$\begin{array}{r} 197 \\ + 65 \\ \hline 262 \end{array} \text{ mod } 256 = 6$$

output is $\boxed{6}$

Question 3. [10 points] Show the 8-bit two's complement representation of the following integer values:

- $51 = 32 + 16 + 2 + 1 \quad \boxed{00110011}$
- $-107 = -128 + 21$
 $= -128 + 16 + 4 + 1 \quad \boxed{10010101}$

Question 4. [10 points] What output is printed by the following code? (Hint: $|$ means bitwise or.) Explain briefly.

```
int16_t x = 32767;  $\rightarrow$  01111111 11111111
printf("%d\n", x);
x = x |  $\boxed{0x8000}$ ; sets sign bit to 1 (adding  $-2^{15} = -32768$ )
printf("%d\n", x);
```

output is $\boxed{-1}$

Question 5. [10 points] A 32-bit IEEE 754 single precision floating point value has the following representation:

Sign	Exponent	Fraction
1 bit	8 bits	23 bits

Recall that normalized floating point numbers have values $\pm 1.x \times 2^y$, where x is specified by the fraction bits, and y is value of the exponent (which has a value between -126 and 127 .)

This format allows all integer values in the range $-q$ to q (inclusive) to be represented exactly.

State the value of $q = 1.x \times 2^y$. First, specify the fraction (x) in base 2 (i.e., a sequence of 23 0s and 1s):

$x =$

Next, specify the exponent (y) in base 10:

$$q = 1.0 \times 2^{24}$$

$y =$

Optional: explain briefly.

an exponent of 23 allows the decimal point to be moved to the right past all digits of the fraction. So, all integers with magnitudes 2^{24} and less can be represented exactly.

Question 6. [10 points] What output is printed by the following C program? Assume that `sizeof(int) == 4`. Explain briefly.

```
int a[4] = { 6, 7, 8, 9 };
printf("%d\n", (int) (&a[2] - &a[0])); 
printf("%d\n", (int) ((char *)&a[2] - (char *)&a[0])); 
```

Arithmetic on C pointers is in terms of number of array elements, regardless of element size. Because char elements are 1 byte, arithmetic on char pointers is in terms of bytes. Elements 0 and 2 of an array of int elements are 8 bytes apart.

Question 7. [40 points] Consider the following C function prototype:

```
void add_to_vec_if_even(int32_t *vec, unsigned len, int32_t value);
```

This function takes an array of `len` `int32_t` values and adds `value` to each of the *even* values in the array. Its behavior is described by the following unit test:

```
int32_t data[] = { 247, -550, 582, 181 };
add_to_vec_if_even(data, 4, 10);
ASSERT(247 == data[0]); // original value was odd
ASSERT(-540 == data[1]); // original value was even
ASSERT(592 == data[2]); // original value was even
ASSERT(181 == data[3]); // original value was odd
```

Show an x86-64 assembly language implementation of the `add_to_vec_if_even` function. (Continue on next page if necessary.) **Hint:** `andl $1, %reg` is a useful way to check whether the 32-bit value in `Reg` is odd.

```
.globl add_to_vec_if_even:
add_to_vec_if_even:
    subq $8, %rsp // not really necessary
    leaq (%rdi, %rsi, 4), %r10 // %r10 points one element past
                                // end of array

.Lloop:
    cmpq %r10, %rdi // reached end?
    jae .Ldone // if so, done

    movl (%rdi), %r11d // get array element
    andl $1, %r11d // see if low bit is set
    jne .Ladvance // if it is, don't modify (b/c value is odd)
    addl %edx, (%rdi) // add value to array element

.Ladvance:
    addq $4, %rdi // advance to next element
    jmp .Lloop // continue loop

.Ldone:
    addq $8, %rsp
    ret
```

[Continue your answer to Question 7 here if necessary.]