

# Midterm Exam 3

## 601.229 Computer Systems Fundamentals

December 16, 2021

Complete all questions.

Time: 90 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 ( $y = 2^x$ ):

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

$x$	13	14	15	16
$y$	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.** [20 points] On the 32-bit x86 architecture, the page size is  $2^{12} = 4096$  bytes, and there are two levels of page tables. Each page table (both the root page table and the second-level page tables) have 1024 page table entries. Virtual addresses are 32 bits, and each virtual page in the entire 32-bit address space can be mapped to a physical page. Assume the bits in a virtual address are numbered 0–31, with 0 being the least significant bit, and 31 being the most significant bit.

(a) Which bits of a virtual address are the *page offset*?

0–11 inclusive

(b) Which bits of a virtual address are used as the index in the “root” (level 1) page table, in order to find the page table entry leading to the level 2 page table?

22–31 inclusive

(c) Which bits of a virtual address are the index in the level 2 page table, in order to find the page table entry leading to the mapped physical page?

12–21 inclusive

(d) How many virtual pages are there in the overall address space? You may express this as a power of 2 or sum of powers of 2.

$$2^{32} \text{ bytes in addr space} \times \frac{1 \text{ page}}{2^{12} \text{ bytes}} = 2^{20} \text{ pages}$$

(e) Assume that a virtual address space maps every virtual page to a corresponding physical page. How many page tables (at both levels) are needed? You may express this as a power of 2 or sum of powers of 2.

$$\begin{array}{r} 1 \text{ level 1 page table} \\ + 1024 \text{ level 2 page tables} \\ \hline 1025 \text{ page tables} \end{array}$$

**Question 2.** [10 points] Consider the following server loop which uses processes to allow concurrent client connections:

```
1:  while (1) {
2:      int clientfd = accept(serverfd, NULL, NULL);
3:      pid_t pid = fork();
4:      if (pid == 0) {
5:          chat_with_client(clientfd);
6:          exit(0);
7:      }
8:      close(clientfd);
9:  }
```

Briefly explain the reason why the call to close is needed at line 8.

After the fork, both parent and child processes have the client socket open. (The child process duplicates the parent's file descriptors.) The parent process must close the file descriptor, or else the TCP socket will stay open, even though the client has finished.

**Question 3.** [10 points] Consider the following function, which is meant to write the contents of a buffer in memory to a file descriptor:

```
// Returns 1 if successful, 0 if unsuccessful
int send_data(const void *buf, unsigned num_bytes, int fd) {
    ssize_t bytes_sent = write(fd, buf, num_bytes);
    if (bytes_sent >= 0 && (unsigned)bytes_sent == num_bytes) {
        return 1;
    } else {
        return 0;
    }
}
```

Briefly explain the most important flaw in this function, and how to fix it. (You don't need to show code for a fixed version.) Hint: consider that fd might refer to a TCP socket.

The call to write might not be able to send all of the bytes requested, i.e., a short write could occur. The function should call write in a loop until all bytes are written.

**Question 4.** [20 points] Consider the following partially-specified multithreaded server implementation (note that error handling is omitted, and assume appropriate system headers are #included):

```
void chat_with_client(int fd); // defined elsewhere

struct ConnInfo {
    HERE 1
};

void *worker(void *arg) {
    HERE 2
}

int main(int argc, char **argv) {
    int serverfd = Open_listenfd(argv[1]);
    while (1) {
        int clientfd = accept(serverfd, NULL, NULL);
        struct ConnInfo *info = malloc(sizeof(struct ConnInfo));
        HERE 3
        pthread_t thr;
        pthread_create(&thr, NULL, worker, HERE 4);
    }
}
```

Indicate what code should be substituted for the missing code labeled `HERE 1`, `HERE 2`, `HERE 3`, and `HERE 4`. Assume that the `chat_with_client` function implements sending data to and receiving data from the remote client.

HERE 1

int fd;

HERE 2

struct ConnInfo \*info = arg;  
pthread\_detach(  
pthread\_self());

HERE 3

info->fd = clientfd;

HERE 4

info

chat\_with\_client(info->fd);  
close(info->fd);  
free(info);  
return NULL;

**Question 5.** [15 points] Assume that the message format for a network protocol is defined as follows. Each message is a single line of text terminated by a newline ('\n') character. The content of a line is a *code* specified by an upper case letter ('A' through 'Z'), immediately followed by an integer *value* specified as 1 and 6 digit characters ('0' through '9').

Examples of messages:      Struct data type to represent a message:

```
Q9                            struct Message {
B55                            char code;
Y90125                        int value;
                              };
```

Implement the following `recv_msg` function so that it reads a single message from the specified file descriptor and uses the received data to fill in the contents of the `struct Message` instance pointed-to by the parameter `p`.

Hints and specifications:

- The function should read *only* data that is part of *one* message
- It will probably be easiest to read one character at a time
- The function prototype for the read system call is  
`int read(int fd, void *buf, size_t n);`
- The function prototype for the `atoi` function (to convert a NUL-terminated string of digits to an `int` value) is  
`int atoi(const char *str);`
- You may use the `isalpha` and/or `isdigit` functions
- You may assume that the received data is properly formatted *and that no errors will occur*

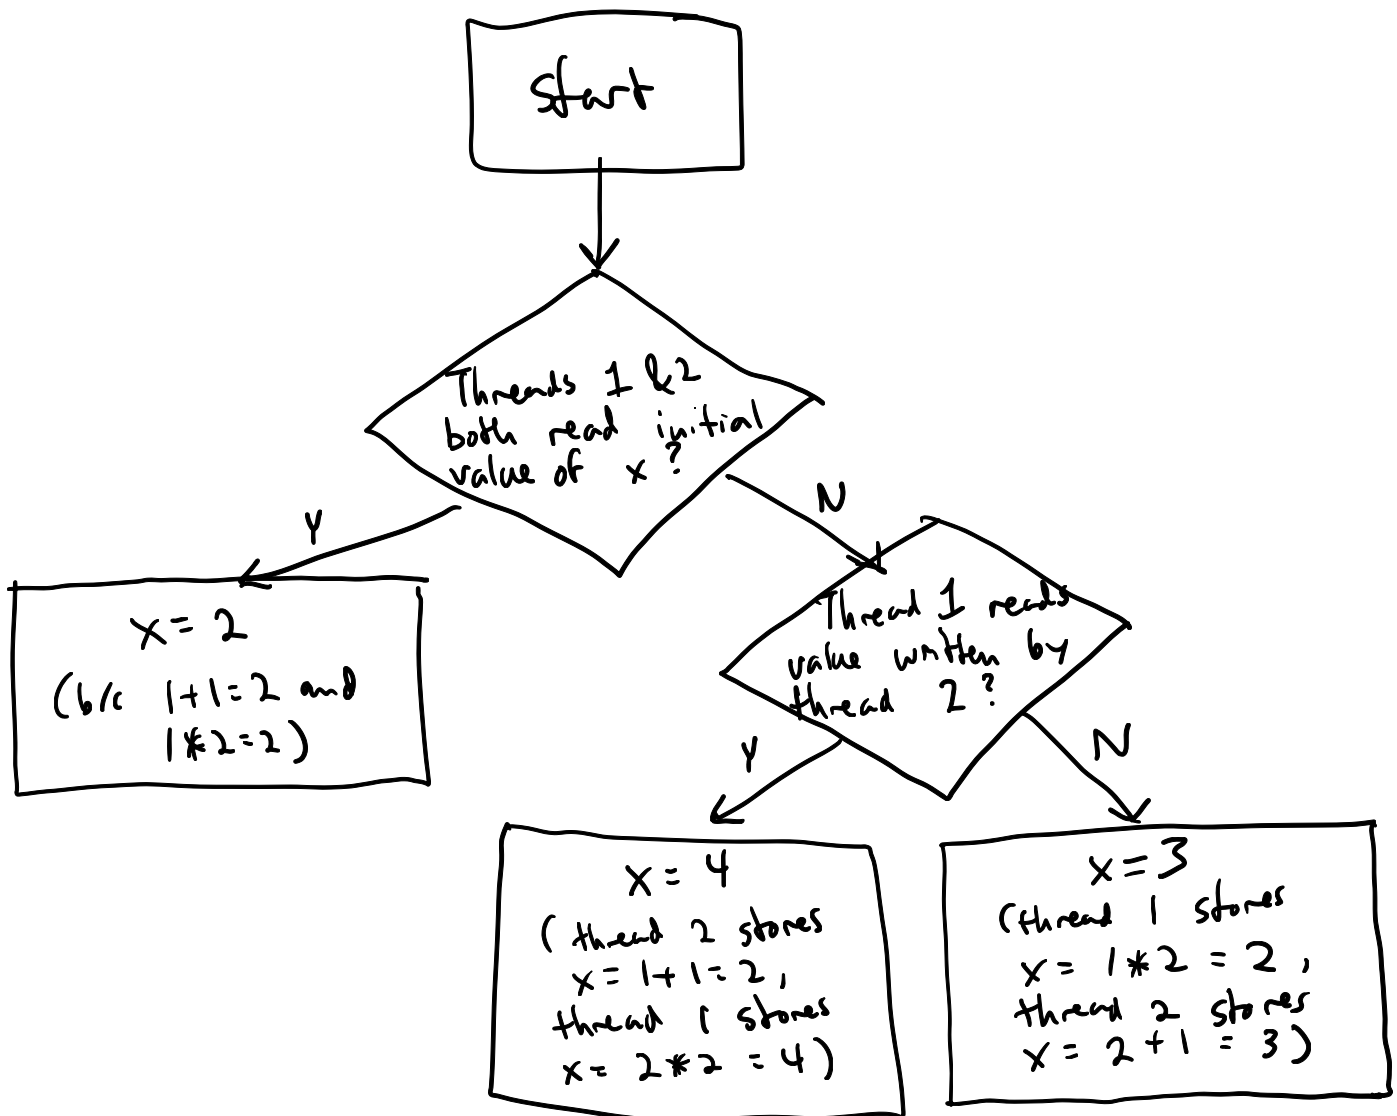
```
void recv_msg(int fd, struct Message *p) {
    char buf;
    read(fd, &buf, 1);
    p->code = buf;
    p->value = 0;
    read(fd, &buf, 1);
    while (buf != '\n') { p->value *= 10;
                        p->value += (buf - '0');
                        read(fd, &buf, 1); }
}
```

**Question 6.** [15 points] Consider the following operations performed by two different threads without synchronization:

```
// Thread 1  
foo = x  
foo = foo * 2  
x = foo
```

```
// Thread 2  
bar = x  
bar = bar + 1  
x = bar
```

Assume that  $x$  is a shared variable accessible by both threads, and that its initial value (before either thread starts) is 1. What possible final value(s) could  $x$  have after both threads finish their operations? Explain briefly.



Question 7. [15 points] Consider the following C data type and functions:

```
struct SharedVec3 {  
  
    float data[3];  
    pthread_mutex_t lock;  
};  
  
void svec3_init(struct SharedVec3 *sv) {  
  
    for (int i = 0; i < 3; i++) { sv->data[i] = 0.0f; }  
    pthread_mutex_init(&sv->lock, NULL);  
}  
  
void svec3_addto(struct SharedVec3 *sv, int index, float val) {  
    pthread_mutex_lock(&sv->lock);  
    sv->data[index] += val;  
    pthread_mutex_unlock(&sv->lock);  
}  
  
float svec3_get(struct SharedVec3 *sv, int index) {  
    pthread_mutex_lock(&sv->lock);  
    float val = return sv->data[index];  
    pthread_mutex_unlock(&sv->lock);  
    return val;  
}
```

Show how to add synchronization to the data type and functions so that it is safe for concurrent use by multiple threads. Indicate your changes above.



[Extra page for answers and/or scratch work.]

[Extra page for answers and/or scratch work.]

[Extra page for answers and/or scratch work.]