

# Section 4: Synchronization, Locks

CS 162

July 6, 2020

## Contents

<b>1</b>	<b>Vocabulary</b>	<b>2</b>
<b>2</b>	<b>Synchronization</b>	<b>3</b>
2.1	Locking via Disabling Interrupts . . . . .	3
2.2	Counting Semaphores . . . . .	3
2.3	The Central Galactic Floopy Corporation . . . . .	4
2.4	Crowded Video Games . . . . .	6
<b>3</b>	<b>Additional Questions</b>	<b>7</b>
3.1	Signals . . . . .	7
3.2	Signal Handlers . . . . .	7
3.3	Exec . . . . .	8
3.4	Exec + Fork . . . . .	8
3.5	RPC Consistency . . . . .	8
3.6	Socket Choices . . . . .	9

# 1 Vocabulary

- **Lock** - Synchronization variables that provide mutual exclusion. Threads may acquire or release a lock. Only one thread may hold a lock at a time. If a thread attempts to acquire a lock that is held by some other thread, it will block at that line of code until the lock is released and it successfully acquires it. Implementations can vary.
- **test\_and\_set** - An atomic operation implemented in hardware. Often used to implement locks and other synchronization primitives. In this handout, assume the following implementation.

```
int test_and_set(int *value) {
    int result = *value;
    *value = 1;
    return result;
}
```

This is more expensive than most other instructions, and it is not preferable to repeatedly execute this instruction.

- **semaphore** - Synchronization primitives that are used to control access to a shared variable in a more general way than locks. A semaphore is simply an integer with restrictions on how it can be modified:
  - When a semaphore is initialized, the integer is set to a specified starting value.
  - A thread can call **down()** (also known as **P**) to attempt to decrement the integer. If the integer is zero, the thread will block until it is positive, and then unblock and decrement the integer.
  - A thread can call **up()** (also known as **V**) to increment the integer, which will always succeed.

Unlike locks, semaphores have no concept of "ownership", and any thread can call **down()** or **up()** on any semaphore at any time.

- **Race Condition** - A state of execution that causes multiple threads to access the same shared variable (heap or global data segment) with at least one thread attempting to execute a write without enforcing mutual exclusion. The result is not necessarily garbage but is treated as being undefined since there is no guarantee as to what will actually happen. Note that multiple reads do not need to be mutexed.
- **Critical section** - A section of code that accesses a shared resource and must not be concurrently run by more than a single thread.
- **TCP** - Transmission Control Protocol (TCP) is a common L4 (transport layer) protocol that guarantees reliable in-order delivery. In-order delivery is accomplished through the use of sequence numbers attached to every data packet, and reliable delivery is accomplished through the use of ACKs (acknowledgements).
- **Device Driver** - Device-specific code in the kernel that interacts directly with the device hardware. They support a standard, internal interface so the same kernel I/O system can interact easily with different hardware. The top half of a device driver is used by the kernel to start I/O operations. The bottom half of a device driver services interrupts produced by the device. You should know that Linux has different definitions for "top half" and "bottom half", which are essentially the reverse of these definitions (top half in Linux is the interrupt service routine, whereas the bottom half is the kernel-level bookkeeping).

## 2 Synchronization

### 2.1 Locking via Disabling Interrupts

Consider the following implementation of Locks:

```
Lock::Acquire() {                               Lock::Release() {
    disable_interrupts();                         enable_interrupts();
}                                                  }
```

1. For a single-processor system state whether this implementation is incorrect.

- Does not work for multiple locks
- Process holding lock may not release for a long time, effectively halting the machine. Can deadlock if program holds the lock and waits for some I/O, which requires an interrupt.
- Does not maintain the "acquired" state of the lock across a context switch (`yield()`).

2. For a multiprocessor system, explain what additional reason(s) might make this implementation incorrect?

Does not block other processors from accessing the critical section.

### 2.2 Counting Semaphores

We have a limited number of resources such that only N threads can use them at any given time. Explain how we can use a "counting semaphore" to control access to these resources.

- Set the initial value of the semaphore to N (number of concurrent accesses allowed).
- P() decrements the semaphore's counter and either causes the process to wait until the resource is available or allocates the process the resource.
- V() increments the semaphore's counter, releasing a waiting process (if any is waiting).

## 2.3 The Central Galactic Floopy Corporation

It's the year 3162. Floopies are the widely recognized galactic currency. Floopies are represented in digital form only, at the Central Galactic Floopy Corporation (CGFC).

You receive some inside intel from the CGFC that they have a GalaxyNet server running on some old OS called x86 Ubuntu 14.04 LTS. Anyone can send requests to it. Upon receiving a request, the server forks a POSIX thread to handle the request. In particular, you are told that sending a transfer request will create a thread that will run the following function immediately, for speedy service.

```
void transfer(account_t *donor, account_t *recipient, float amount) {
    assert (donor != recipient); // Thanks CS161

    if (donor->balance < amount) {
        printf("Insufficient funds.\n");
        return;
    }
    donor->balance -= amount;
    recipient->balance += amount;
}
```

Assume that there is some struct with a member `balance` that is typedef-ed as `account_t`. Describe how a malicious user might exploit some unintended behavior.

There are multiple race conditions here.

Suppose Alice and Bob have 5 floopies each. We send two quick requests: `transfer(&alice, &bob, 5)` and `transfer(&bob, &alice, 5)`. The first call decrements Alice's balance to 0, adds 5 to Bob's balance, but before storing 10 in Bob's balance, the next call comes in and executes to completion, decrementing Bob's balance to 0 and making Alice's balance 5. Finally we return to the first call, which just has to store 10 into Bob's balance. In the end, Alice has 5, but Bob now has 10. We have effectively duplicated 5 floopies.

Graphically:

### Thread 1

```
temp1 = Alice's balance (== 5)
temp1 = temp1 - 5 (== 0)
Alice's balance = temp1 (== 0)
temp1 = Bob's balance (== 5)
temp1 = temp1 + 5 (== 10)
INTERRUPTED BY THREAD 2
```

### RESUME THREAD 1

```
Bob's balance = temp1 (== 10)
THREAD 1 COMPLETE
```

### Thread 2

```
temp2 = Bob's balance (== 5)
temp2 = temp2 - 5 (== 0)
Bob's balance = temp2 (== 0)
temp2 = Alice's balance (== 0)
temp2 = temp2 + 5 (== 5)
Alice's balance = temp2 (== 5)
THREAD 2 COMPLETE
```

It is also possible to achieve a negative balance. Suppose at the beginning of the function, the donor has enough money to participate in the transfer, so we pass the conditional check for sufficient funds. Immediately after that, the donor's balance is reduced below the required amount by some

other running thread. Then the transfer will go through, resulting in a negative balance for the donor.

Sending two identical `transfer(&alice, &bob, 2)` may also cause unintended behavior, since the increment/decrement operations are not atomic (though it is arguably harder to exploit for profit).

Since you're a good person who wouldn't steal floopies from a galactic corporation, what changes would you suggest to the CGFC to defend against this exploit?

The entire function must be made atomic. One could do this by disabling interrupts for that period of time (if there is a single processor), or by acquiring a lock beforehand and releasing the lock afterwards. Alternatively, you could have a lock for each account. In order to prevent deadlocks, you will have to acquire locks in some predetermined order, such as lowest account number first.

## 2.4 Crowded Video Games

A recent popular game is having issues with its servers lagging heavily due to too many players being connected at a time. Below is the code that a player runs to play on a server:

```
void play_session(struct server s) {
    connect(s);
    play();
    disconnect(s);
}
```

After testing, it turns out that the servers can run without lagging for a max of up to 1000 players concurrently connected.

How can you add semaphores to the above code to enforce a strict limit of 1000 players connected at a time? Assume that a game server can create semaphores and share them amongst the player threads.

Introduce a semaphore for each server, initialized to 1000, to control the ability to connect to the game. A player will `down()` the semaphore **before** connecting, and `up()` the semaphore **after** disconnecting.

The order here is important - downing the semaphore after connecting but before playing means that there is no block on the `connect()` call, and upping the semaphore before disconnecting could lead to "zombie" players, who were pre-empted before disconnecting. Both of these cases mean that the limit of 1000 could be violated.

### 3 Additional Questions

These are questions we may not get to in discussion but you may find useful when studying for the exam.

#### 3.1 Signals

The following is a list of standard Linux signals:

Signal	Value	Action	Comment
SIGHUP	1	Terminate	Hangup detected on controlling terminal or death of controlling process
SIGINT	2	Terminate	Interrupt from keyboard (Ctrl - c)
SIGQUIT	3	Core Dump	Quit from keyboard (Ctrl - \)
SIGILL	4	Core Dump	Illegal Instruction
SIGABRT	6	Core Dump	Abort signal from abort(3)
SIGFPE	8	Core Dump	Floating point exception
SIGKILL	9	Terminate	Kill signal
SIGSEGV	11	Core Dump	Invalid memory reference
SIGPIPE	13	Terminate	Broken pipe: write to pipe with no readers
SIGALRM	14	Terminate	Timer signal from alarm(2)
SIGTERM	15	Terminate	Termination signal
SIGUSR1	30,10,16	Terminate	User-defined signal 1
SIGUSR2	31,12,17	Terminate	User-defined signal 2
SIGCHLD	20,17,18	Ignore	Child stopped or terminated
SIGCONT	19,18,25	Continue	Continue if stopped
SIGSTOP	17,19,23	Stop	Stop process
SIGTSTP	18,20,24	Stop	Stop typed at tty
SIGTTIN	21,21,26	Stop	tty input for background process
SIGTTOU	22,22,27	Stop	tty output for background process

#### 3.2 Signal Handlers

Fill in the blanks for the following function using syscalls such that when we type Ctrl-C, the user is prompted with a message: “Do you really want to quit [y/n]? ”, and if “y” is typed, the program quits. Otherwise, it continues along.

```

void sigint_handler(int sig)
{
    char c;
    printf("\0uch, you just hit Ctrl-C?. Do you really want to quit [y/n]?");
    c = getchar();
    if (c == "y" || c = "Y")
        exit(0);
}

int main() {
    signal(SIGINT, sigint_handler);
    ...
}

```

```
}

```

### 3.3 Exec

What will C print?

```
int main(void)
{
    char** argv = (char**) malloc(3*sizeof(char*));
    argv[0] = "/bin/ls";
    argv[1] = ".";
    argv[2] = NULL;
    for (int i = 0; i < 10; i++) {
        printf("%d\n", i);
        if (i == 3)
            execv("/bin/ls", argv);
    }
}
```

```
0
1
2
3
<output of ls>
```

### 3.4 Exec + Fork

How would I modify the above program using fork so it both prints the output of `ls` and all the numbers from 0 to 9 (order does not matter)? You may not remove lines from the original program; only add statements (and use `fork!`).

```
int main(void)
{
    char** argv = (char**) malloc(3*sizeof(char*));
    argv[0] = "/bin/ls";
    argv[1] = ".";
    argv[2] = NULL;
    for (int i = 0; i < 10; i++){
        printf("%d\n", i);
        if (i == 3) {
            pid_t pid = fork();
            if (pid == 0)
                execv("/bin/ls", argv);
        }
    }
}
```

### 3.5 RPC Consistency

In our RPC example (Section 3) we opted to use the type `uint32_t`. What is a potentially issue that could result if we opted to use an `unsigned int` instead? If a value can only take either 0 or 1 is could

your previous answer still be an issue?

`uint32_t` is defined to be exactly 4 bytes whereas an `unsigned int` has a machine specific size. This is a problem because our server and client may not agree on the size which could cause data to be incorrectly unmarshaled. This is still an issue if the value are only 0 and 1 because size issues could incorrectly cause us to read part of another argument for example or reproduce us waiting on data like in Section 3 4.2.

### 3.6 Socket Choices

RPC calls are implemented with socket implementations. In lecture you saw running a server with 1 thread and last week in discussion you looked at sockets with multiple threads. Below is the code for running a server where each request is its own process.

```
struct addrinfo *setup_address (char *port) {
    struct addrinfo *server;
    struct addrinfo hints;
    memset (&hints, 0, sizeof(hints));
    hints.ai_family = AF_UNSPEC;
    hints.ai_socktype = SOCK_STREAM;
    hints.ai_flags = AI_PASSIVE;

    if (getaddrinfo(NULL, port, &hints, &server) != 0) {
        perror ("Address");
        return NULL;
    }
    return server;
}
```

```

int setup_server_socket (struct addrinfo *server) {
    // Setup the addr socket
    bool connected = false;
    int sock = -1;
    for (struct addrinfo *addr = server; addr != NULL && !connected; addr = addr->ai_next) {
        int sock = socket (addr->ai_family, addr->ai_socktype, addr->ai_protocol);
        if (sock != -1) {
            if (bind(sock, addr->ai_addr, addr->ai_addrlen) == -1)
                close (sock);
                sock = -1;
            } else {
                return sock;
            }
        }
    }
    return sock;
}

void run_server (char *port) {
    struct addrinfo *server = setup_address(port);
    int server_fd = setup_server_socket (server);
    if (server_fd == -1) {
        perror ("socket");
    } else {
        int rv = listen (server_fd, BACKLOG);
        if (rv < 0) {
            perror ("socket");
        } else {
            while (1) {
                int connection_socket = accept (server_fd, NULL, NULL);
                int pid = fork ();
                if (pid < 0) {
                    perror ("fork");
                } else if (pid == 0) {
                    receive_rpc (connection_socket);
                    exit (0);
                } else {
                    close (connection_socket);
                }
            }
        }
    }
}

```

What's a possible advantage of playing each RPC request in a separate process instead of a separate thread?

The biggest advantage for us here is safety. Because we are exposing function calls it is possible we could be fed an input that doesn't meet our assumptions which if not properly handled could cause a segfault. If we have this code in a separate thread it will crash the process while if its in a separate process it won't crash the whole server.