

# Review Session: Threads, Synchronization, and File I/O

CS 162

September 25, 2019

## Contents

<b>1</b>	<b>Threads, Processes</b>	<b>2</b>
1.1	Process vs Thread . . . . .	2
1.2	Context Switch . . . . .	2
1.3	Overuse of Threads . . . . .	2
1.4	Critical Section . . . . .	2
1.5	Mode Switch . . . . .	2
1.6	Threads . . . . .	3
<b>2</b>	<b>Synchronization</b>	<b>3</b>
2.1	Locking via Disabling Interrupts . . . . .	3
2.2	Counting Semaphores . . . . .	4
2.3	Bounded Buffer with Locks and Semaphores . . . . .	4
<b>3</b>	<b>File I/O</b>	<b>6</b>
3.1	Benefits of Concurrency . . . . .	6
3.2	Storing Integers . . . . .	6

# 1 Threads, Processes

## 1.1 Process vs Thread

How is a process different from a thread?

## 1.2 Context Switch

What state information do you need to save/restore about threads when performing a context switch?

## 1.3 Overuse of Threads

List two reasons why overuse of threads is bad (i.e., using too many threads for different tasks). Be explicit in your answers.

## 1.4 Critical Section

Suppose a thread is running in a critical section of code, meaning that it has acquired the locks through proper arbitration. Can it get context switched? Why or why not?

## 1.5 Mode Switch

Name three ways in which the processor can transition from user mode to kernel mode. Can the user execute arbitrary code after transitioning?

## 1.6 Threads

What are the possible outputs of this program?

```
void* threadfun(void* arg) {
    print("hello\n");
}

int main(int argc, char** argv) {
    pthread_t t;
    pthread_create(&t, NULL, threadfun, NULL);
    sleep(1);
    printf("world\n");
}
```

## 2 Synchronization

### 2.1 Locking via Disabling Interrupts

Consider the following implementation of Locks:

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

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

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

## 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.

## 2.3 Bounded Buffer with Locks and Semaphores

In class we discussed a solution to the bounded-buffer problem for multiple producers and multiple consumers using semaphores and locks. In this problem, there is a fixed size buffer. Producers add to the buffer, but only if there is room in the buffer. Consumers remove from the buffer, but only if there are items. Producers should sleep until there is space in the buffer and consumers should sleep until there is something in the buffer.

1. Reproduce this solution from lecture. Think about how we can use semaphores track how many empty and full spots are currently in the buffer. Do not go onto the next part until you are done with this, as it will give away the answer.

<pre> Producer () {     -----     -----     queue.Enqueue()     -----     ----- }                 </pre>	<pre> Consumer () {     -----     -----     queue.Dequeue()     -----     ----- }                 </pre>
--	--

2. Examine the following implementation where the lock operations have been moved. Explain it is correct. If not, explain why.

```

Producer () {
    mutex.Acquire()
    emptyBuffers.P()
    queue.Enqueue()
    fullBuffers.V()
    mutex.Release()
}

Consumer () {
    mutex.Acquire()
    fullBuffers.P()
    queue.Dequeue()
    emptyBuffers.V()
    mutex.Release()
}

```

3. Examine the following implementation where the lock operations have been moved. Explain it is correct. If not, explain why.

```

Producer () {
    mutex.Acquire()
    emptyBuffers.P()
    queue.Enqueue()
    fullBuffers.V()
    mutex.Release()
}

Consumer () {
    fullBuffers.P()
    mutex.Acquire()
    queue.Dequeue()
    mutex.Release()
    emptyBuffers.V()
}

```

4. Examine the following implementation where the lock operations have been moved. Explain it is correct. If not, explain why.

```

Producer () {
    emptyBuffers.P()
    mutex.Acquire()
    queue.Enqueue()
    fullBuffers.V()
    mutex.Release()
}

Consumer () {
    fullBuffers.P()
    mutex.Acquire()
    queue.Dequeue()
    emptyBuffers.V()
    mutex.Release()
}

```

### 3 File I/O

#### 3.1 Benefits of Concurrency

On a single processor machine, under what circumstances can a multithreaded program complete more quickly than a non-multithreaded program? Keep in mind that multithreading has context-switch overhead associated with it.

#### 3.2 Storing Integers

You are working for BigStore and your boss has tasked you with writing a function that takes an array of ints and writes it to a specified file for later use. He also wants you to use file descriptors (no `fopen`, etc.). Fill in the following function:

```
void write_to_file(const char *file, int *a, int size) {

    int write_fd = open(_____, _____);

    char *write_buf = _____

    int buf_size = _____
    int bytes_written = 0;
    // Write a to file.

    _____
    _____

    _____
    close(write_fd);
}
```

Now, write the function that retrieves previously saved integers and places them in a int array.

```
void read_from_file(const char *file, int *a, int size) {

    int read_fd = open(_____, _____);
```

```
char *read_buf = -----  
  
int buf_size = -----  
// Read a from a file.  
  
-----  
-----  
-----  
-----  
  
-----  
close(read_fd);  
}
```

Your coworker opens up one of the files that you used to store ints on his text editor and complains its full of junk! Explain to him why this might be the case.