

# Section 4: Sockets and Scheduling

February 21

## Contents

<b>1</b>	<b>Vocabulary</b>	<b>2</b>
<b>2</b>	<b>Scheduling</b>	<b>3</b>
2.1	Round Robin Scheduling . . . . .	3
2.2	Life Ain't Fair . . . . .	3
2.3	All Threads Must Die . . . . .	4
<b>3</b>	<b>Socket Programming</b>	<b>7</b>
3.1	Multi-threaded Echo Server . . . . .	7
<b>4</b>	<b>Files</b>	<b>9</b>
4.1	Files vs File Descriptor . . . . .	9
4.2	Reading and Writing with File Pointers vs. Descriptors . . . . .	9

# 1 Vocabulary

- **Socket** - Sockets are an abstraction of a bidirectional network I/O queue. It embodies one side of a communication channel, meaning that two must be required for a communication channel to form. The two ends of the communication channel may be local to the same machine, or they may span across different machines through the Internet. Most functions that operate on file descriptors like `read()` or `write()` work on sockets. but certain operations like `lseek()` do not.
- **file descriptors** - File descriptors are an index into a file-descriptor table stored by the kernel. The kernel creates a file-descriptor in response to an `open` call and associates the file-descriptor with some abstraction of an underlying file-like object; be that an actual hardware device, or a file-system or something else entirely. Using file descriptors, a process's `read` or `write` calls are routed to the correct place by the kernel. When your program starts you have 3 file descriptors.

File Descriptor	File
0	stdin
1	stdout
2	stderr

- **Scheduler** - Routine in the kernel that picks which thread to run next given a vacant CPU and a ready queue of unblocked threads.
- **Preemption** - The process of interrupting a running thread to allow for the scheduler to decide which thread runs next.
- **FIFO Scheduling** - First-In-First-Out (aka First-Come-First-Serve) scheduling runs jobs as they arrive. Turnaround time can degrade if short jobs get stuck behind long ones (convoy effect);
- **round-robin Scheduling** - Round-Robin scheduling runs each job in fixed-length time slices (quanta). The scheduler preempts a job that exceeds its quantum and moves on, cycling through the jobs. It avoids starvation and is good for short jobs, but context switching overhead can become important depending on quanta length;
- **Shortest Time Remaining First Scheduling** - A scheduling algorithm where the thread that runs is the one with the least time remaining. This is ideal for throughput but also must be approximated in practice.
- **Linux CFS** - Linux scheduling algorithm designed to optimize for fairness. It gives each thread a weighted share of some target latency and then ensures that each thread receives that much virtual CPU time in its scheduling decisions.
- **Earliest Deadline First** - Scheduling algorithm used in real time systems. It attempts to meet deadlines of threads that must be processed in real time by selecting the thread that has the closest deadline to complete first.

## 2 Scheduling

### 2.1 Round Robin Scheduling

Which of the following are true about Round Robin Scheduling?

1. The average wait time is less than that of FCFS for the same workload.
2. It requires pre-emption to maintain uniform quanta.
3. If quanta is constantly updated to become the # of cpu ticks since boot, Round Robin becomes FIFO.
4. If all threads in the system have the same priority, Priority Schedulers **must** behave like round robin.
5. Cache performance is likely to improve relative to FCFS.
6. If no new threads are entering the system all threads will get a chance to run in the cpu every  $QUANTA * SECONDS\_PER\_TICK * NUMTHREADS$  seconds. (Assuming  $QUANTA$  is in ticks).
7. It is the fairest scheduler

### 2.2 Life Ain't Fair

Suppose the following threads denoted by `THREADNAME : PRIORITY` pairs arrive in the ready queue at the clock ticks shown. Assume all threads arrive unblocked and that each takes 5 clock ticks to finish executing. Assume threads arrive in the queue at the beginning of the time slices shown and are ready to be scheduled in that same clock tick. (This means you update the ready queue with the arrival before you schedule/execute that clock tick.) Assume you only have one physical CPU.

```
0  Yiming : 7
1
2  Alan : 1
3  Sarah: 3
4
5  Annie : 5
6
7  Neil: 11
8
9  Alex: 14
```

Determine the order and time allocations of execution for the following scheduler scenarios:

- Round Robin with time slice 3
- Shortest Time Remaining First (SRTF/SJF) WITH preemptions
- Preemptive priority (higher is more important)

Write answers in the form of vertical columns with one name per row, each denoting one clock tick of execution. For example, allowing Yiming 3 units at first looks like:

```
0  Yiming
1  Yiming
2  Yiming
```

It will probably help you to draw a diagram of the ready queue at each tick for this problem.



### 2.3 All Threads Must Die

You have three threads with the associated priorities shown below. They each run the functions with their respective names. Assume upon execution all threads are initially unblocked and begin at the top of their code blocks. The operating system runs with a preemptive priority scheduler. You may assume that `set_priority` commands are atomic.

Tyrion : 4  
Ned: 5  
Gandalf: 11

Note: The following uses references to Pintos locks and data structures.

```
struct list braceYourself;    // pintos list. Assume it's already initialized and populated.  
struct lock midTerm;        // pintos lock. Already initialized.  
struct lock isComing;
```

```

void tyrion(){
    thread_set_priority(12);
    lock_acquire(&midTerm);
    lock_release(&midTerm);
    thread_exit();
}


void ned(){
    lock_acquire(&midTerm);
    lock_acquire(&isComing);
    list_remove(list_head(braceYourself));
    lock_release(&midTerm);
    lock_release(&isComing);
    thread_exit();
}

void gandalf(){
    lock_acquire(&isComing);
    thread_set_priority(3);
    while (thread_get_priority() < 11) {
        printf("YOU .. SHALL NOT .. PAAASS!!!!!!");
        timer_sleep(20);
    }
    lock_release(&isComing);
    thread_exit();
}

```

What is the output of this program when there is no priority donation? Trace the program execution and number the lines in the order in which they are executed.

What is the output and order of line execution if priority donation was implemented? Draw a diagram of the three threads and two locks that shows how you would use data structures and struct members (variables and pointers, etc) to implement priority donation for this example.



### 3 Socket Programming

#### 3.1 Multi-threaded Echo Server

Please look at the three versions of server code provided with Lecture 8. The first version uses a single process and single thread, the second version sequentially handles each client connection in a child process, and the third version allows child processes to handle connections concurrently.

Write a fourth version of the server implementation that uses multiple threads in a single process. Each connection is handled in its own thread, and threads should be allowed to handle connections concurrently.

```

#define BUF_SIZE 1024

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;

    int rv = getaddrinfo(NULL, port, &hints, &server);
    if (rv != 0) {
        printf("getaddrinfo failed: %s\n", gai_strerror(rv));
        return NULL;
    }
    return server;
}

void *serve_client(void *client_socket_arg) {
    int client_socket = (int)client_socket_arg;
    char buf[BUF_SIZE];
    ssize_t n;

    while ((n = read(client_socket, buf, BUF_SIZE)) > 0) {
        buf[n] = '\0';
        printf("Client Sent: %s\n", buf);

        if (write(client_socket, buf, n) == -1) {
            -----;
            -----;
        }
    }

    if (n == -1) {
        -----;
        -----;
    }

    -----;
}

```

```

    -----;
}

int main(int argc, char **argv) {
    if (argc < 2) {
        printf("Usage: %s <port>\n", argv[0]);
        return 1;
    }

    struct addrinfo *server = setup_address(argv[1]);
    if (server == NULL) {
        return 1;
    }
    int server_socket = -----(server->ai_family,
                                server->ai_socktype, server->ai_protocol);
    if (server_socket == -1) {
        return 1;
    }

    if (-----(server_socket, server->ai_addr,
                server->ai_addrlen) == -1) {
        return 1;
    }
    if (-----(server_socket, 1) == -1) {
        return 1;
    }

    while (1) {
        int connection_socket = accept(server_socket, NULL, NULL);
        if (connection_socket == -1) {
            perror("accept");
            pthread_exit(NULL);
        }

        pthread_t handler_thread;
        int err = pthread_create(-----);
        if (err != 0) {
            -----;
        }
        pthread_detach(handler_thread);
    }
    -----;
}

```



## 4 Files

### 4.1 Files vs File Descriptor

What's the difference between `fopen` and `open`?

### 4.2 Reading and Writing with File Pointers vs. Descriptors

Write a utility function, `void copy(const char *src, const char *dest)`, that simply copies the file contents from `src` and places it in `dest`. You can assume both files are already created. Also assume that the `src` file is at most 100 bytes long. First, use the file pointer library to implement this. Fill in the code given below:

```
void copy(const char *src, const char *dest) {
    char buffer [100];
    FILE* read_file = fopen(_____, ____);
    int buf_size = fread(_____, ____, _____, _____);
    fclose(read_file);
    FILE* write_file = fopen(_____, ____);
    fwrite(_____, ____, _____, _____);
    fclose(write_file);
}
```

Next, use file descriptors to implement the same thing.

```
void copy(const char *src, const char *dest) {
    char buffer [100];
    int read_fd = open(_____, _____);
    int bytes_read = 0;
    int buf_size = 0;
    while ((bytes_read = read(_____, _____, _____)) > 0) {
        _____
    }
    close(read_fd);
    int bytes_written = 0;
    int write_fd = open(_____, _____);
    while (_____) {
        _____ += write(_____, _____, _____);
    }
    close(write_fd);
}
```

Compare the file pointer implementation to the file descriptor implementation. In the file descriptor implementation, why does **read** and **write** need to be called in a loop?