

Section 4: Threads

September 16, 2015

Contents

1	Warmup	2
1.1	Hello World	2
2	Vocabulary	2
3	Problems	3
3.1	Join	3
3.2	Stack Allocation	4
3.3	Heap Allocation	4
3.4	Threads and Processes	5
3.5	Atomicity	5
3.6	Synchronization	6
3.7	Who's Your Daddy	7
3.8	Simple HTTP Server with Threads	7

1 Warmup

1.1 Hello World

What does C print in the following code?

```
void print_hello_world() {
    pid_t pid = getpid();
    printf("Hello world %d\n", pid);
    pthread_exit(0);
}

void main() {
    pthread_t thread;
    pthread_create(&thread, NULL, (void *) &print_hello_world, NULL);
    print_hello_world();
}
```

2 Vocabulary

- **thread** - a thread of execution is the smallest unit of sequential instructions that can be scheduled for execution by the operating system. Multiple threads can share the same address space, and each thread independently operates using its own program counter.
- **pthreads** - A POSIX-compliant (standard specified by IEEE) implementation of threads.
- **pthread_join** - Waits for a specific thread to terminate, similar to `waitpid(3)`. (Hint: `man pthread_join`)
- **pthread_create** - Creates and immediately starts a child thread running in the same address space of the thread that spawned it. The child executes starting from the function specified. Internally, this is implemented by calling the `clone` syscall. (Hint: `man pthread_create`)
- **atomic** - An operation is deemed to be atomic if it can be executed without interruption or interference from other threads/processes).
- **critical section** - A section of code that accesses a shared resource that must not be concurrently accessed by more than a single thread.
- **semaphore** - A synchronization primitive that can be used to protect a shared resource. Semaphores contain an integer value and support two operations:
 - 1) Increment: atomically increments the value of the semaphore. (Hint: `man sem_post`)
 - 2) Decrement: waits for the value of the semaphore to become positive, and then atomically decrements the value of the semaphore. (Hint: `man sem_wait`)
- **pthread_yield** - Equivalent to `thread_yield()` in Pintos. Causes the calling thread to vacate the CPU and go back into the ready queue without blocking. The calling thread is able to be scheduled again immediately. This is not the same as an interrupt and will succeed in Pintos even if interrupts are disabled.

- **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.

3 Problems

3.1 Join

What does C print in the following code?

(Hint: There may be zero, one, or multiple answers.)

```
void main() {
    pthread_t thread;
    pthread_create(&thread, NULL, &helper, NULL);
    printf("Hello World! 2\n");
    exit(0);
}

void *helper(void* arg) {
    printf("Hello World! 1\n");
    pthread_exit(0);
}
```

How can we modify the code above to always print out "Hello World! 1" followed by "Hello World! 2"?

3.2 Stack Allocation

What does C print in the following code?

```
void main() {
    int i = 0;
    pthread_t thread;
    pthread_create(&thread, NULL, &helper, &i);
    pthread_join(thread, NULL);
    printf("i is %d\n", i);
}
```

```
void *helper(void *arg) {
    int *num = (int*) arg;
    *num = 2;
    pthread_exit(0);
}
```

3.3 Heap Allocation

What does C print in the following code?

```
void main() {
    char *message = malloc(100);
    strcpy(message, "I am the parent");
    pthread_t thread;
    pthread_create(&thread, NULL, &helper, message);
    pthread_join(thread, NULL);
    printf("%s\n", message);
}
```

```
void *helper(void *arg) {
    char *message = (char *) arg;
    strcpy(message, "I am the child");
    pthread_exit(0);
}
```

3.4 Threads and Processes

What does C print in the following code?

(Hint: There may be zero, one, or multiple answers.)

```
void *worker(void *arg) {
    int *data = (int *) arg;
    *data = *data + 1;
    printf("Data is %d\n", *data);
    pthread_exit(0);
}

int data;
void main() {
    int status;
    data = 0;
    pthread_t thread;

    pid_t pid = fork();
    if (pid == 0) {
        pthread_create(&thread, NULL, &worker, &data);
        pthread_join(thread, NULL);
    } else {
        pthread_create(&thread, NULL, &worker, &data);
        pthread_join(thread, NULL);
        pthread_create(&thread, NULL, &worker, &data);
        pthread_join(thread, NULL);
        wait(&status);
    }
}
```

3.5 Atomicity

Given:

```
int x = 0;
```

Circle all the atomic operations below:

```
printf("x is %d\n", x);

x = malloc(sizeof(int));

int y = x;

x++;
```

3.6 Synchronization

What does C print in the following code?

```
void main() {
    pthread_t threads[5];
    int i, j = 0;
    for (i = 0; i < 5; i++) {
        pthread_create(&threads[i], NULL, &helper, &j);
    }
    for (i = 0; i < 5; i++) {
        pthread_join(threads[i], NULL);
    }
    printf("j is %d\n", j);
}

void *helper(void *arg) {
    int *num = (int *) arg;
    *num = (*num) + 1;
    pthread_exit(0);
}
```

How can we modify the code in the previous page to always print out "i is 5" without modifying/using `pthread_join`?

3.7 Who's Your Daddy

What does C print in the following code? Assume the new thread's ID is 0x4c554b45. Assume Round Robin Scheduling and no other threads/processes are running.

```
pthread_t thread_id = 0x0;

void print_hello_world() {
    thread_id = pthread_self(); //like getpid but for threads
    printf("You deallocated my father!\n");
    pthread_exit(0);
}

void main() {
    pthread_t thread;
    pthread_create(&thread, NULL, (void *) &print_hello_world, NULL);
    pthread_yield();
    printf("No, thread %#010x, I AM your father!!"\n, thread_id);
}
```



3.8 Simple HTTP Server with Threads

Implement a crude HTTP server by filling in the blanks with the necessary operations. This HTTP server spawns worker threads that are responsible for printing out the first 256 bytes from a socket. For this question, you can assume the following structures and functions are implemented appropriately.

```
// a FIFO queue consisting of nodes.
struct list;
// A single element in the FIFO queue, containing an integer.
struct node {
    int client_sock;
};

// This function sets up a socket, binds the socket to a specific port,
// and listens for connections. It returns the file descriptor
// corresponding to the socket setup.
int setupSocket();
```

```

// initializes a FIFO queue
void list_init(struct list *l);
// Pops and returns the earliest enqueued node from a list
struct node *list_pop(struct list *l);
// Appends a new file descriptor as a struct node in the list.
struct node *list_append(struct list *l, int client_sock);

```

```

// Global variables
struct list      clients;
sem_t           worker_sem;
sem_t           list_sem;

// Prints out the first 256 bytes from a socket, and then closes it.
void accept_request(void * unused) {
    char buf[256];
    (void) unused;

    while(1) {

        -----;

        -----;
        int sockfd = list_pop(&clients)->client_sock;

        if (----- != 256) {
            fprintf(stderr, "Failed to read from client %d\n", sockfd);
        } else {
            fprintf(stdout, "Received from client %d: %s\n", sockfd, buf);
        }

        -----;
        close(sockfd);
    }
}

void main() {
    pthread_t workers[10];
    int client_sock;
    struct sockaddr_in client_name;
    socklen_t size;
    int sockfd;

    // initialize variables
    list_init(&clients);
    sem_init(&worker_sem, 0, 0);

    -----;

    -----;

```



```
// create worker threads
for (int i = 0; i < 10; i++) {

    -----;
}

while (1) {
    client_sock = -----;

    -----;
    list_append(&clients, client_sock);

    -----;

    -----;
}
}
```