

# CSE 333 Section 9 - pthreads

Welcome back to section! We're glad that you're here :)

## Process

- A process has a virtual address space. Each process is started with a single thread, but can create additional threads.

## Threads

- A thread contains a sequential execution of a program.
- Contained within a process.
- Threads of the same process share a memory/address space: see the same heap and globals, but each thread has its own stack.

## POSIX threads (Pthreads)

- The POSIX standard provides APIs for dealing with threads.
- Part of the standard C/C++ libraries, declared in `pthread.h`.
- **Compile and link with** `-pthread`.
- Core pthread functions: `pthread_create`, `pthread_exit`, `pthread_join`
- ```
int pthread_create( pthread_t *thread,
                  const pthread_attr_t *attr,
                  void *(*start_routine) (void *),
                  void *arg);
```

`thread`: output parameter.

`attr`: used to set thread attributes. Use `NULL` as default.

`start_routine`: function pointer to a C routine that the thread will execute once it is created.

`arg`: a single argument that may be passed to `start_routine`. `NULL` may be used if no argument is to be passed.

Overall, it creates a new thread and calls `start_routine` with `arg` as its parameter.

Returns 0 if successful; otherwise, returns an error number.

- ```
int pthread_join( pthread_t thread, void **retval);
```

Synchronization between threads. It waits for the thread specified by `thread` to terminate. If that thread has already terminated, then it returns immediately. If `retval` is non-`NULL`, then `retval` acts as an output parameter and the address passed to `pthread_exit` by the finished thread is stored in it. For this course we can just set `retval` to `NULL`. It returns 0 if successful; otherwise, returns an error number.
- ```
void pthread_exit(void *retval);
```

It terminates the calling thread and allows the user to specify an optional termination status parameter, `retval`. For this course we can just set `retval` to `NULL`.

## Mutex

- Protect shared data from being simultaneously accessed by multiple threads.
- `pthread_mutex_init`, `pthread_mutex_lock`, `pthread_mutex_unlock`, `pthread_mutex_destroy`

- `int pthread_mutex_init(pthread_mutex_t *mutex, const pthread_mutexattr_t *attr);`  
**mutex: initializes the mutex referenced by mutex.**  
**attr: use NULL for the default values.**
- `int pthread_mutex_destroy(pthread_mutex_t *mutex);`  
**Destroys the mutex object referenced by mutex.**
- `int pthread_mutex_lock(pthread_mutex_t *mutex);`  
`int pthread_mutex_unlock(pthread_mutex_t *mutex);`  
**Use these to let a single thread access/modify shared data while blocking the other threads.**

### Pthread Example

```
// pthread example. This program dispatches some threads to run
// thread_main. thread_main does two separate things: (1) casts the
// argument to an int and prints it out; (2) updates sum_total.
```

```
#include <pthread.h>
#include <iostream>

using std::cout;
using std::cerr;
using std::endl;

const int NUM_THREADS = 50;
const int LOOP_NUM = 10000;

static int sum_total = 0;
static pthread_mutex_t sum_lock;

void *thread_main(void *arg) {
    int *num = reinterpret_cast<int*>(arg);
    cout << "[cthread: " << *num << "]" << endl;

    for (int i = 0; i < LOOP_NUM; i++) {
        pthread_mutex_lock(&sum_lock);
        sum_total++;
        pthread_mutex_unlock(&sum_lock);
    }

    delete num;
    return NULL;
}
```

```

int main(int argc, char** argv) {
    pthread_t thds[NUM_THREADS];
    pthread_mutex_init(&sum_lock, NULL);

    for (int i = 0; i < NUM_THREADS; i++) {
        int *num = new int(i);
        if (pthread_create(&thds[i], NULL, &thread_main, num) != 0) {
            /*report error*/
        }
    }

    for (int i = 0; i < NUM_THREADS; i++) {
        if (pthread_join(thds[i], NULL) != 0) { /*report error*/ }
    }

    cout << "Total: " << sum_total << endl;

    pthread_mutex_destroy(&sum_lock);
    return 0;
}

```

### Question

If we have

```

MyClass onTheStack;
pthread_t child;
pthread_create(&child, nullptr, foo, &onTheStack);

```

`onTheStack` is on the parents stack. However, each thread has its own stack. Can we still access `onTheStack` from the child? Why or why not?

## Exercise

**1) Write code to print "Woof!" "Meow!" and "Ssss!" from three different child threads:**

```
void* woof(void* data) {
    std::cout << "Woof!" << std::endl;
}

void* meow(void* data) { ... }

void* ssss(void* data) { ... }

int main() {
    // create a woof thread

    // create the meow and ssss threads

    // exit without stopping any running threads
}
```

**What are some possible outputs of this program?**

**2) Calculating primes is slow. Use 10 threads to calculate the first 1,000 primes. Then, print them out in ascending order:**

```
#define NTHREAD 10
struct Bounds {
    int lo;
    int hi;
    Bounds(int lo, int hi): lo(lo), hi(hi) {}
};

bool isPrime(int num) { ... }

void* getPrimes(void* data) {
    Bounds* b = reinterpret_cast<Bounds*>(data);
    // setup a way to store the primes we find in order

    // calculate primes

    // ???
    return
}
```

// continued on next page

```
int main() {
    // make space to store our threads and data
    std::vector<std::unique_ptr<Bounds>> bounds;

    // create and run our threads
    int err;
    for (int i = 0; i < NTHREAD; i++) {

    }

    // wait for each thread to finish and get its data
    for (int i = 0; i < NTHREAD; i++) {
        // wait for thread, storing its return value

        // print the data

    }

    return 0;
}
```

## Boost Library

- Very useful for dealing with strings (HW 4!), such as trimming, pattern matching, splitting, replacing, etc.
- `#include <boost/algorithm/string.hpp>`.
- **Examples:** `boost::split`, `boost::to_upper`, etc.
- **View boost libraries online.**