


TCSS 422: OPERATING SYSTEMS

Three Easy Pieces:  
Thread API

Wes J. Lloyd  
Institute of Technology  
University of Washington - Tacoma

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma



OBJECTIVES

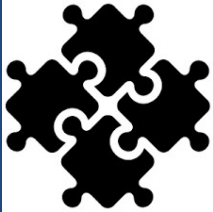
- Linux Thread API – Ch. 27

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.2

# CHAPTER 27 - LINUX THREAD API



April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.3

## THREAD CREATION

- **pthread\_create**

```
#include <pthread.h>

int
pthread_create(      pthread_t*      thread,
                    const pthread_attr_t* attr,
                    void*          (*start_routine) (void*),
                    void*          arg);
```

- **thread**: thread struct
- **attr**: stack size, scheduling priority... (*optional*)
- **start\_routine**: function pointer to thread routine
- **arg**: argument to pass to thread routine (*optional*)

April 18, 2018	TCSS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L7a.4
----------------	--	-------

## PTHREAD\_CREATE – PASS ANY DATA

```
#include <pthread.h>

typedef struct __myarg_t {
    int a;
    int b;
} myarg_t;

void *mythread(void *arg) {
    myarg_t *m = (myarg_t *) arg;
    printf("%d %d\n", m->a, m->b);
    return NULL;
}

int main(int argc, char *argv[]) {
    pthread_t p;
    int rc;

    myarg_t args;
    args.a = 10;
    args.b = 20;
    rc = pthread_create(&p, NULL, mythread, &args);
    ...
}
```

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.5

## PASSING A SINGLE VALUE

Using this approach on your Ubuntu VM,  
How large (in bytes) can the primitive data type be?

How large (in bytes) can the primitive data type  
be on a 32-bit operating system?

```
9     int rc, m;
10    pthread_create(&p, NULL, mythread, (void *) 100);
11    pthread_join(p, (void **) &m);
12    printf("returned %d\n", m);
13    return 0;
14 }
```

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.6

```
int pthread_join(pthread_t thread, void **value_ptr);
```

- April 18, 2018

L7a.7

```
struct myarg {
    int a;
    int b;
};
```

```

struct myarg *input = (struct myarg *) arg;
printf("a=%d b=%d\n", input->a, input->b);
struct myarg output;
output.a = 1;
output.b = 2;
return (void *) &output;

```

**Data on thread s**

**\$/pt**  
**a=10**

```
int main (int argc, char * argv[])
{
    pthread_t p1;
    struct myarg args;
    struct myarg *ret_args;
    args.a = 10;
    args.b = 20;
    pthread_
    pthread_
    printf("
    return 0;
}
```

```
$ ./pthread_struct
a=10 b=20
Segmentation fault (core dumped)
```

## How can this code be fixed?

L7a.8

```

struct myarg {
    int a;
    int b;
};

void *worker(void *arg)
{
    struct myarg *input = (struct myarg *) arg;
    printf("a=%d b=%d\n", input->a, input->b);
    input->a = 1;
    input->b = 2;
    return (void *) &input;
}

int main (int argc, char * argv[])
{
    pthread_t p1;
    struct myarg args;
    struct myarg *ret_args;
    args.a = 10;
    args.b = 20;
    pthread_create(&p1, NULL, worker, &args);
    pthread_join(p1, (void *)&ret_args);
    printf("returned %d %d\n", ret_args->a, ret_args->b);
    return 0;
}

```

## How about this code?

**\$ ./pthread\_struct  
a=10 b=20  
returned 1 2**

April 18, 2018
TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma
L7a.9

## ADDING CASTS

- Casting
- Suppresses compiler warnings when passing “typed” data where (void) or (void \*) is called for
- Example: uncasted capture in pthread\_join

```

pthread_int.c: In function 'main':
pthread_int.c:34:20: warning: passing argument 2 of 'pthread_join'
from incompatible pointer type [-Wincompatible-pointer-types]
    pthread_join(p1, &p1val);

```
- Example: uncasted return

```

In file included from pthread_int.c:3:0:
/usr/include/pthread.h:250:12: note: expected 'void **' but argument
is of type 'int **'
extern int pthread_join (pthread_t __th, void **__thread_return);

```

April 18, 2018
TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma
L7a.10

## ADDING CASTS - 2

- **pthread\_join**

```
int * p1val;
int * p2val;
pthread_join(p1, (void *)&p1val);
pthread_join(p2, (void *)&p2val);
```
- **return from thread function**

```
int * counterval = malloc(sizeof(int));
*counterval = counter;
return (void *) counterval;
```

April 18, 2018

TCS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.11

## LOCKS

- **pthread\_mutex\_t data type**
- **/usr/include/bits/pthread\_types.h**

```
// Global Address Space
static volatile int counter = 0;
pthread_mutex_t lock;

void *worker(void *arg)
{
    int i;
    for (i=0;i<10000000;i++) {
        int rc = pthread_mutex_lock(&lock);
        assert(rc==0);
        counter = counter + 1;
        pthread_mutex_unlock(&lock);
    }
    return NULL;
}
```

April 18, 2018

TCS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.12

## LOCKS - 2

- Ensure critical sections are executed atomically-as a *unit*
  - Provides implementation of “**Mutual Exclusion**”

- API

```
int pthread_mutex_lock(pthread_mutex_t *mutex);  
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

- Example w/o initialization & error checking



```
pthread_mutex_t lock;  
pthread_mutex_lock(&lock);  
x = x + 1; // or whatever your critical section is  
pthread_mutex_unlock(&lock);
```

- Blocks forever until lock can be obtained
- Enters critical section once lock is obtained
- Releases lock

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.13

## LOCK INITIALIZATION

- Assigning the constant

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
```

- API call:

```
int rc = pthread_mutex_init(&lock, NULL);  
assert(rc == 0); // always check success!
```

- Initializes mutex with attributes specified by 2<sup>nd</sup> argument
- If NULL, then default attributes are used
- Upon initialization, the mutex is initialized and unlocked

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.14

## LOCKS - 3

### ■ Error checking wrapper

```
// Use this to keep your code clean but check for failures
// Only use if exiting program is OK upon failure
void Pthread_mutex_lock(pthread_mutex_t *mutex) {
    int rc = pthread_mutex_lock(mutex);
    assert(rc == 0);
}
```

### ■ What if lock can't be obtained?

```
int pthread_mutex_trylock(pthread_mutex_t *mutex);
int pthread_mutex_timelock(pthread_mutex_t *mutex,
                          struct timespec *abs_timeout);
```

- trylock – returns immediately (fails) if lock is unavailable
- timelock – tries to obtain a lock for a specified duration

April 18, 2018

TCCS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.15

## CONDITIONS AND SIGNALS

### ■ Condition variables support “signaling” between threads

```
int pthread_cond_wait(pthread_cond_t *cond,
                     pthread_mutex_t *mutex);
int pthread_cond_signal(pthread_cond_t *cond);
```



### ■ pthread\_cond\_t datatype

### ■ pthread\_cond\_wait()

- Puts thread to “sleep” (waits) (THREAD is BLOCKED)
- Threads added to FIFO queue, lock is released
- Waits (*listens*) for a “signal” (NON-BUSY WAITING, no polling)
- When signal occurs, interrupt fires, wakes up first thread, (THREAD is RUNNING), lock is provided to thread

April 18, 2018

TCCS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.16



## CONDITIONS AND SIGNALS - 2

```
int pthread_cond_signal(pthread_cond_t * cond);
int pthread_cond_broadcast(pthread_cond_t * cond);
```

- `pthread_cond_signal()`
  - Called to send a “signal” to wake-up first thread in FIFO “wait” queue
  - The goal is to unblock a thread to respond to the signal
- `pthread_cond_broadcast()`
  - Unblocks all threads in FIFO “wait” queue, currently blocked on the specified condition variable
  - Broadcast is used when all threads should wake-up for the signal
- Which thread is unblocked first?
  - Determined by OS scheduler (based on priority)
  - Thread(s) awoken based on placement order in FIFO wait queue
  - When awoken threads acquire lock as in `pthread_mutex_lock()`

April 18, 2018

TCS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.17

## CONDITIONS AND SIGNALS - 3

- Wait example:

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;

pthread_mutex_lock(&lock);
while (initialized == 0)
    pthread_cond_wait(&cond, &lock);
// Perform work that requires lock
a = a + b;
pthread_mutex_unlock(&lock);
```

- wait puts thread to sleep, releases lock
- when awoken, lock reacquired (but then released by this code)
- When initialized, another thread signals

```
pthread_mutex_lock(&lock);
initialized = 1;
pthread_cond_signal(&cond);
pthread_mutex_unlock(&lock);
```

State variable set,  
Enables other thread(s)  
to proceed above.

April 18, 2018

TCS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.18

## CONDITION AND SIGNALS - 4

```
pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;  
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;  
  
pthread_mutex_lock(&lock);  
while (initialized == 0)  
    pthread_cond_wait(&cond, &lock);  
// Perform work that requires lock  
a = a + b;  
pthread_mutex_unlock(&lock);
```

- Why do we wait inside a while loop?
- The while ensures upon awakening the condition is rechecked
  - A signal is raised, but the pre-conditions required to proceed may have not been met. **\*\*MUST CHECK STATE VARIABLE\*\***
  - Without checking the state variable the thread may proceed to execute when it should not. (e.g. too early)

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.19

## PTHREADS LIBRARY

- Compilation
  - `gcc -pthread pthread.c -o pthread`
  - Requires explicitly linking the library with compiler flag
  - Use makefile to provide compiler arguments
- List of pthread manpages
  - `man -k pthread`

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.20

## SAMPLE MAKEFILE

```
CC=gcc
CFLAGS=-pthread -I. -Wall

binaries=pthread pthread_int pthread_lock_cond pthread_struct

all: $(binaries)

pthread_mult: pthread.c pthread_int.c
    $(CC) $(CFLAGS) $^ -o $@

clean:
    $(RM) -f $(binaries) *.o
```

- **Example builds multiple single file programs**
  - All target
- **pthread\_mult**
  - Example if multiple source files should produce a single executable
- **clean target**

April 18, 2018

TCSS422: Operating Systems [Spring 2018]  
Institute of Technology, University of Washington - Tacoma

L7a.21

## QUESTIONS

