


TCCS 422: OPERATING SYSTEMS

Locks API, Introduction to Locks, Lock-Based Data Structures

Wes J. Lloyd
 School of Engineering and Technology,
 University of Washington - Tacoma



October 22, 2018
TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma
Tacoma

FEEDBACK FROM 10/17

- **What resources do threads share?**
- **pthread.c example:**
 - **Do locks cause the two threads to both ping back and forth until both reach desired count?**
 - **Why do locks cause both thread values to not be overridden with each other?**
 - **Why did the worker function in the pthread.c example have an asterisk before it?**
 - void * - is a void pointer – essentially an untyped pointer
 - Worker function uses this to avoid compiler warning to match typing of the pthread_create() function signature

October 22, 2018
TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma
L8.2

FEEDBACK - 2

- **What is the purpose of joining threads?**
 - When a thread exits, the parent can join to receive return results from the worker method
- **Is the only purpose for locking to protect variables from outside manipulation?**
 - Locks can also be used to order the sequence of execution
 - Who goes first...
 (though condition variables are technically better...)
- **Could you simply poll a variable (e.g. int ready) to enforce sequence of execution?**
- **How is using pthread_mutex_t() better than polling?**

October 22, 2018
TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma
L8.3

Why does this code seg fault?

```

struct my {
  int a;
  int b;
};

void *worker(void *arg)
{
  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;
}

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

← Data on thread stack

```

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

What would be another example where joining would cause a seg fault?


October 22, 2018
TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma
L8.4

OBJECTIVES

- Program 1 – MASH Shell (Friday 10/26)
- Midterm – (Wed 10/31)
- **Multi-threaded Programming**
 - Chapter 27 – Linux Thread API
 - Chapter 28 – Introduction to Locks
 - Chapter 29 – Lock-based Data Structures


October 22, 2018
TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma
L8.5

CHAPTER 9 - PROPORTIONAL SHARE SCHEDULER



October 22, 2018
TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma
L8.6

CHAPTER 27 - LINUX THREAD API



October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.7

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);
```

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.8

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;
```

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.9

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;
}
```

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.10

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

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.11

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 2nd argument
- If NULL, then default attributes are used
- Upon initialization, the mutex is initialized and unlocked

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.12

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

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.13

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

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.14

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()

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.15

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(&init);
                pthread_mutex_unlock(&lock);
            
```

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

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.16

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)

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L8.17

If multiple threads are blocked waiting for a signal after calling pthread_cond_wait(), when pthread_cond_signal() is called which thread will execute next?

W

- The thread that most recently called pthread_cond_signal()
- The thread that most recently called pthread_cond_wait()
 - The thread that first called pthread_cond_signal()
 - The thread that first called pthread_cond_wait()
- The thread that most recently returned from pthread_join()

Start the presentation to see live content. Still no live content? Install the app or get help at PallEx.com/app Total Results

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`

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L8.19

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


October 22, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L8.20

CHAPTER 28 – LOCKS



October 22, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L8.21

LOCKS



- Ensure critical section(s) are executed atomically-as a unit
 - Only one thread is allowed to execute a critical section at any given time
 - Ensures the code snippets are “mutually exclusive”
- Protect a global counter:

```
balance = balance + 1;
```
- A “critical section”:

```
1 lock_t mutex; // some globally-allocated lock 'mutex'
2 ..
3 lock(&mutex);
4 balance = balance + 1;
5 unlock(&mutex);
```

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L8.22

LOCKS - 2

- Lock variables are called “MUTEX”
 - Short for mutual exclusion (that’s what they guarantee)
- Lock variables store the state of the lock
- States
 - **Locked** (acquired or held)
 - **Unlocked** (available or free)
- Only 1 thread can hold a lock

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L8.23

LOCKS - 3

- `pthread_mutex_lock(&lock)`
 - Try to acquire lock
 - If lock is free, calling thread will acquire the lock
 - Thread with lock enters critical section
 - Thread “owns” the lock
- No other thread can acquire the lock before the owner releases it.

October 22, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L8.24

LOCKS - 4

- Program can have many mutex (lock) variables to “serialize” many critical sections
- Locks are also used to protect data structures
 - Prevent multiple threads from changing the same data simultaneously
 - Programmer can make sections of code “granular”
 - Fine grained – means just one grain of sand at a time through an hour glass
 - Similar to relational database transactions
 - DB transactions prevent multiple users from modifying a table, row, field


October 22, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L8.25

FINE GRAINED?

- Is this code a good example of “fine grained parallelism”?

```
pthread_mutex_lock(&lock);
a = b++;
b = a * c;
*d = a + b + c;
FILE * fp = fopen ("file.txt", "r");
fscanf(fp, "%s %s %d", str1, str2, str3, &e);
ListNode *node = mylist->head;
int i=0
while (r
    node->
    node->subheading = str2;
    node->desc = str3;
    node->end = *e;
    node = node->next;
    i++;
}
e = e - i;
pthread_mutex_unlock(&lock);
```

Example of coarse-grained parallelism



October 22, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L8.26

FINE GRAINED PARALLELISM


```
pthread_mutex_lock(&lock_a);
pthread_mutex_lock(&lock_b);
a = b++;
pthread_mutex_unlock(&lock_b);
pthread_mutex_unlock(&lock_a);

pthread_mutex_lock(&lock_b);
b = a * c;
pthread_mutex_unlock(&lock_b);

pthread_mutex_lock(&lock_d);
*d = a + b + c;
pthread_mutex_unlock(&lock_d);

FILE * fp = fopen ("file.txt", "r");
pthread_mutex_lock(&lock_e);
fscanf(fp, "%s %s %d", str1, str2, str3, &e);
pthread_mutex_unlock(&lock_e);


ListNode *node = mylist->head;
int i=0 . . .
```



October 22, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L8.27

EVALUATING LOCK IMPLEMENTATIONS

- Correctness
 - Does the lock work?
 - Are critical sections mutually exclusive? (atomic-as a unit?)
- Fairness
 - Are threads competing for a lock have a fair chance of acquiring it?
- Overhead



October 22, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L8.28

BUILDING LOCKS

- Locks require hardware support
 - To minimize overhead, ensure fairness and correctness
 - Special “atomic-as a unit” instructions to support lock implementation
 - Atomic-as a unit exchange instruction
 - XCHG
 - Compare and exchange instruction
 - CMPXCHG
 - CMPXCHGB
 - CMPXCHG16B

October 22, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L8.29

QUESTIONS

