


TCSS 422: OPERATING SYSTEMS

Concurrency: An Introduction



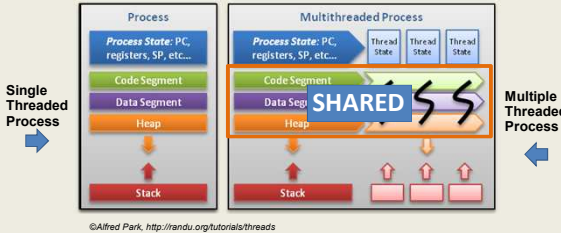
Wes J. Lloyd
Institute of Technology
University of Washington - Tacoma

OBJECTIVES

- Introduction to threads
- Race condition
- Critical section
- Thread API

October 12, 2016 TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma L6.2

THREADS



©Alfred Park, <http://randu.org/tutorials/threads>

October 12, 2016 TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma L6.3

THREADS - 2

- Enables a single process (program) to have multiple “workers”
- Supports independent path(s) of execution within a program
- Each thread has its own Thread Control Block (TCB)
 - PC, registers, SP, and stack
- Code segment, memory, and heap are shared

October 12, 2016 TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma L6.4

PROCESS AND THREAD METADATA

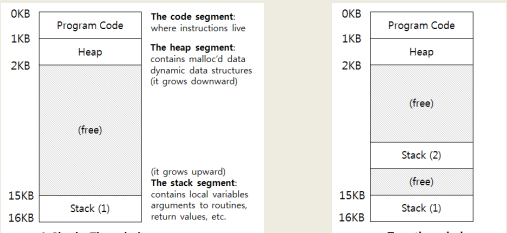
- Thread Control Block vs. Process Control Block

Thread Identification	Process Identification
Thread state	Process status
CPU information:	Process state:
Program counter	Process status word
Register contents	Register contents
Thread priority	Main memory
Pointer to process that created this thread	Resources
Pointers to all other threads created by this thread	Process priority
	Accounting

October 12, 2016 TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma L6.5

SHARED ADDRESS SPACE

- Every thread has its own stack / PC



October 12, 2016 TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma L6.6

THREAD CREATION EXAMPLE

```
#include <stdio.h>
#include <assert.h>
#include <pthread.h>

void *mythread(void *arg) {
    printf("%s\n", (char *) arg);
    return NULL;
}

int
main(int argc, char *argv[]) {
    pthread_t p1, p2;
    int rc;
    printf("main: begin\n");
    rc = pthread_create(&p1, NULL, mythread, "A"); assert(rc == 0);
    rc = pthread_create(&p2, NULL, mythread, "B"); assert(rc == 0);
    // join waits for the threads to finish
    rc = pthread_join(p1, NULL); assert(rc == 0);
    rc = pthread_join(p2, NULL); assert(rc == 0);
    printf("main: end\n");
    return 0;
}
```

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.7

POSSIBLE ORDERINGS OF EVENTS

int main()	Thread 1	Thread 2
Starts running		
Prints 'main: begin'		
Creates Thread 1		
Creates Thread 2		
Waits for T1		
	Runs	
	Prints 'A'	
	Returns	
Waits for T2		
		Runs
		Prints 'B'
		Returns
Prints 'main: end'		

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.8

POSSIBLE ORDERINGS OF EVENTS - 2

int main()	Thread 1	Thread 2
Starts running		
Prints 'main: begin'		
Creates Thread 1		
	Runs	
	Prints 'A'	
	Returns	
Creates Thread 2		
		Runs
		Prints 'B'
		Returns
Waits for T1	Returns immediately	
Waits for T2		Returns immediately
Prints 'main: end'		

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.9

POSSIBLE ORDERINGS OF EVENTS - 3

int main()	Thread 1	Thread 2
Starts running		
Prints 'main: begin'		
Creates Thread 1		
Creates Thread 2		
Waits for T1		
	Runs	
	Prints 'A'	
	Returns	
Waits for T2		
		Immediately returns
Prints 'main: end'		

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.10

What if execution order of events in the program matters?

COUNTER EXAMPLE

- Show example
- A + B : ordering
- Counter: incrementing global variable by two threads

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.11

RACE CONDITION

- Example when counter=50
 - Counter = counter + 1
 - If synchronized, counter will = 52

OS	Thread1	Thread2	(after instruction) PC %eax counter
	before critical section		100 0 50
	mov 0x8049a1c, %eax		105 50 50
	add \$0x1, %eax		108 51 50
interrupt	save T1's state		
	restore T2's state		100 0 50
		mov 0x8049a1c, %eax	105 50 50
		add \$0x1, %eax	108 51 50
		mov %eax, 0x8049a1c	113 51 51
interrupt	save T2's state		
	restore T1's state		108 51 50
	mov %eax, 0x8049a1c		113 51 51

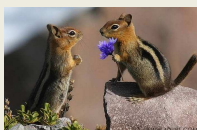
October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.12

CRITICAL SECTION

- Code that accesses a shared variable must not be **concurrently** executed by more than one thread
- Multiple **active** threads inside a critical section produces a **race condition**.
- Atomicity** of execution must be ensured in **critical** sections
 - These sections must be **mutually exclusive**



October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.13

LOCKS

- To demonstrate how critical section(s) can be executed "atomically" Chapter 27 & beyond introduce locks

```
1 lock_t mutex;
2 . . .
3 lock(&mutex);
4 balance = balance + 1;
5 unlock(&mutex);
```

Critical section

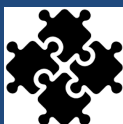
- Counter example revisited

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.14

LINUX THREAD API



October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.15

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...
- start_routine**: function pointer to thread routine
- arg**: argument to pass to thread routine

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.16

THREAD CREATION - 2

- 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 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.17

ADDING CASTS

- 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 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.18

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);
    ...
}
```

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.19

WAITING FOR THREADS TO FINISH

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

- thread: which thread?
- value_ptr: pointer to return value
type is dynamic / agnostic
- Returned values *must* be on the heap
- Thread stacks destroyed upon thread termination (join)
- Pointers to thread stack memory addresses are invalid
 - May appear as gibberish or lead to crash (seg fault)
- Not all threads join - **What would be Examples ??**

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.20

What will this code do?

```
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);
    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;
    pthread_create(&p1, NULL, worker, &args);
    pthread_join(p1, (void **)&ret_args);
    printf("returned %d\n", ret_args->a);
    return 0;
}
```

Data on thread stack

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

How can this code be fixed?

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.21

How about this code?

```
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\n", ret_args->a);
    return 0;
}
```

```
$ ./pthread_struct
a=10 b=20
returned 1 2
```

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.22

PASSING A SINGLE VALUE

Using this approach on your CentOS 7 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 }
```

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.23

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<100000000; i++) {
        int rc = pthread_mutex_lock(&lock);
        assert(rc==0);
        counter = counter + 1;
        pthread_mutex_unlock(&lock);
    }
    return NULL;
}
```

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.24

LOCKS - 2

- Ensure critical sections are executed atomically
 - 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 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.25

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 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.26

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 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.27

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()

- Waits (sleeps)
- Listens for a "signal"
- Releases the lock until signaled



October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.28

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 all listeners → to wake them up
 - The goal is to unblock (at least one) to respond to the signal
- pthread_cond_broadcast()
 - Unblocks all threads currently blocked on the specified condition
 - Used when all threads should respond to the signal
- Which thread is unblocked first?
 - Determined by OS scheduler (based on priority)
 - Thread(s) gain the lock individually (based on priority) as if they called pthread_mutex_lock()

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.29

CONDITIONS AND SIGNALS - 3

- Wait example

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

pthread_mutex_lock(&lock);
while (initialized == 0)
    pthread_cond_wait(&init, &lock);
pthread_mutex_unlock(&lock);
```

- wait puts thread to sleep, releases lock
- when awoken, lock reacquired (and released by this code)
- Another thread signals the thread

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

Code performs required work before other thread(s) can continue

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.30

CONDITION AND SIGNALS - 4

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

pthread_mutex_lock(&lock);
while (initialized == 0)
    pthread_cond_wait(&init, &lock);
pthread_mutex_unlock(&lock);
```

- Why do we wait inside a while loop?
 - The while ensures upon awakening the condition is rechecked
 - A signal may have been raised, but the condition to proceed has not been satisfied.
 - Without checking the condition the thread may proceed to execute when it should not.

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.31

DO-IT-YOURSELF LOCK

- "initialized" -- global variable shared by multiple threads

- Wait (client thread):

```
while(initialized == 0)
    ; // spin
```

- Signal (parent thread): *when ready...*

```
initialized = 1;
```

- How is this "wait" different than pthread_cond_wait() ?

- Wastes CPU cycles → effectively pegs a core at 100%
- Potential synchronization errors with changing the value of "initialized"
- Thread API is provided to advance the DO-IT-YOURSELF approach

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.32

PTHREADS LIBRARY

- Compilation
 - gcc -pthread pthread.c -o pthread
 - Requires explicitly linking the library with compiler flag
- List of pthread manpages
 - man -k pthread

October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.33

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 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.34

QUESTIONS



October 12, 2016

TCSS422: Operating Systems [Fall 2016]
Institute of Technology, University of Washington - Tacoma

L6.35