


TCCS 422: OPERATING SYSTEMS

Linux Thread API,
Locks,
Lock-Based Data Structures



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

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

Tacoma

OBJECTIVES – 4/23

■ Questions from 4/21

■ Assignment 0 (Apr 23 11:59p AOE)

■ C Tutorial (Apr 30 11:59p AOE)

■ Assignment 1 (May 7 11:59p AOE)

■ Chapter 27: Linux Thread API

- pthread_create/_join
- pthread_mutex_lock/_unlock/_trylock/_timelock
- pthread_cond_wait/_signal/_broadcast

■ Chapter 28: Locks

- Introduction, Lock Granularity
- Spin Locks, Test and Set, Compare and Swap

■ Chapter 29: Lock Based Data Structures

- Concurrent Data Structures

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.2

MATERIAL / PACE

■ Please classify your perspective on material covered in today's class (48 respondents):

■ 1-mostly review, 5-equal new/review, 10-mostly new

■ Average – 7.32 (↓ from 7.6)

■ Please rate the pace of today's class:

■ 1-slow, 5-just right, 10-fast

■ Average – 5.63 (↑ from 5.45)

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.3

FEEDBACK FROM 4/23

■ *I'm confused on how the scheduling question was done but I'm hoping the lecture tonight (Wed 4/22) will clear up any questions I have.*

■ 7 scheduling examples, about an hour

■ Video posted

■ 1 FIFO, 1 SJF, 1 STCF

■ 4 MLFQ

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.4

FEEDBACK - 2

■ *I would like to have more clarity on lock and why is it slow?*

```
// Global Address Space
#define COUNT 8000000
static volatile int counter = 0;
pthread_mutex_t lock;

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

#1: 8,000,000 function calls

#2: Calls are system calls (kernel API)
Context switch req'd to kernel worker process to perform requested work with privileged access to the HW

#3: Mutual Exclusion:
If another thread is already executing inside the Critical Section, then it blocks (running → blocked) and waits for the lock to become available.

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.5

RACE CONDITION

■ What is happening with our counter?

■ When counter=50, consider code: 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

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma


L8.6

Slides by Wes J. Lloyd

L8.1

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** produce a **race condition**.
- Atomic execution** (all code executed as a unit) must be ensured in **critical** sections
 - These sections must be **mutually exclusive**




April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.7

ANYWHERE ON EARTH (AOE)
ASSIGNMENTS REVISED SUBMISSION TIME

- TCSS 422 B Sp 2020
- Now due at 11:59 pm **Anywhere On Earth** (AOE) TIME ZONE
- Pacific Daylight Time minus 5 hours
- Last time zone before international date line
- 11:59 pm = ~4:59 am PDT



April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.8

OBJECTIVES – 4/23

- Questions from 4/21
- Assignment 0 (Apr 23 11:59p AOE)
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 27: Linux Thread API
 - pthread_create/_join
 - pthread_mutex_lock/_unlock/_trylock/_timelock
 - pthread_cond_wait/_signal/_broadcast
- Chapter 28: Locks
 - Introduction, Lock Granularity
 - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
 - Concurrent Data Structures

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.9

OBJECTIVES – 4/23

- Questions from 4/21
- Assignment 0 (Apr 23 11:59p AOE)
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 27: Linux Thread API
 - pthread_create/_join
 - pthread_mutex_lock/_unlock/_trylock/_timelock
 - pthread_cond_wait/_signal/_broadcast
- Chapter 28: Locks
 - Introduction, Lock Granularity
 - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
 - Concurrent Data Structures

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.10

OBJECTIVES – 4/23


- Questions from 4/21
- Assignment 0 (Apr 23 11:59p AOE)
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 27: Linux Thread API
 - pthread_create/_join
 - pthread_mutex_lock/_unlock/_trylock/_timelock
 - pthread_cond_wait/_signal/_broadcast
- Chapter 28: Locks
 - Introduction, Lock Granularity
 - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
 - Concurrent Data Structures

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.11

CHAPTER 27 -
LINUX
THREAD API



April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.12

OBJECTIVES – 4/23

- Questions from 4/21
- Assignment 0 (Apr 23 11:59p AOE)
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 27: Linux Thread API
 - pthread_create/_join
 - pthread_mutex_lock/_unlock/_trylock/_timelock
 - pthread_cond_wait/_signal/_broadcast
- Chapter 28: Locks
 - Introduction, Lock Granularity
 - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
 - Concurrent Data Structures

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.13

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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.14

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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.15

PASSING A SINGLE VALUE

■ Here we “cast” the pointer to pass/return a primitive data type

```
1 void *mythread(void *arg) {
2     int m = (int) arg;
3     printf("%d\n", m);
4     return (void *) (arg + 1);
5 }
6
7 int main(int argc, char *argv[]) {
8     pthread_t p;
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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.16

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
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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.17

For pthread_create(), how large (in bytes) can a casted primitive data type be that is passed in as a replacement for (void *) on a 32-bit operating system?

- 1 byte
- 2 bytes
- 3 bytes
- 4 bytes

Unlimited size

Start the presentation to see live content. Still no live content? Install the app or get help at [pellix.com/app](https://www.pellix.com/app)

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 ??*

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.19

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

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.20

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

Data on thread stack

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

How can this code be fixed?

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.21

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

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

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.22

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 1 is of type 'int **'
extern int pthread_join (pthread_t __th, void **__thread_return);

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.23

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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.24

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;
}
```

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.25

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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.26

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

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.27

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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.28

Which Linux Thread API function *is* a blocking function? A blocking function causes the thread to go from READY --> BLOCKED.



April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.29

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 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.30



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 23, 2020

TCCS422: Operating Systems (Spring 2020)
School of Engineering and Technology, University of Washington - Tacoma

L8.31

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.

April 23, 2020

TCCS422: Operating Systems (Spring 2020)
School of Engineering and Technology, University of Washington - Tacoma

L8.32

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 23, 2020

TCCS422: Operating Systems (Spring 2020)
School of Engineering and Technology, University of Washington - Tacoma

L8.33

Five pthreads wait using a condition variable. Which pthread is woken and provided the lock FIRST when a signal or broadcast occurs?

The first thread to block and wait on the condition variable.

A

100%

The last thread to block and wait on the condition variable.

B

All threads that have blocked and are waiting on the condition variable are awoken at the same time.

C

None of the Above

D

All of the Above

E

April 23, 2020

TCCS422: Operating Systems (Spring 2020)
Start the pre-school of Engineering and Technology, University of Washington - Tacoma

L8.34

PTHREADS LIBRARY

- **Compilation:**
gcc requires special option to require programs with pthreads:
 - `gcc -pthread pthread.c -o pthread`
 - Explicitly links library with compiler flag
 - **RECOMMEND:** using makefile to provide compiler arguments
- List of pthread manpages
 - `man -k pthread`

April 23, 2020

TCCS422: Operating Systems (Spring 2020)
School of Engineering and Technology, University of Washington - Tacoma

L8.35

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 23, 2020

TCCS422: Operating Systems (Spring 2020)
School of Engineering and Technology, University of Washington - Tacoma

L8.36

TCSS 422 WILL RETURN AT ~2:50PM

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.37

CHAPTER 28 – LOCKS

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.38

OBJECTIVES – 4/23

- Questions from 4/21
- Assignment 0 (Apr 23 11:59p AOE)
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 27: Linux Thread API
 - pthread_create/_join
 - pthread_mutex_lock/_unlock/_trylock/_timelock
 - pthread_cond_wait/_signal/_broadcast
- Chapter 28: Locks
 - Introduction, Lock Granularity
 - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
 - Concurrent Data Structures

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.39

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);
```

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.40

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

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.41

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.

April 23, 2020

TCSS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.42

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

April 23, 2020


TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.43

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 %s %d", str1, str2, str3, &e);
ListNode *node = mylist->head;
int i=0
while (node) {
    node->title = str1;
    node->subheading = str2;
    node->desc = str3;
    node->end = *e;
    node = node->next;
    i++;
}
e = e - i;
pthread_mutex_unlock(&lock);
```



April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.44

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 %s %d", str1, str2, str3, &e);
pthread_mutex_unlock(&lock_e);

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



April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.45

EVALUATING LOCK IMPLEMENTATIONS

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



April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.46

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
 - CMPPXCHG
 - CMPPXCHG8B
 - CMPPXCHG16B

April 23, 2020

TCCS422: Operating Systems [Spring 2020]
School of Engineering and Technology, University of Washington - Tacoma

L8.47

HISTORICAL IMPLEMENTATION

- To implement mutual exclusion
 - Disable interrupts upon entering critical sections
- ```
1 void lock() {
2 DisableInterrupts();
3 }
4 void unlock() {
5 EnableInterrupts();
6 }
```
- Any thread could disable system-wide interrupt
    - What if lock is never released?
  - On a multiprocessor processor each CPU has its own interrupts
    - Do we disable interrupts for all cores simultaneously?
  - While interrupts are disabled, they could be lost
    - If not queued...

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.48



## SPIN LOCK IMPLEMENTATION

- Operate without atomic-as a unit assembly instructions
- "Do-it-yourself" Locks
- Is this lock implementation: (1)Correct? (2)Fair? (3)Performant?



```
1 typedef struct __lock_t { int flag; } lock_t;
2
3 void init(lock_t *mutex) {
4 // 0 → lock is available, 1 → held
5 mutex->flag = 0;
6 }
7
8 void lock(lock_t *mutex) {
9 while (mutex->flag == 1) // TEST the flag
10 ; // spin-wait (do nothing)
11 mutex->flag = 1; // now SET it !
12 }
13
14 void unlock(lock_t *mutex) {
15 mutex->flag = 0;
16 }
```

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.49

## DIY: CORRECT?

- Correctness requires luck... (e.g. DIY lock is incorrect)

| Thread1                           | Thread2                       |
|-----------------------------------|-------------------------------|
| call lock()                       |                               |
| while (flag == 1)                 |                               |
| interrupt: switch to Thread 2     |                               |
|                                   | call lock()                   |
|                                   | while (flag == 1)             |
|                                   | flag = 1;                     |
|                                   | interrupt: switch to Thread 1 |
| flag = 1; // set flag to 1 (too!) |                               |

- Here both threads have "acquired" the lock simultaneously

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.50

## DIY: PERFORMANT?

```
void lock(lock_t *mutex)
{
 while (mutex->flag == 1); // while lock is unavailable, wait...
 mutex->flag = 1;
}
```

- What is wrong with while(<cond>); ?
- Spin-waiting wastes time actively waiting for another thread
- while (1); will "peg" a CPU core at 100%
  - Continuously loops, and evaluates mutex->flag value...
  - Generates heat...

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.51

## TEST-AND-SET INSTRUCTION

- Hardware support required for working locks
- Book presents pseudo code of C implementation
  - TEST-and-SET adds a simple check to the basic spin lock
  - Assumption is this "C code" runs atomically on CPU:

```
1 int TestAndSet(int *ptr, int new) {
2 int old = *ptr; // fetch old value at ptr
3 *ptr = new; // store 'new' into ptr
4 return old; // return the old value
5 }
```

- lock() method checks that TestAndSet doesn't return 1
- Comparison is in the caller
- Can implement the C version (non-atomic) and have some success on a single-core VM

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.52

## DIY: TEST-AND-SET - 2

- C version: requires preemptive scheduler on single core system
- Lock is never released without a context switch
- single-core VM: occasionally will deadlock, doesn't miscount

```
1 typedef struct __lock_t {
2 int flag;
3 } lock_t;
4
5 void init(lock_t *lock) {
6 // 0 indicates that lock is available,
7 // 1 that it is held
8 lock->flag = 0;
9 }
10
11 void lock(lock_t *lock) {
12 while (TestAndSet(&lock->flag, 1) == 1)
13 ; // spin-wait
14 }
15
16 void unlock(lock_t *lock) {
17 lock->flag = 0;
18 }
```

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.53

## SPIN LOCK EVALUATION

- Correctness:**
  - Spin locks with atomic Test-and-Set: Critical sections won't be executed simultaneously by (2) threads
- Fairness:**
  - No fairness guarantee. Once a thread has a lock, nothing forces it to relinquish it...
- Performance:**
  - Spin locks perform "busy waiting"
  - Spin locks are best for short periods of waiting (< 1 time quantum)
  - Performance is slow when multiple threads share a CPU
    - Especially if "spinning" for long periods

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.54

## COMPARE AND SWAP

- Checks that the lock variable has the expected value **FIRST**, before changing its value
  - If so, make assignment
  - Return value at location
- Adds a comparison to TestAndSet
  - Textbook presents C pseudo code
  - Assumption is that the compare-and-swap method runs atomically
- Useful for wait-free synchronization
  - Supports implementation of shared data structures which can be updated atomically (as a *unit*) using the HW support CompareAndSwap instruction
  - Shared data structure updates become “wait-free”
  - Upcoming in Chapter 32

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.55

## COMPARE AND SWAP

### Compare and Swap

```
1 int CompareAndSwap(int *ptr, int expected, int new) {
2 int actual = *ptr;
3 if (actual == expected)
4 *ptr = new;
5 return actual;
6 }
```

### Spin lock

**C implementation 1-core VM:  
Count is correct, no deadlock**

```
1 int spin_lock = 0;
2
3 // spin
4 }
```

- X86 provides “cmpxchg1” compare-and-exchange instruction
  - cmpxchg8b
  - cmpxchg16b

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.56

## TWO MORE “LOCK BUILDING” CPU INSTRUCTIONS

- Cooperative instructions used together to support synchronization on RISC systems
- No support on x86 processors
  - Supported by RISC: Alpha, PowerPC, ARM
- Load-linked (LL)
  - Loads value into register
  - Same as typical load
  - Used as a mechanism to track competition
- Store-conditional (SC)
  - Performs “mutually exclusive” store
  - Allows only one thread to store value

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.57

## LL/SC LOCK

```
1 int LoadLinked(int *ptr) {
2 return *ptr;
3 }
4
5 int StoreConditional(int *ptr, int value) {
6 if (no one has updated *ptr since the LoadLinked to this address) {
7 *ptr = value;
8 return 1; // success!
9 } else {
10 return 0; // failed to update
11 }
12 }
```

- LL instruction loads pointer value (ptr)
- SC only stores if the load link pointer has not changed
- Requires HW support
  - C code is pseudo code

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.58

## LL/SC LOCK - 2

```
1 void lock(lock_t *lock) {
2 while (1) {
3 while (LoadLinked(&lock->flag) == 1)
4 ; // spin until it's zero
5 if (StoreConditional(&lock->flag, 1) == 1)
6 return; // if set-it-to-1 was a success: all done
7 }
8 }
9
10 void unlock(lock_t *lock) {
11 lock->flag = 0;
12 }
13 }
```

- Two instruction lock

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.59

## CHAPTER 29 – LOCK BASED DATA STRUCTURES

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.60

## OBJECTIVES – 4/23

- Questions from 4/21
- Assignment 0 (Apr 23 11:59p AOE)
- C Tutorial (Apr 30 11:59p AOE)
- Assignment 1 (May 7 11:59p AOE)
- Chapter 27: Linux Thread API
  - pthread\_create/\_join
  - pthread\_mutex\_lock/\_unlock/\_trylock/\_timelock
  - pthread\_cond\_wait/\_signal/\_broadcast
- Chapter 28: Locks
  - Introduction, Lock Granularity
  - Spin Locks, Test and Set, Compare and Swap
- Chapter 29: Lock Based Data Structures
  - Concurrent Data Structures

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.61

## LOCK-BASED CONCURRENT DATA STRUCTURES

- Adding locks to data structures make them **thread safe**.
- Considerations:
  - Correctness
  - Performance
  - Lock granularity

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.62

## COUNTER STRUCTURE W/O LOCK

- Synchronization weary --- not thread safe

```

1 typedef struct __counter_t {
2 int value;
3 } counter_t;
4
5 void init(counter_t *c) {
6 c->value = 0;
7 }
8
9 void increment(counter_t *c) {
10 c->value++;
11 }
12
13 void decrement(counter_t *c) {
14 c->value--;
15 }
16
17 int get(counter_t *c) {
18 return c->value;
19 }
```

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.63

## CONCURRENT COUNTER

```

1 typedef struct __counter_t {
2 int value;
3 pthread_mutex_t lock;
4 } counter_t;
5
6 void init(counter_t *c) {
7 c->value = 0;
8 pthread_mutex_init(&c->lock, NULL);
9 }
10
11 void increment(counter_t *c) {
12 pthread_mutex_lock(&c->lock);
13 c->value++;
14 pthread_mutex_unlock(&c->lock);
15 }
16
```

- Add lock to the counter
- Require lock to change data

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.64

## CONCURRENT COUNTER - 2

- Decrease counter
- Get value

```

(Cont.)
17 void decrement(counter_t *c) {
18 pthread_mutex_lock(&c->lock);
19 c->value--;
20 pthread_mutex_unlock(&c->lock);
21 }
22
23 int get(counter_t *c) {
24 pthread_mutex_lock(&c->lock);
25 int rc = c->value;
26 pthread_mutex_unlock(&c->lock);
27 return rc;
28 }
```


April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.65

## CONCURRENT COUNTERS - PERFORMANCE

- iMac: four core Intel 2.7 GHz i5 CPU
- Each thread increments counter 1,000,000 times



Traditional vs. sloppy counter  
Sloppy Threshold (S) = 1024

**Synchronized counter scales poorly.**

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.66

PERFECT SCALING

- Achieve (N) performance gain with (N) additional resources
- Throughput:
  - Transactions per second
- 1 core
  - N = 100 tps
- 10 core
  - N = 1000 tps

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.67

SLOPPY COUNTER

- Provides single logical shared counter
  - Implemented using local counters for each ~CPU core
    - 4 CPU cores = 4 local counters & 1 global counter
    - Local counters are synchronized via local locks
  - Global counter is updated periodically
    - Global counter has lock to protect global counter value
    - Sloppiness threshold (S):
      - Update threshold of global counter with local values
    - Small (S): more updates, more overhead
    - Large (S): fewer updates, more performant, less synchronized
- Why this implementation?
  - Why do we want counters local to each CPU Core?

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.68

SLOPPY COUNTER - 2

- Update threshold (S) = 5
- Synchronized across four CPU cores
- Threads update local CPU counters

| Time | L <sub>1</sub> | L <sub>2</sub> | L <sub>3</sub> | L <sub>4</sub> | G                         |
|------|----------------|----------------|----------------|----------------|---------------------------|
| 0    | 0              | 0              | 0              | 0              | 0                         |
| 1    | 0              | 0              | 1              | 1              | 0                         |
| 2    | 1              | 0              | 2              | 1              | 0                         |
| 3    | 2              | 0              | 3              | 1              | 0                         |
| 4    | 3              | 0              | 3              | 2              | 0                         |
| 5    | 4              | 1              | 3              | 3              | 0                         |
| 6    | 5 → 0          | 1              | 3              | 4              | 5 (from L <sub>1</sub> )  |
| 7    | 0              | 2              | 4              | 5 → 0          | 10 (from L <sub>4</sub> ) |

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.69

THRESHOLD VALUE S

- Consider 4 threads increment a counter 1000000 times each
- Low S → What is the consequence?
- High S → What is the consequence?

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.70

SLOPPY COUNTER - EXAMPLE

- Example implementation
- Also with CPU affinity

April 23, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L8.71

QUESTIONS

