

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

L7.2

L7.3

L7.4

L7.5

L7.6

NICE VALUES AND PROCESS PRIORITIES

- **Nice values** (*user space suggestions for process priority*) :
 - -20 HIGH, 0 NORMAL, 19 LOW
- **Process priorities** (*actual kernel space value*):
 - 0 LOW (user), 39 HIGH (user), >40 to 139 (MAX) Realtime
- To check priorities:
 - `ps ax -o pid,ni,pri,cmd,%cpu`
 - `top` (PR NI columns)

January 29, 2018

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

L7.7

FEEDBACK - 6

- **What is an example of a program that benefits from multi-threading?**
 - Parts of the computation must be separable to run at same time (parallel)
 - Embarrassingly parallel:
Separate parts of computation can run independently without communication
- **Is multithreading overhead offset by the performance gained from parallel processing?**
 - It is entirely dependent in what the program is doing...

January 29, 2018

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

L7.8

FEEDBACK - 7

- **Is there better ways then locking code to make It run slightly faster If assembly code is known?**
 - Question refers to the fact that incrementing a variable in C requires three non-atomic lines of assembly code
 - Coming soon
- **Blue comment bars in downloadable slides are obscuring content**
 - Actually no, slides are duplicated with and without messages...

January 29, 2018

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

L7.9

CHAPTER 27 - LINUX THREAD API



January 29, 2018

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

L7.10

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*)

January 29, 2018

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

L7.11

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);
    ...
}
```

January 29, 2018

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

L7.12

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 }
```

January 29, 2018

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

L7.13

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 ??*

January 29, 2018

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

L7.14

```
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;
}
```

January 29, 2018

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

L7.15

```
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;
}
```

What will this code do?

Data on thread stack

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

How can this code be fixed?

January 29, 2018

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

L7.16

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**

January 29, 2018

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

L7.17

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);

January 29, 2018

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

L7.18

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;
```

January 29, 2018

TCS5422: Operating Systems [Winter 2018]
Institute of Technology, University of Washington - Tacoma

L7.19

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;
}
```

January 29, 2018

TCS5422: Operating Systems [Winter 2018]
Institute of Technology, University of Washington - Tacoma

L7.20

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

January 29, 2018

TCS5422: Operating Systems [Winter 2018]
Institute of Technology, University of Washington - Tacoma

L7.21

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**

January 29, 2018

TCS5422: Operating Systems [Winter 2018]
Institute of Technology, University of Washington - Tacoma

L7.22

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

January 29, 2018

TCS5422: Operating Systems [Winter 2018]
Institute of Technology, University of Washington - Tacoma

L7.23

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



January 29, 2018

TCS5422: Operating Systems [Winter 2018]
Institute of Technology, University of Washington - Tacoma

L7.24

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()`

January 29, 2018

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

L7.25

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

January 29, 2018

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

L7.26

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.

January 29, 2018

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

L7.27

PTHREADS LIBRARY

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

January 29, 2018

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

L7.28

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

January 29, 2018

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

L7.29


CHAPTER 28 – LOCKS

January 29, 2018

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

L7.30

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);
```

January 29, 2018	TCCS422: Operating Systems [Winter 2018] Institute of Technology, University of Washington - Tacoma	L7.31
------------------	--	-------

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

January 29, 2018	TCCS422: Operating Systems [Winter 2018] Institute of Technology, University of Washington - Tacoma	L7.32
------------------	--	-------

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.

January 29, 2018	TCCS422: Operating Systems [Winter 2018] Institute of Technology, University of Washington - Tacoma	L7.33
------------------	--	-------

LOCKS - 4

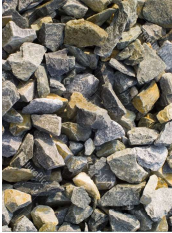
- Programs can have many mutex (lock) variables to "serialize" many critical sections
- Locks are also used to protect data structures
 - Prevent multiple threads from changing the same data simultaneously
 - Programmer can make sections of code "granular"
 - Fine grained – means just one grain of sand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

January 29, 2018	TCCS422: Operating Systems [Winter 2018] Institute of Technology, University of Washington - Tacoma	L7.34
------------------	--	-------

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);
```



January 29, 2018	TCCS422: Operating Systems [Winter 2018] Institute of Technology, University of Washington - Tacoma	L7.35
------------------	--	-------

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 . . .
```




January 29, 2018	TCCS422: Operating Systems [Winter 2018] Institute of Technology, University of Washington - Tacoma	L7.36
------------------	--	-------

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

January 29, 2018

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

L7.37

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
 - CMPXCHG8B
 - CMPXCHG16B

January 29, 2018

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

L7.38

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...

January 29, 2018

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


L7.39

SPIN LOCK IMPLEMENTATION

■ Operate without atomic-as a unit assembly instructions

■ “Do-it-yourself” Locks

■ Is this lock implementation: Correct? Fair? 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 }
```

January 29, 2018

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

L7.40

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

January 29, 2018

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

L7.41

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...

January 29, 2018

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

L7.42

Slides by Wes J. Lloyd

L7.7

TEST-AND-SET INSTRUCTION

- C implementation: not atomic
 - Adds a simple check to basic spin lock
 - One a single core CPU system with preemptive scheduler:
 - Try this...

```
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
- Single core systems are becoming scarce
- Try on a one-core VM

January 29, 2018

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

L7.43

DIY: TEST-AND-SET - 2

- Requires a preemptive scheduler on single CPU core system
- Lock is never released without a context switch
- 1-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 }
```

January 29, 2018

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

L7.44

SPIN LOCK EVALUATION

- **Correctness:**
 - Spin locks guarantee: 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
 - Performance is slow when multiple threads share a CPU
 - Especially for long periods

January 29, 2018

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

L7.45

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
- 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

January 29, 2018

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

L7.46

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 usage

```
1 void lock(lock_t *lock) {
2     while (CompareAndSwap(&lock->flag, 0, 1) == 1)
3         ; // spin
4 }
```
- X86 provides "cmpxchg1" compare-and-exchange instruction
 - cmpxchg8b
 - cmpxchg16b

January 29, 2018

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

L7.47

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

```
1 void lock(lock_t *lock) {
2     while (CompareAndSwap(&lock->flag, 0, 1) == 1)
3         ; // spin
4 }
```
- X86 provides "cmpxchg1" compare-and-exchange instruction
 - cmpxchg8b
 - cmpxchg16b

January 29, 2018

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

L7.48

