


TCSS 422: OPERATING SYSTEMS

Three Easy Pieces:
Ch. 28 - Locks

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

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma



OBJECTIVES

■ Locks – Ch. 28

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.2

CHAPTER 28 – LOCKS

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.3



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);
```

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.4

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

Chapter 28

TCCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.5

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.

Chapter 28

TCCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.6

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

Chapter 28


TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.7

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);
```



Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.8

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 . . .
```



Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.9

EVALUATING LOCK IMPLEMENTATIONS

■ Correctness

- Does the lock work?
- Are critical sections mutually exclusive?
(atomic-as a unit?)



■ Fairness

- Do threads competing for a lock have a fair chance of acquiring it?

■ Overhead

Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.10

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

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.11

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...

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.12

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 }
    
```

Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.13

DIY: CORRECT?

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

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

- Here both threads have “acquired” the lock simultaneously

Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.14

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...

Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.15

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

Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.16

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 }

```

Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.17

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

Chapter 28

TCS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.18

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

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.19

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  while (!lock)
2      ;
3      ; // spin
4  }
```

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

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

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.20

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

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.21

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 psuedo code

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.22

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                     otherwise: try it all over again  
8     }  
9 }  
10  
11 void unlock(lock_t *lock) {  
12     lock->flag = 0;  
13 }
```

■ Two instruction lock

Chapter 28

TCSS422: Operating Systems
Institute of Technology, University of Washington - Tacoma

L7b.23

QUESTIONS

