

## Mutual Exclusion: Primitives and Implementation Considerations

- ### Too Much Milk: Lessons
- ◆ Software solution (Peterson's algorithm) works, but it is unsatisfactory
    - Solution is complicated; proving correctness is tricky even for the simple example
    - While thread is waiting, it is consuming CPU time
    - Asymmetric solution exists for 2 processes.
  - ◆ How can we do better?
    - Use hardware features to eliminate busy waiting
    - Define higher-level programming abstractions to simplify concurrent programming

### Concurrency Quiz

If two threads execute this program concurrently, how many different final values of X are there?  
Initially, X == 0.

Thread 1

```
void increment() {
  int temp = X;
  temp = temp + 1;
  X = temp;
}
```

Thread 2

```
void increment() {
  int temp = X;
  temp = temp + 1;
  X = temp;
}
```

Answer:  
A. 0  
B. 1  
C. 2  
D. More than 2

### Schedules/Interleavings

- ◆ Model of concurrent execution
- ◆ Interleave statements from each thread into a single thread
- ◆ If **any** interleaving yields incorrect results, some synchronization is needed

Thread 1

```
tmp1 = X;
tmp1 = tmp1 + 1;
X = tmp1;
```

Thread 2

```
X;
tmp2 = X;
tmp2 = tmp2 + 1;
X = tmp2;
```

If X==0 initially, X == 1 at the end. WRONG result!

### Locks fix this with Mutual Exclusion

```
void increment() {
  lock.acquire();
  int temp = X;
  temp = temp + 1;
  X = temp;
  lock.release();
}
```

- ◆ Mutual exclusion ensures only safe interleavings
  - When is mutual exclusion too safe?

### Introducing Locks

- ◆ Locks – implement mutual exclusion
  - Two methods
    - ◊ Lock::Acquire() – wait until lock is free, then grab it
    - ◊ Lock::Release() – release the lock, waking up a waiter, if any
- ◆ With locks, too much milk problem is very easy!
  - Check and update happen as one unit (exclusive access)

```
Lock.Acquire();
if (noMilk) {
  buy milk;
}
Lock.Release();
```

```
Lock.Acquire();
x++;
Lock.Release();
```

How can we implement locks?

How to think about synchronization code	
<ul style="list-style-type: none"> <li>Every thread has the same pattern           <ul style="list-style-type: none"> <li>Entry section: code to attempt entry to critical section</li> <li>Critical section: code that requires isolation (e.g., with mutual exclusion)</li> <li>Exit section: cleanup code after execution of critical region</li> <li>Non-critical section: everything else</li> </ul> </li> <li>There can be multiple critical regions in a program           <ul style="list-style-type: none"> <li>Only critical regions that access the same resource (e.g., data structure) need to synchronize with each other</li> </ul> </li> </ul> <pre> while(1) {     Entry section     Critical section     Exit section     Non-critical section } </pre>	7

The correctness conditions	
<ul style="list-style-type: none"> <li><b>Safety</b> <ul style="list-style-type: none"> <li>Only one thread in the critical region</li> </ul> </li> <li><b>Liveness</b> <ul style="list-style-type: none"> <li>Some thread that enters the entry section eventually enters the critical region</li> <li>Even if other thread takes forever in non-critical region</li> </ul> </li> <li><b>Bounded waiting</b> <ul style="list-style-type: none"> <li>A thread that enters the entry section enters the critical section within some bounded number of operations.</li> </ul> </li> <li><b>Failure atomicity</b> <ul style="list-style-type: none"> <li>It is OK for a thread to die in the critical region</li> <li>Many techniques do not provide failure atomicity</li> </ul> </li> </ul> <pre> while(1) {     Entry section     Critical section     Exit section     Non-critical section } </pre>	8

Read-Modify-Write (RMW)	
<ul style="list-style-type: none"> <li>Implement locks using read-modify-write instructions           <ul style="list-style-type: none"> <li>As an atomic and isolated action               <ol style="list-style-type: none"> <li>read a memory location into a register, AND</li> <li>write a new value to the location</li> </ol> </li> <li>Implementing RMW is tricky in multi-processors               <ul style="list-style-type: none"> <li>Requires cache coherence hardware. Caches snoop the memory bus.</li> </ul> </li> </ul> </li> <li>Examples:           <ul style="list-style-type: none"> <li>Test&amp;set instructions (most architectures)               <ul style="list-style-type: none"> <li>Reads a value from memory</li> <li>Write "1" back to memory location</li> </ul> </li> <li>Compare &amp; swap (a.k.a. cmpxchg on x86)               <ul style="list-style-type: none"> <li>Test the value against some constant</li> <li>If the test returns true, set value in memory to different value</li> <li>Report the result of the test in a flag</li> <li>if [addr] == r1 then [addr] = r2;</li> </ul> </li> <li>Double Compare &amp; Swap (68000)               <ul style="list-style-type: none"> <li>Variant: if [addr1] == r1 then [addr2] = r2</li> </ul> </li> <li>Exchange, locked increment, locked decrement (x86)</li> <li>Load linked/store conditional (PowerPC, Alpha, MIPS)</li> </ul> </li> </ul>	9

Implementing Locks with Test&set		
<pre> int lock_value = 0; int* lock = &amp;lock_value; </pre> <pre> Lock::Acquire() {     while (test&amp;set(lock) == 1)         ; // spin } </pre> <pre> Lock::Release() {     *lock = 0; } </pre>	<ul style="list-style-type: none"> <li>If lock is free (lock_value == 0), then test&amp;set reads 0 and sets value to 1 → lock is set to busy and Acquire completes</li> <li>If lock is busy, the test&amp;set reads 1 and sets value to 1 → no change in lock's status and Acquire loops</li> <li>Does this lock have bounded waiting?</li> </ul>	10

Locks and Busy Waiting	
<pre> Lock::Acquire() {     while (test&amp;set(lock) == 1)         ; // spin } </pre> <ul style="list-style-type: none"> <li><b>Busy-waiting:</b> <ul style="list-style-type: none"> <li>Threads consume CPU cycles while waiting</li> <li>Low latency to acquire</li> </ul> </li> <li><b>Limitations</b> <ul style="list-style-type: none"> <li>Occupies a CPU core</li> <li>What happens if threads have different priorities?           <ul style="list-style-type: none"> <li>Busy-waiting thread remains runnable</li> <li>If the thread waiting for a lock has higher priority than the thread occupying the lock, then ?</li> <li>Ugh, I just wanted to lock a data structure, but now I'm involved with the scheduler!</li> </ul> </li> <li>What if programmer forgets to unlock?</li> </ul> </li> </ul>	11

Remember to always release locks	
<ul style="list-style-type: none"> <li>Java provides a convenient mechanism.</li> </ul> <pre> import     java.util.concurrent.locks.ReentrantLock; public static final aLock = new     ReentrantLock();  aLock.lock(); try {     ... } finally {     aLock.unlock(); } return 0; </pre>	12

### Remember to always release locks

◆ Java also has implicit locks:

```
synchronized void method(void) {
    XXX
}
```

is short for

```
void method(void) {
    synchronized(this) {
        XXX }
}
```

is short for

```
void method(void) {
    this.l.lock();
    try {
        XXX } finally {
            this.l.unlock();
        }
}
```

13

### Cheaper Locks with Cheaper busy waiting

Using Test&Set

```
Lock::Acquire() {
    while (test&set(lock) == 1);
}
```

```
Lock::Acquire() {
    while(1) {
        if (test&set(lock) == 0) break;
        else sleep(1);
    }
}
```

With busy-waiting                      With voluntary yield of CPU

```
Lock::Release() {
    *lock = 0;
}
```

```
Lock::Release() {
    *lock = 0;
}
```

◆ What is the problem with this?

- A. CPU usage B. Memory usage C. Lock::Acquire() latency
- D. Memory bus usage E. Messes up interrupt handling

14

### Test & Set with Memory Hierarchies

What happens to lock variable's cache line when different cpu's contend for the same lock?

Load can stall

15

### Cheap Locks with Cheap busy waiting

Using Test&Test&Set

```
Lock::Acquire() {
    while (test&set(lock) == 1);
}
```

```
Lock::Acquire() {
    while(1) {
        while (*lock == 1); // spin just reading
        if (test&set(lock) == 0) break;
    }
}
```

Busy-wait on in-memory copy                      Busy-wait on cached copy

```
Lock::Release() {
    *lock = 0;
}
```

```
Lock::Release() {
    *lock = 0;
}
```

◆ What is the problem with this?

- A. CPU usage B. Memory usage C. Lock::Acquire() latency
- D. Memory bus usage E. Does not work

16

### Test & Set with Memory Hierarchies

What happens to lock variable's cache line when different cpu's contend for the same lock?

17

### Test & Set with Memory Hierarchies

What happens to lock variable's cache line when different cpu's contend for the same lock?

18

Implementing Locks: Summary	
<ul style="list-style-type: none"> <li>◆ Locks are higher-level programming abstraction               <ul style="list-style-type: none"> <li>➢ Mutual exclusion can be implemented using locks</li> </ul> </li> <li>◆ Lock implementation generally requires some level of hardware support               <ul style="list-style-type: none"> <li>➢ Details of hardware support affects efficiency of locking</li> </ul> </li> <li>◆ Locks can busy-wait, and busy-waiting cheaply is important               <ul style="list-style-type: none"> <li>➢ Soon come primitives that block rather than busy-wait</li> </ul> </li> </ul>	19

Best Practices for Lock Programming (So Far...)	
<ul style="list-style-type: none"> <li>◆ When you enter a critical region, check what may have changed while you were spinning               <ul style="list-style-type: none"> <li>➢ Did Jill get milk while I was waiting on the lock?</li> </ul> </li> <li>◆ Always unlock any locks you acquire</li> </ul>	20

Implementing Locks without Busy Waiting (blocking) Using Test&Set	
<pre>Lock::Acquire() {   while (test&amp;set(lock) == 1)     ; // spin }</pre> <p style="text-align: center;">With busy-waiting</p> <pre>Lock::Release() {   *lock := 0; }</pre> <pre>Lock::Switch() {   q_lock = 0;   pid = schedule();   if (waited_on_lock(pid))     while (test&amp;set(q_lock) == 1)       dispatch pid; }</pre>	<pre>Lock::Acquire() {   if (test&amp;set(q_lock) == 1) {     Put TCB on wait queue for lock;     Lock::Switch(); // dispatch thread   } }</pre> <p style="text-align: center;">Without busy-waiting, use a queue</p> <pre>Lock::Release() {   if (wait queue is not empty) {     Move 1 (or all?) waiting threads to ready queue;   }   *q_lock = 0; }</pre> <p style="text-align: center;">Must only 1 thread be awakened?</p>
21	21

Implementing Locks: Summary	
<ul style="list-style-type: none"> <li>◆ Locks are higher-level programming abstraction               <ul style="list-style-type: none"> <li>➢ Mutual exclusion can be implemented using locks</li> </ul> </li> <li>◆ Lock implementations have 2 key ingredients:               <ul style="list-style-type: none"> <li>➢ Hardware instruction that does atomic read-modify-write                   <ul style="list-style-type: none"> <li>✦ Uni- and multi-processor architectures</li> </ul> </li> <li>➢ Blocking mechanism                   <ul style="list-style-type: none"> <li>✦ Busy waiting, or</li> <li>✦ Block on a scheduler queue in the OS</li> </ul> </li> </ul> </li> <li>◆ Locks are good for mutual exclusion but weak for coordination, e.g., producer/consumer patterns.</li> </ul>	22

Why Locks are Hard (Preview)	
<ul style="list-style-type: none"> <li>◆ Coarse-grain locks               <ul style="list-style-type: none"> <li>➢ Simple to develop</li> <li>➢ Easy to avoid deadlock</li> <li>➢ Few data races</li> <li>➢ Limited concurrency</li> </ul> </li> <li>◆ Fine-grain locks               <ul style="list-style-type: none"> <li>➢ Greater concurrency</li> <li>➢ Greater code complexity</li> <li>➢ Potential deadlocks                   <ul style="list-style-type: none"> <li>✦ Not composable</li> </ul> </li> <li>➢ Potential data races                   <ul style="list-style-type: none"> <li>✦ Which lock to lock?</li> </ul> </li> </ul> </li> </ul>	<pre>// WITH FINE-GRAIN LOCKS void move(T s, T d, Obj key) {   LOCK(s);   LOCK(d);   tmp = s.remove(key);   d.insert(key, tmp);   UNLOCK(d);   UNLOCK(s); }</pre> <pre>Thread 0      Thread 1 move(a, b, key1);  move(b, a, key2);</pre> <p style="text-align: center;">DEADLOCK!</p>
23	23