

CS162  
Operating Systems and  
Systems Programming  
Lecture 4

Synchronization, Atomic operations,  
Locks, Semaphores

September 12, 2011  
Anthony D. Joseph and Ion Stoica  
<http://inst.eecs.berkeley.edu/~cs162>

Review: Event Programming – Analogy

- Once upon a time (before cell-phones)



- Ann calls Bank of America to verify her account...
- ... she is put on hold...
- ... Mary, her roommate, waits for an important phone call
  
- Is there a better solution?
- Yes! Ann can ask the clerk to call her back ;-)
  
- Event programming:
  - Non-blocking I/O call: returns immediately after I/O call
  - Program periodically checks or interrupted when I/O completes

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.2

Review: Another Concurrent  
Program Example

- Two threads, A and B, compete with each other
  - One tries to increment a shared counter
  - The other tries to decrement the counter

<u>Thread A</u>	<u>Thread B</u>
<code>i = 0;</code>	<code>i = 0;</code>
<code>while (i &lt; 10)</code>	<code>while (i &gt; -10)</code>
<code>  i = i + 1;</code>	<code>  i = i - 1;</code>
<code>  printf("A wins!");</code>	<code>  printf("B wins!");</code>

- Assume that memory loads and stores are atomic, but incrementing and decrementing are *not* atomic
- Who wins?
- Is it guaranteed that someone wins? Why or why not?
- What if both threads have their own CPU running at same speed? Is it guaranteed that it goes on forever?

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.3

Goals for Today

- Synchronization
- Hardware Support for Synchronization
- Higher-level Synchronization Abstractions
  - Semaphores

**Note: Some slides and/or pictures in the following are adapted from slides ©2005 Silberschatz, Galvin, and Gagne. Many slides generated by Kubiawicz.**

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.4

## Motivation: “Too much milk”

- Great thing about OS’s – analogy between problems in OS and problems in real life
  - Help you understand real life problems better
  - But, computers are much stupider than people
- Example: People need to coordinate:



Time	Person A	Person B
3:00	Look in Fridge. Out of milk	
3:05	Leave for store	
3:10	Arrive at store	Look in Fridge. Out of milk
3:15	Buy milk	Leave for store
3:20	Arrive home, put milk away	Arrive at store
3:25		Buy milk
3:30		Arrive home, put milk away

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.5

## Definitions

- **Synchronization**: using atomic operations to ensure cooperation between threads
  - For now, only loads and stores are atomic
  - We’ll show its hard to build anything useful with only reads and writes
- **Mutual Exclusion**: ensuring that only one thread does a particular thing at a time
  - One thread *excludes* the other while doing its task
- **Critical Section**: piece of code that only one thread can execute at once
  - Critical section is the result of mutual exclusion
  - Critical section and mutual exclusion are two ways of describing the same thing

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.6

## More Definitions

- **Lock**: prevents someone from doing something
  - Lock before entering critical section and before accessing shared data
  - Unlock when leaving, after accessing shared data
  - Wait if locked
  - » **Important idea**: all synchronization involves waiting
- For example: fix the milk problem by putting a key on the refrigerator
  - Lock it and take key if you are going to go buy milk
  - Fixes too much: roommate angry if only wants orange juice



– Of Course – We don’t know how to make a lock yet

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.7

## Too Much Milk: Correctness Properties

- Need to be careful about correctness of concurrent programs, since non-deterministic
  - Always write down behavior first
  - Impulse is to start coding first, then when it doesn’t work, pull hair out
  - Instead, think first, then code
- What are the correctness properties for the “Too much milk” problem?
  - Never more than one person buys
  - Someone buys if needed
- Restrict ourselves to use only atomic load and store operations as building blocks

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.8

## Too Much Milk: Solution #1

- Use a note to avoid buying too much milk:
  - Leave a note before buying (kind of “lock”)
  - Remove note after buying (kind of “unlock”)
  - Don’t buy if note (wait)
- Suppose a computer tries this (remember, only memory read/write are atomic):

```
if (noMilk) {
  if (noNote) {
    leave Note;
    buy milk;
    remove note;
  }
}
```



- Result?

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.9

## Too Much Milk: Solution #1

- Still too much milk **but only occasionally!**

```
Thread A          Thread B
if (noMilk)
  if (noNote) {
    if (noMilk)
      if (noNote) {
        leave Note;
        buy milk;
        remove note;
      }
    }
}

    leave Note;
    buy milk;
    ...
```

- Thread can get context switched after checking milk and note but before buying milk!
- Solution makes problem worse since fails **intermittently**
  - Makes it really hard to debug...
  - Must work despite what the thread dispatcher does!

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.10

## Too Much Milk: Solution #1½

- Clearly the Note is not quite blocking enough
  - Let’s try to fix this by placing note first
- Another try at previous solution:

```
leave Note;
if (noMilk) {
  if (noNote) {
    buy milk;
  }
}
remove note;
```



- What happens here?
  - Well, with human, probably nothing bad
  - With computer: no one ever buys milk

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.11

## Too Much Milk Solution #2

- How about labeled notes?
  - Now we can leave note before checking
- Algorithm looks like this:

```
Thread A          Thread B
leave note A;     leave note B;
if (noNote B) {  if (noNote A) {
  if (noMilk) {   if (noMilk) {
    buy Milk;     buy Milk;
  }               }
}                 }
remove note A;   remove note B;
```

- Does this work?

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.12

## Too Much Milk Solution #2

- Possible for neither thread to buy milk!

Thread A	Thread B
<pre>leave note A;</pre>	<pre>leave note B; if (noNote A) {   if (noMilk) {     buy Milk;   } }</pre>
<pre>if (noNote B) {   if (noMilk) {     buy Milk;     ...</pre>	<pre>remove note B;</pre>

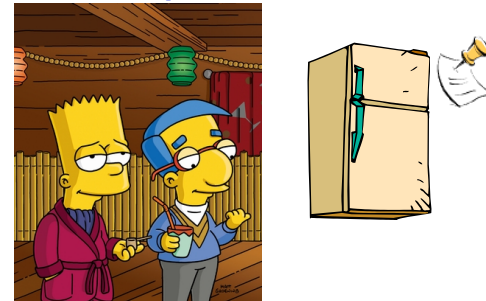
- Really insidious:
  - **Unlikely** that this would happen, but will at worst possible time

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.13

## Too Much Milk Solution #2: problem!



- *I'm not getting milk, You're getting milk*
- This kind of lockup is called "starvation!"

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.14

## Review: Too Much Milk Solution #3

- Here is a possible two-note solution:

Thread A	Thread B
<pre>leave note A; while (note B) {\X   do nothing; } if (noMilk) {   buy milk; } remove note A;</pre>	<pre>leave note B; if (noNote A) {\Y   if (noMilk) {     buy milk;   } } remove note B;</pre>

- Does this work? Yes. Both can guarantee that:
  - It is safe to buy, or
  - Other will buy, ok to quit
- At X:
  - if no note B, safe for A to buy,
  - otherwise wait to find out what will happen
- At Y:
  - if no note A, safe for B to buy
  - Otherwise, A is either buying or waiting for B to quit

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.15

## Review: Solution #3 discussion

- Our solution protects a single "Critical-Section" piece of code for each thread:

```
if (noMilk) {
  buy milk;
}
```

- Solution #3 works, but it's really unsatisfactory
  - Really complex – even for this simple an example
    - » Hard to convince yourself that this really works
  - A's code is different from B's – what if lots of threads?
    - » Code would have to be slightly different for each thread
  - While A is waiting, it is consuming CPU time
    - » This is called "busy-waiting"
- There's a better way
  - Have hardware provide better (higher-level) primitives than atomic load and store
  - Build even higher-level programming abstractions on this new hardware support

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.16

## High-Level Picture

- The abstraction of threads is good:
  - Maintains sequential execution model
  - Allows simple parallelism to overlap I/O and computation
- Unfortunately, still too complicated to access state shared between threads
  - Consider “too much milk” example
  - Implementing a concurrent program with only loads and stores would be tricky and error-prone
- Today, we’ll start implementing higher-level operations on top of atomic operations provided by hardware
  - Develop a “synchronization toolbox”
  - Explore some common programming paradigms



9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.17

## Too Much Milk: Solution #4

- Suppose we have some sort of implementation of a lock (more in a moment).
  - `Lock.Acquire()` – wait until lock is free, then grab
  - `Lock.Release()` – unlock, waking up anyone waiting
  - These must be atomic operations – if two threads are waiting for the lock and both see it’s free, only one succeeds to grab the lock

- Then, our milk problem is easy:

```
millock.Acquire();
if (nomilk)
    buy milk;
millock.Release();
```

- Once again, section of code between `Acquire()` and `Release()` called a “**Critical Section**”

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.18

## How to implement Locks?

- **Lock**: prevents someone from doing something
  - Lock before entering critical section and before accessing shared data
  - Unlock when leaving, after accessing shared data
  - Wait if locked
    - » Important idea: all synchronization involves waiting
    - » Should *sleep* if waiting for a long time
- Atomic Load/Store: get solution like Milk #3
  - Pretty complex and error prone
- Hardware Lock instruction
  - Is this a good idea?
  - What about putting a task to sleep?
    - » How do you handle the interface between the hardware and scheduler?
  - Complexity?
    - » Each feature makes hardware more complex and slow



9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.19

## Naïve use of Interrupt Enable/Disable

- How can we build multi-instruction atomic operations?
  - Recall: dispatcher gets control in two ways.
    - » Internal: Thread does something to relinquish the CPU
    - » External: Interrupts cause dispatcher to take CPU
  - On a uniprocessor, can avoid context-switching by:
    - » Avoiding internal events (although virtual memory tricky)
    - » Preventing external events by disabling interrupts
- Consequently, naïve Implementation of locks:

```
LockAcquire { disable Ints; }
LockRelease { enable Ints; }
```

- Problems with this approach:

- **Can’t let user do this!** Consider following:

```
LockAcquire();
While(TRUE) {;
```

- Real-Time system—no guarantees on timing!
  - » Critical Sections might be arbitrarily long
- What happens with I/O or other important events?
  - » “Reactor about to meltdown. Help?”



9/12/11


Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.20

## Better Implementation of Locks by Disabling Interrupts

- Key idea: maintain a lock variable and impose mutual exclusion only during operations on that variable

```

int value = FREE; 

Acquire() {
  disable interrupts;
  if (value == BUSY) {
    put thread on wait queue;
    Go to sleep();
    // Enable interrupts?
  } else {
    value = BUSY;
  }
  enable interrupts;
}

Release() {
  disable interrupts;
  if (anyone on wait queue) {
    take thread off wait queue;
    Place on ready queue;
  } else {
    value = FREE;
  }
  enable interrupts;
}

```

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.21

5min Break

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.22

## New Lock Implementation: Discussion

- Why do we need to disable interrupts at all?
  - Avoid interruption between checking and setting lock value
  - Otherwise two threads could think that they both have lock

```

Acquire() {
  disable interrupts;
  if (value == BUSY) {
    put thread on wait queue;
    Go to sleep();
    // Enable interrupts?
  } else {
    value = BUSY;
  }
  enable interrupts;
}

```

} Critical Section

- Note: unlike previous solution, the critical section (inside `Acquire()`) is very short
  - User of lock can take as long as they like in their own critical section: doesn't impact global machine behavior
  - Critical interrupts taken in time

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.23

## Interrupt re-enable in going to sleep

- What about re-enabling ints when going to sleep?

```

Acquire() {
  disable interrupts;
  if (value == BUSY) {
    put thread on wait queue;
    Go to sleep();
  } else {
    value = BUSY;
  }
  enable interrupts;
}

```

Enable Position             
 Enable Position             
 Enable Position           

- Before Putting thread on the wait queue?
  - Release can check the queue and not wake up thread
- After putting the thread on the wait queue
  - Release puts the thread on the ready queue, but the thread still thinks it needs to go to sleep
  - Misses wakeup and still holds lock (deadlock!)
- Want to put it after `sleep()`. But – how?

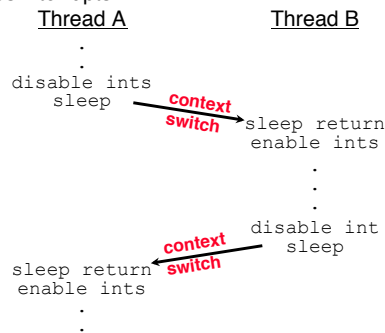
9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.24

## How to Re-enable After Sleep()?

- Since ints are disabled when you call sleep:
  - Responsibility of the next thread to re-enable ints
  - When the sleeping thread wakes up, returns to acquire and re-enables interrupts



9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.25

## Atomic Read-Modify-Write instructions

- Problems with previous solution:
  - Can't give lock implementation to users
  - Doesn't work well on multiprocessor
    - » Disabling interrupts on all processors requires messages and would be very time consuming
- Alternative: atomic instruction sequences
  - These instructions read a value from memory and write a new value atomically
  - Hardware is responsible for implementing this correctly
    - » on both uniprocessors (not too hard)
    - » and multiprocessors (requires help from cache coherence protocol)
  - Unlike disabling interrupts, can be used on both uniprocessors and multiprocessors

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.26

## Examples of Read-Modify-Write

- `test&set (&address) { /* most architectures */`  
`result = M[address];`  
`M[address] = 1;`  
`return result;`  
`}`
- `swap (&address, register) { /* x86 */`  
`temp = M[address];`  
`M[address] = register;`  
`register = temp;`  
`}`
- `compare&swap (&address, reg1, reg2) { /* 68000 */`  
`if (reg1 == M[address]) {`  
`M[address] = reg2;`  
`return success;`  
`} else {`  
`return failure;`  
`}`  
`}`

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.27

## Implementing Locks with test&set

- Simple solution:

```

int value = 0; // Free
Acquire() {
  while (test&set(value));
}
Release() {
  value = 0;
}

```

```

test&set (&address) {
  result = M[address];
  M[address] = 1;
  return result;
}

```

- Simple explanation:

- If lock is free, test&set reads 0 and sets value=1, so lock is now busy. It returns 0 so while exits.
- If lock is busy, test&set reads 1 and sets value=1 (no change). It returns 1, so while loop continues
- When we set value = 0, someone else can get lock

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.28

## Problem: Busy-Waiting for Lock

- Positives for this solution
  - Machine can receive interrupts
  - User code can use this lock
  - Works on a multiprocessor
- Negatives
  - This is very inefficient because the busy-waiting thread will consume cycles waiting
  - Waiting thread may take cycles away from thread holding lock (no one wins!)
  - **Priority Inversion:** If busy-waiting thread has higher priority than thread holding lock  $\Rightarrow$  no progress!
- Priority Inversion problem with original Martian rover
- For semaphores and monitors, waiting thread may wait for an arbitrary length of time!
  - Thus even if busy-waiting was OK for locks, definitely not ok for other primitives
  - Homework/exam solutions should not have busy-waiting!



9/12/11


Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.29

## Better Locks using test&set

- Can we build test&set locks without busy-waiting?
  - Can't entirely, but can minimize!
  - Idea: only busy-wait to atomically check lock value

```
int guard = 0;
int value = FREE;
```



```
Acquire() {
    // Short busy-wait time
    while (test&set(guard));
    if (value == BUSY) {
        put thread on wait queue;
        go to sleep() & guard = 0;
    } else {
        value = BUSY;
        guard = 0;
    }
}

Release() {
    // Short busy-wait time
    while (test&set(guard));
    if anyone on wait queue {
        take thread off wait queue;
        Place on ready queue;
    } else {
        value = FREE;
    }
    guard = 0;
}
```

- Note: sleep has to be sure to reset the guard variable
  - Why can't we do it just before or just after the sleep?

9/12/11


Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.30

## Better Locks using test&set

- Compare to “disable interrupt” solution

```
int value = FREE;
```



```
Acquire() {
    disable interrupts;
    if (value == BUSY) {
        put thread on wait queue;
        Go to sleep();
        // Enable interrupts?
    } else {
        value = BUSY;
    }
    enable interrupts;
}

Release() {
    disable interrupts;
    if (anyone on wait queue) {
        take thread off wait queue;
        Place on ready queue;
    } else {
        value = FREE;
    }
    enable interrupts;
}
```

- Basically replace
  - `disable interrupts`  $\rightarrow$  `while (test&set(guard));`
  - `enable interrupts`  $\rightarrow$  `guard = 0;`

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.31

## Higher-level Primitives than Locks

- Synchronization is a way of coordinating multiple concurrent activities that are using shared state
  - This lecture and the next presents a couple of ways of structuring the sharing
  - What is the right abstraction for synchronizing threads that share memory?
  - Want as high a level primitive as possible
- Good primitives and practices important!
  - Since execution is not entirely sequential, really hard to find bugs, since they happen rarely
  - UNIX is pretty stable now, but up until about mid-80s (10 years after started), systems running UNIX would crash every week or so – concurrency bugs

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.32



## Semaphores



- Semaphores are a kind of generalized locks
  - First defined by Dijkstra in late 60s
  - Main synchronization primitive used in original UNIX
- Definition: a Semaphore has a non-negative integer value and supports the following two operations:
  - **P()**: an atomic operation that waits for semaphore to become positive, then decrements it by 1
    - » Think of this as the wait() operation
  - **V()**: an atomic operation that increments the semaphore by 1, waking up a waiting P, if any
    - » Think of this as the signal() operation
  - Note that **P()** stands for “*proberen*” (to test) and **V()** stands for “*verhogen*” (to increment) in Dutch

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.33

## Semaphores Like Integers Except

- Semaphores are like integers, except
  - No negative values
  - Only operations allowed are P and V – can’t read or write value, except to set it initially
  - Operations must be atomic
    - » Two P’s together can’t decrement value below zero
    - » Similarly, thread going to sleep in P won’t miss wakeup from V – even if they both happen at same time
- Semaphore from railway analogy
  - Here is a semaphore initialized to 2 for resource control:



9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.34

## Two Uses of Semaphores

- Mutual Exclusion (initial value = 1)
  - Also called “Binary Semaphore”.
  - Can be used for mutual exclusion:

```
semaphore.P();
// Critical section goes here
semaphore.V();
```
- Scheduling Constraints (initial value = 0)
  - Locks are fine for mutual exclusion, but what if you want a thread to wait for something?
  - Example: suppose you had to implement ThreadJoin which must wait for thread to terminate:

```
Initial value of semaphore = 0
ThreadJoin {
    semaphore.P();
}
ThreadFinish {
    semaphore.V();
}
```

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.35

## Summary

- Important concept: Atomic Operations
  - An operation that runs to completion or not at all
  - These are the primitives on which to construct various synchronization primitives
- Talked about hardware atomicity primitives:
  - Disabling of Interrupts, test&set
- Showed several constructions of Locks
  - Must be very careful not to waste/tie up machine resources
    - » Shouldn’t disable interrupts for long
    - » Shouldn’t spin wait for long
  - Key idea: Separate lock variable, use hardware mechanisms to protect modifications of that variable
- Semaphores: Higher level constructs that are harder to “screw up”

9/12/11

Anthony D. Joseph and Ion Stoica CS162 ©UCB Fall 2011

Lec 4.36