

CS162 Operating Systems and Systems Programming Lecture 2

Synchronization

January 26th, 2011
Ion Stoica
<http://inst.eecs.berkeley.edu/~cs162>

Execution Stack Example

```

addrX: A(int tmp) {
.       if (tmp<2)
.       B();
addrY: printf(tmp);
.       }
.       B() {
.       C();
addrU: }
.       C() {
.       A(2);
addrV: }
.       A(1);
addrZ: exit;

```

- Stack holds function arguments, return address
- Permits recursive execution
- Crucial to modern languages

1/19/10

Ion Stoica CS162 ©UCB Spring 2011


Lec 1.2

Execution Stack Example

```

addrX: A(int tmp) {
.       if (tmp<2)
.       B();
addrY: printf(tmp);
.       }
.       B() {
.       C();
addrU: }
.       C() {
.       A(2);
addrV: }
.       A(1);
addrZ: exit;

```



- Stack holds function arguments, return address
- Permits recursive execution
- Crucial to modern languages

1/19/10

Ion Stoica CS162 ©UCB Spring 2011


Lec 1.3

Execution Stack Example

```

addrX: A(int tmp) {
.       if (tmp<2)
.       B();
addrY: printf(tmp);
.       }
.       B() {
.       C();
addrU: }
.       C() {
.       A(2);
addrV: }
.       A(1);
addrZ: exit;

```



- Stack holds function arguments, return address
- Permits recursive execution
- Crucial to modern languages

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.4

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY:   printf(tmp);
      .   }
      .   B() {
      .   C();
addrU:   }
      .   C() {
      .   A(2);
addrV:   }
      .   A(1);
addrZ:   exit;

```

Stack Growth

- Stack holds function arguments, return address
- Permits recursive execution
- Crucial to modern languages

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.5

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY:   printf(tmp);
      .   }
      .   B() {
      .   C();
addrU:   }
      .   C() {
      .   A(2);
addrV:   }
      .   A(1);
addrZ:   exit;

```

Stack Growth

- Stack holds function arguments, return address
- Permits recursive execution
- Crucial to modern languages

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.6

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY:   printf(tmp);
      .   }
      .   B() {
      .   C();
addrU:   }
      .   C() {
      .   A(2);
addrV:   }
      .   A(1);
addrZ:   exit;

```

Stack Growth

- Stack holds function arguments, return address
- Permits recursive execution
- Crucial to modern languages

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.7

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY:   printf(tmp);
      .   }
      .   B() {
      .   C();
addrU:   }
      .   C() {
      .   A(2);
addrV:   }
      .   A(1);
addrZ:   exit;

```

Stack Growth

Output:

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.8

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY: printf(tmp);
      .   }
      .   B() {
      .   C();
addrU: }
      .   C() {
      .   A(2);
addrV: }
      .   A(1);
addrZ: exit;

```

Stack Pointer →

Stack Growth ↓

Output:
2

1/19/10
Ion Stoica CS162 ©UCB Spring 2011
Lec 1.9

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY: printf(tmp);
      .   }
      .   B() {
      .   C();
addrU: }
      .   C() {
      .   A(2);
addrV: }
      .   A(1);
addrZ: exit;

```

Stack Pointer →

Stack Growth ↓

Output:
2

1/19/10
Ion Stoica CS162 ©UCB Spring 2011
Lec 1.10

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY: printf(tmp);
      .   }
      .   B() {
      .   C();
addrU: }
      .   C() {
      .   A(2);
addrV: }
      .   A(1);
addrZ: exit;

```

Stack Pointer →

Stack Growth ↓

Output:
2

1/19/10
Ion Stoica CS162 ©UCB Spring 2011
Lec 1.11

Execution Stack Example

```

addrX: A(int tmp) {
      .   if (tmp<2)
      .   B();
addrY: printf(tmp);
      .   }
      .   B() {
      .   C();
addrU: }
      .   C() {
      .   A(2);
addrV: }
      .   A(1);
addrZ: exit;

```

Stack Pointer →

Stack Growth ↓

Output:
2
1

1/19/10
Ion Stoica CS162 ©UCB Spring 2011
Lec 1.12

Execution Stack Example

```

addrX: A(int tmp) {
        .   if (tmp<2)
        .   .   B();
addrY:   printf(tmp);
        .   }
        .   B() {
        .   .   C();
addrU:   }
        .   C() {
        .   .   A(2);
addrV:   }
        .   A(1);
addrZ:   exit;
    
```

Output:
2
1

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.13

Classification

# threads Per AS:	# of addr spaces:	One	Many
One		MS/DOS, early Macintosh	Traditional UNIX
Many		Embedded systems (Geoworks, VxWorks, JavaOS, etc) JavaOS, Pilot(PC)	Mach, OS/2, Linux, Win 95?, Mac OS X, Win NT to XP, Solaris, HP-UX

- Real operating systems have either
 - One or many address spaces
 - One or many threads per address space
- Did Windows 95/98/ME have real memory protection?
 - No: Users could overwrite process tables / system DLLs

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.14

Goals for Today

- Thread Dispatching
- Cooperating Threads
- Concurrency examples
- Need for synchronization

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

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.15

Single-Threaded Example

- Imagine the following C program:

```

main() {
    ComputePI("pi.txt");
    PrintClassList("clist.txt");
}
    
```

- What is the behavior here?
 - Program would never print out class list
 - Why? ComputePI would never finish

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.16

Use of Threads

- Version of program with Threads:

```
main() {
    CreateThread(ComputePI("pi.txt"));
    CreateThread(PrintClassList("clist.txt"));
}
```

- What does "CreateThread" do?
 - Start independent thread running given procedure
- What is the behavior here?
 - Now, you would actually see the class list
 - This *should* behave as if there are two separate CPUs




1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.17

Memory Footprint of Two-Thread Example

- If we stopped this program and examined it with a debugger, we would see
 - Two sets of CPU registers
 - Two sets of Stacks



- Questions:
 - How do we position stacks relative to each other?
 - What maximum size should we choose for the stacks?
 - What happens if threads violate this?
 - How might you catch violations?

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.18

Per Thread State


- Each Thread has a *Thread Control Block* (TCB)
 - Execution State: CPU registers, program counter, pointer to stack
 - Scheduling info: State, priority, CPU time
 - Accounting Info
 - Various Pointers (for implementing scheduling queues)
 - Pointer to enclosing process? (PCB)?
 - Etc (add stuff as you find a need)
- OS Keeps track of TCBs in protected memory
 - In Array, or Linked List, or ...

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.19

Lifecycle of a Thread (or Process)



- As a thread executes, it changes state:
 - new**: The thread is being created
 - ready**: The thread is waiting to run
 - running**: Instructions are being executed
 - waiting**: Thread waiting for some event to occur
 - terminated**: The thread has finished execution
- "Active" threads are represented by their TCBs
 - TCBs organized into queues based on their state

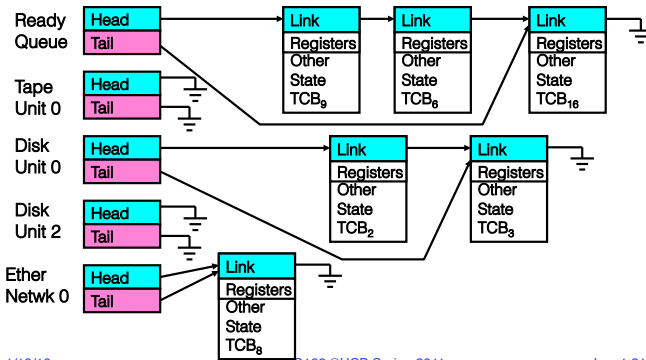
1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.20

Ready Queue And Various I/O Device Queues

- Thread not running ⇒ TCB is in some scheduler queue
 - Separate queue for each device/signal/condition
 - Each queue can have a different scheduler policy



1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.21

Dispatch Loop

- Conceptually, the dispatching loop of the operating system looks as follows:

```

Loop {
    RunThread();
    ChooseNextThread();
    SaveStateOfCPU(curTCB);
    LoadStateOfCPU(newTCB);
}
    
```

- This is an *infinite* loop
 - One could argue that this is all that the OS does

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.22

Running a thread

Consider first portion: `RunThread()`

- How do I run a thread?
 - Load its state (registers, PC, stack pointer) into CPU
 - Load environment (virtual memory space, etc)
 - Jump to the PC
- How does the dispatcher get control back?
 - Internal events: thread returns control voluntarily
 - External events: thread gets *preempted*

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.23

Review: Yielding through Internal Events

- Blocking on I/O
 - The act of requesting I/O implicitly yields the CPU
 - Waiting on a “signal” from other thread
 - Thread asks to wait and thus yields the CPU
 - Thread executes a `yield()`
 - Thread volunteers to give up CPU
- ```

computePI() {
 while(TRUE) {
 ComputeNextDigit();
 yield();
 }
}


```
- Note that `yield()` must be called by programmer frequently enough!

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.24

## Review: Stack for Yielding Thread



- How do we run a new thread?

```
run_new_thread() {
 newThread = PickNewThread();
 switch(curThread, newThread);
 ThreadHouseKeeping(); /* deallocates finished threads */
}
```


- Finished thread not killed right away. Why?
  - Move them in "exit" state
  - ThreadHouseKeeping() deallocates finished threads

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.25

## Review: Stack for Yielding Thread



- How do we run a new thread?

```
run_new_thread() {
 newThread = PickNewThread();
 switch(curThread, newThread);
 ThreadHouseKeeping(); /* deallocates finished threads */
}
```

- How does dispatcher switch to a new thread?
  - Save anything next thread may trash: PC, regs, stack
  - Maintain isolation for each thread

1/19/10

Ion Stoica CS162 ©UCB Spring 2011


Lec 1.26

## Review: Two Thread Yield Example

- Consider the following code blocks:

```
proc A() {
 B();
}
proc B() {
 while(TRUE) {
 yield();
 }
}
```

- Suppose we have 2 threads:
  - Threads S and T




1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.27

## Detour: Interrupt Controller



- Interrupts invoked with interrupt lines from devices
- Interrupt controller chooses interrupt request to honor
  - Mask enables/disables interrupts
  - Priority encoder picks highest enabled interrupt
  - Software Interrupt Set/Cleared by Software
  - Interrupt identity specified with ID line
- CPU can disable all interrupts with internal flag
- Non-maskable interrupt line (NMI) can't be disabled


1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.28

## Review: Preemptive Multithreading

- Use the timer interrupt to force scheduling decisions



- Timer Interrupt routine:
 

```
TimerInterrupt() {
 DoPeriodicHouseKeeping();
 run_new_thread();
}
```
- This is often called **preemptive multithreading**, since threads are preempted for better scheduling
  - Solves problem of user who doesn't insert `yield()`;

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.29

## Project Signup

- Project Signup: Watch "Group/Section Signup" Link
  - 4-5 members to a group
    - Everyone in group must be able to *actually* attend same section
    - The sections assigned to you by Telebears are temporary!
  - Only submit once per group!
    - Everyone in group must have logged into their cs162-xx accounts once before you register the group
    - Make sure that you select at least 2 potential sections
      - Hard deadline: due Friday (1/28) by 11:59pm**
- Sections:
  - Watch for section assignments next Monday/Tuesday
  - Attend new sections next week

| Section | Time             | Location        | TA                      |
|---------|------------------|-----------------|-------------------------|
| 101     | Th 10:00A-11:00A | 3105 Etcheverry | Jorge Ortiz             |
| 102     | Th 11:00A-12:00P | 4 Evans         |                         |
| 104     | Th 1:00P-2:00P   | 85 Evans        | Stephen Dawson-Haggerty |
| 105     | Th 2:00P-3:00P   | B56 Hildebrand  |                         |
| 103     | Th 3:00P-4:00P   | 4 Evans         | David Zhu               |
| 106     | Th 4:00P-5:00P   | 320 Soda        |                         |

## Other Announcements...

- Exam date: May 13, 8-11am
- A few students that take CS 184 have conflicts...
- We'll accommodate these conflicts
- Exam **only** for people that have conflict (tentative): May 13, 1-4pm

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.31

**5min Break**

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.32



## Why allow cooperating threads?


- People cooperate; computers help/enhance people's lives, so computers must cooperate
  - By analogy, the non-reproducibility/non-determinism of people is a notable problem for "carefully laid plans"
- Advantage 1: Share resources
  - One computer, many users
  - One bank balance, many ATMs
    - » What if ATMs were only updated at night?
  - Embedded systems (robot control: coordinate arm & hand)
- Advantage 2: Speedup
  - Overlap I/O and computation
  - Multiprocessors – chop up program into parallel pieces
- Advantage 3: Modularity
  - Chop large problem up into simpler pieces
    - » To compile, for instance, gcc calls `cpp | cc1 | cc2 | as | ld`
    - » Makes system easier to extend

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.33

## Threaded Web Server



- Multithreaded version:
 

```
serverLoop() {
 connection = AcceptCon();
 ThreadCreate(ServiceWebPage(), connection);
}
```
- Advantages of threaded version:
  - Can share file caches kept in memory, results of CGI scripts, other things
  - Threads are *much* cheaper to create than processes, so this has a lower per-request overhead
- What if too many requests come in at once?


1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.34

## Thread Pools

- Problem with previous version: Unbounded Threads
  - When web-site becomes too popular – throughput sinks
- Instead, allocate a bounded "pool" of threads, representing the maximum level of multiprogramming



```

master() {
 allocThreads(slave, queue);
 while(TRUE) {
 con=AcceptCon();
 Enqueue(queue, con);
 wakeUp(queue);
 }
}

slave(queue) {
 while(TRUE) {
 con=Dequeue(queue);
 if (con==null)
 sleepOn(queue);
 else
 ServiceWebPage(con);
 }
}


```

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.35

## ATM Bank Server



- ATM server problem:
  - Service a set of requests
  - Do so without corrupting database
  - Don't hand out too much money

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.36

## ATM bank server example

- Suppose we wanted to implement a server process to handle requests from an ATM network:

```
BankServer() {
 while (TRUE) {
 ReceiveRequest(&op, &acctId, &amount);
 ProcessRequest(op, acctId, amount);
 }
}

ProcessRequest(op, acctId, amount) {
 if (op == deposit) Deposit(acctId, amount);
 else if ...
}

Deposit(acctId, amount) {
 acct = GetAccount(acctId); /* may use disk I/O */
 acct->balance += amount;
 StoreAccount(acct); /* Involves disk I/O */
}
```

- How could we speed this up?
  - More than one request being processed at once
  - Event driven (overlap computation and I/O)
  - Multiple threads (multi-proc, or overlap comp and I/O)

1/19/10

Ion Stoica CS162 @UCB Spring 2011

Lec 1.37

## Event Driven Version of ATM server

- Suppose we only had one CPU
  - Still like to overlap I/O with computation
  - Without threads, we would have to rewrite in event-driven style
- Example

```
BankServer() {
 while(TRUE) {
 event = WaitForNextEvent();
 if (event == ATMRequest)
 StartOnRequest();
 else if (event == AcctAvail)
 ContinueRequest();
 else if (event == AcctStored)
 FinishRequest();
 }
}
```

- What if we missed a blocking I/O step?
- What if we have to split code into hundreds of pieces which could be blocking?
- This technique is used for graphical programming

1/19/10

Ion Stoica CS162 @UCB Spring 2011

Lec 1.38

## Can Threads Make This Easier?

- Threads yield overlapped I/O and computation without “deconstructing” code into non-blocking fragments
  - One thread per request
- Requests proceeds to completion, blocking as required:

```
Deposit(acctId, amount) {
 acct = GetAccount(acctId); /* May use disk I/O */
 acct->balance += amount;
 StoreAccount(acct); /* Involves disk I/O */
}
```

- Unfortunately, shared state can get corrupted:

| <u>Thread 1</u>         | <u>Thread 2</u>         |
|-------------------------|-------------------------|
| load r1, acct->balance  | load r1, acct->balance  |
|                         | add r1, amount2         |
|                         | store r1, acct->balance |
| add r1, amount1         |                         |
| store r1, acct->balance |                         |

1/19/10

Ion Stoica CS162 @UCB Spring 2011

Lec 1.39

## Problem is at the lowest level

- Most of the time, threads are working on separate data, so scheduling doesn't matter:

| <u>Thread A</u> | <u>Thread B</u> |
|-----------------|-----------------|
| x = 1;          | y = 2;          |

- However, What about (Initially, y = 12):

| <u>Thread A</u> | <u>Thread B</u> |
|-----------------|-----------------|
| x = 1;          | y = 2;          |
| x = y+1;        | y = y*2;        |

- What are the possible values of x?

| <u>Thread A</u> | <u>Thread B</u> |
|-----------------|-----------------|
| x = 1;          |                 |
| x = y+1;        |                 |
|                 | y = 2;          |
|                 | y = y*2         |

x=13

1/19/10

Ion Stoica CS162 @UCB Spring 2011

Lec 1.40

### Problem is at the lowest level

- Most of the time, threads are working on separate data, so scheduling doesn't matter:

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| x = 1;          | y = 2;          |

- However, What about (Initially, y = 12):

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| x = 1;          | y = 2;          |
| x = y+1;        | y = y*2;        |

- What are the possible values of x?

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| y = 2;          |                 |
| y = y*2;        |                 |
| x = 1;          |                 |
| x = y+1;        |                 |

x=5

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.41

### Problem is at the lowest level

- Most of the time, threads are working on separate data, so scheduling doesn't matter:

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| x = 1;          | y = 2;          |

- However, What about (Initially, y = 12):

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| x = 1;          | y = 2;          |
| x = y+1;        | y = y*2;        |

- What are the possible values of x?

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| y = 2;          |                 |
| x = 1;          |                 |
| x = y+1;        |                 |
| y = y*2;        |                 |

x=3

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.42

### Problem is at the lowest level

- Most of the time, threads are working on separate data, so scheduling doesn't matter:

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| x = 1;          | y = 2;          |

- However, What about (Initially, y = 12):

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| x = 1;          | y = 2;          |
| x = y+1;        | y = y*2;        |

- What are the possible values of x?
- Or, what are the possible values of x below?

|                 |                 |
|-----------------|-----------------|
| <u>Thread A</u> | <u>Thread B</u> |
| x = 1;          | x = 2;          |

- x could be 1 or 2 (non-deterministic!)

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.43

### Atomic Operations

- To understand a concurrent program, we need to know what the underlying indivisible operations are!
- Atomic Operation:** an operation that always runs to completion or not at all
  - It is *indivisible*: it cannot be stopped in the middle and state cannot be modified by someone else in the middle
  - Fundamental building block – if no atomic operations, then have no way for threads to work together
- On most machines, memory references and assignments (i.e. loads and stores) of words are atomic
- Many instructions are not atomic
  - Double-precision floating point store often not atomic
  - VAX and IBM 360 had an instruction to copy a whole array

1/19/10 Ion Stoica CS162 ©UCB Spring 2011 Lec 1.44

## Correctness Requirements

- Threaded programs must work for all interleavings of thread instruction sequences
  - Cooperating threads inherently non-deterministic and non-reproducible
  - Really hard to debug unless carefully designed!
- Example: Therac-25
  - Machine for radiation therapy
    - Software control of electron accelerator and electron beam/X-ray production
    - Software control of dosage
  - Software errors caused the death of several patients
    - A series of race conditions on shared variables and poor software design
    - "They determined that data entry speed during editing was the key factor in producing the error condition: If the prescription data was edited at a fast pace, the overdose occurred."




Figure 1. Typical Therac-25 facility


1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.45

## Space Shuttle Example

- Original Space Shuttle launch aborted 20 minutes before scheduled launch
- Shuttle has five computers:
  - Four run the "Primary Avionics Software System" (PASS)
    - Asynchronous and real-time
    - Runs all of the control systems
    - Results synchronized and compared every 3 to 4 ms
  - The Fifth computer is the "Backup Flight System" (BFS)
    - stays synchronized in case it is needed
    - Written by completely different team than PASS
- Countdown aborted because BFS disagreed with PASS
  - A 1/67 chance that PASS was out of sync one cycle
  - Bug due to modifications in **initialization** code of PASS
    - A delayed init request placed into timer queue
    - As a result, timer queue not empty at expected time to force use of hardware clock
  - Bug not found during extensive simulation



1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.46

## 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? Could be either
- 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?

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.47

## Summary

- Concurrent threads are a very useful abstraction
  - Allow transparent overlapping of computation and I/O
  - Allow use of parallel processing when available
- Concurrent threads introduce problems when accessing shared data
  - Programs must be insensitive to arbitrary interleavings
  - Without careful design, shared variables can become completely inconsistent
- 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

1/19/10

Ion Stoica CS162 ©UCB Spring 2011

Lec 1.48