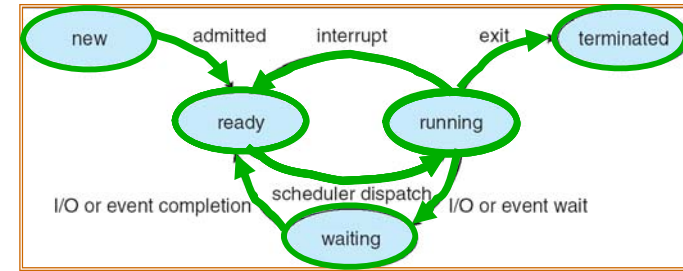


# CS162 Operating Systems and Systems Programming Lecture 6

## Concurrency (Continued), Synchronization (Start)

February 9<sup>th</sup>, 2015  
Prof. John Kubiatowicz  
<http://cs162.eecs.Berkeley.edu>

### Recall: Lifecycle of a Process



- As a process executes, it changes state:
  - new**: The process is being created
  - ready**: The process is waiting to run
  - running**: Instructions are being executed
  - waiting**: Process waiting for some event to occur
  - terminated**: The process has finished execution


### Recall: Use of Threads

- Version of program with Threads (loose syntax):


```

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

- What does "ThreadFork()" 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



### Recall: Thread Abstraction




- Infinite number of processors
- Threads execute with variable speed
  - Programs must be designed to work with any schedule

## Recall: Execution Stack Example

```

A(int tmp) {
  if (tmp<2)
    B();
  printf(tmp);
}
B() {
  C();
}
C() {
  A(2);
}
A(1);
    
```



- Stack holds temporary results
- Permits recursive execution
- Crucial to modern languages

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.5

## MIPS: Software conventions for Registers

0	zero	constant 0	16	s0	callee saves
1	at	reserved for assembler	... (callee must save)		
2	v0	expression evaluation &	23	s7	
3	v1	function results	24	t8	temporary (cont'd)
4	a0	arguments	25	t9	
5	a1		26	k0	reserved for OS kernel
6	a2		27	k1	
7	a3		28	gp	Pointer to global area
8	t0	temporary: caller saves	29	sp	Stack pointer
...		(callee can clobber)	30	fp	frame pointer
15	t7		31	ra	Return Address (HW)

- Before calling procedure:
  - Save caller-saves regs
  - Save v0, v1
  - Save ra
- After return, assume
  - Callee-saves reg OK
  - gp, sp, fp OK (restored!)
  - Other things trashed

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015


Lec 6.6

## Recall: Multithreaded stack switching

- Consider the following code blocks:

```

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



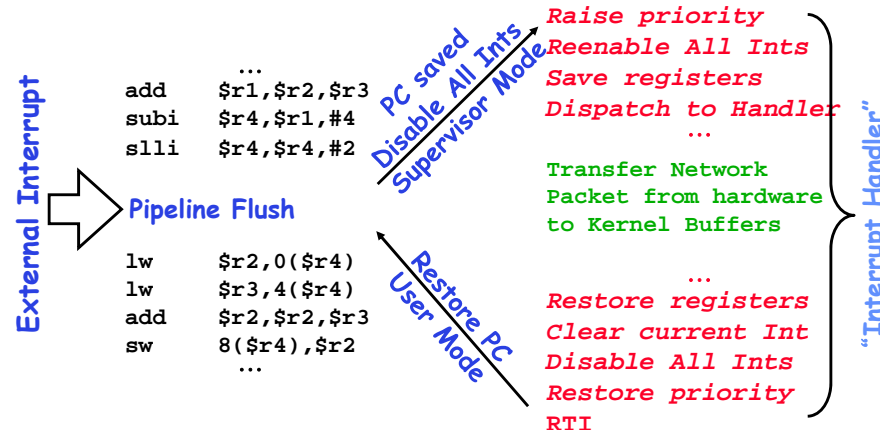
- Suppose we have 2 threads:
  - Threads S and T

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.7

## Example: Network Interrupt



- An interrupt is a hardware-invoked context switch
  - No separate step to choose what to run next
  - Always run the interrupt handler immediately


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.8

## Use of Timer Interrupt to Return Control

- Solution to our dispatcher problem
  - Use the timer interrupt to force scheduling decisions



- Timer Interrupt routine:
 

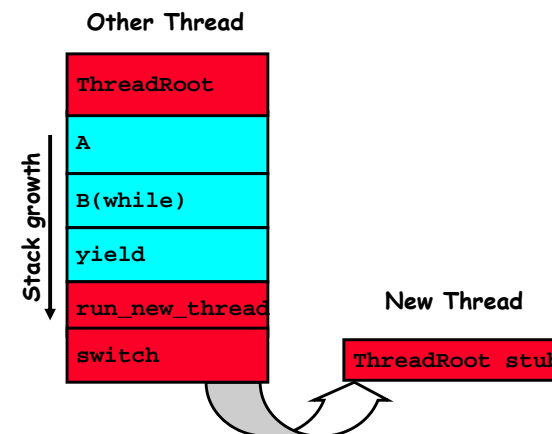
```
TimerInterrupt() {
    DoPeriodicHouseKeeping();
    run_new_thread();
}
```
- I/O interrupt: same as timer interrupt except that DoHousekeeping() replaced by ServiceIO().

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.9

## How does Thread get started?



- Eventually, run\_new\_thread() will select this TCB and return into beginning of ThreadRoot()
  - This really starts the new thread

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015


Lec 6.10

## What does ThreadRoot() look like?

- ThreadRoot() is the root for the thread routine:

```
ThreadRoot() {
    DoStartupHousekeeping();
    UserModeSwitch(); /* enter user mode */
    Call fcnPtr(fcnArgPtr);
    ThreadFinish();
}
```

- Startup Housekeeping
  - Includes things like recording start time of thread
  - Other Statistics



Running Stack

- Stack will grow and shrink with execution of thread
- Final return from thread returns into ThreadRoot() which calls ThreadFinish()
  - ThreadFinish() wake up sleeping threads

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.11

## Examples multithreaded programs

- Embedded systems
  - Elevators, Planes, Medical systems, Wristwatches
  - Single Program, concurrent operations
- Most modern OS kernels
  - Internally concurrent because have to deal with concurrent requests by multiple users
  - But no protection needed within kernel
- Database Servers
  - Access to shared data by many concurrent users
  - Also background utility processing must be done

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.12

## Example multithreaded programs (con't)


- **Network Servers**
  - Concurrent requests from network
  - Again, single program, multiple concurrent operations
  - File server, Web server, and airline reservation systems
- **Parallel Programming (More than one physical CPU)**
  - Split program into multiple threads for parallelism
  - This is called Multiprocessing
- **Some multiprocessors are actually uniprogrammed:**
  - Multiple threads in one address space but one program at a time

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.13

## A typical use case



2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.14

## Some Actual Numbers

- Many processes are multi-threaded, so thread context switches may be either **within-process** or **across-processes**.

Image Name	PID	User Name	CPU	Memory (Private Workin...	Threads	Description
thunderbird.exe *32	5544	jfc	00	422,212 K	28	Thunderbird
firefox.exe *32	6064	jfc	00	362,048 K	49	Firefox
BCU.exe *32	4752	jfc	00	109,012 K	6	Browser Configuration Utility
dwm.exe	4036	jfc	00	105,676 K	5	Desktop Window Manager
POWERPNT.EXE	140	jfc	00	102,204 K	12	Microsoft PowerPoint
explorer.exe	1780	jfc	00	73,244 K	36	Windows Explorer
Dropbox.exe *32	3380	jfc	00	56,792 K	34	Dropbox
CameraHelperShell.exe...	4892	jfc	00	15,068 K	9	Webcam Controller
emacs.exe *32	4856	jfc	00	12,996 K	3	GNU Emacs: The extensible self-doc
FlashPlayerPlugin_11_8...	4260	jfc	00	10,820 K	12	Adobe Flash Player 11.8 r800
nvxdsync.exe	3420	00	00	10,192 K	10	
emacs.exe *32	2736	jfc	00	10,000 K	3	GNU Emacs: The extensible self-doc
BtvStack.exe	2708	ifc	00	9,444 K	43	Bluetooth Stack Server

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.15

## Kernel Use Cases

- Thread for each user process
- Thread for sequence of steps in processing I/O
- Threads for device drivers
- ...

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.16

## Administrivia

- **Group formation: should be completed**
  - Will handle stragglers tonight
- **Project #1: Released!**
  - Technically starts today
  - Autograder should be up by tomorrow.
- **HW1 due today**
  - Must be submitted via the recommended "push" mechanism through git

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.17

## Famous Quote WRT Scheduling: Dennis Richie

```

Dennis Richie,      2230  /*
Unix V6, slp.c:    2231  * If the new process paused because it was
                   2232  * swapped out, set the stack level to the last call
                   2233  * to savu(u_ssav). This means that the return
                   2234  * which is executed immediately after the call to aretu
                   2235  * actually returns from the last routine which did
                   2236  * the savu.
                   2237  *
                   2238  * You are not expected to understand this.
                   2239  */
    
```

*"If the new process paused because it was swapped out, set the stack level to the last call to savu(u\_ssav). This means that the return which is executed immediately after the call to aretu actually returns from the last routine which did the savu."*

*"You are not expected to understand this."*

Source: Dennis Ritchie, Unix V6 slp.c (context-switching code) as per The Unix Heritage Society(tuhs.org); gif by Eddie Koehler.

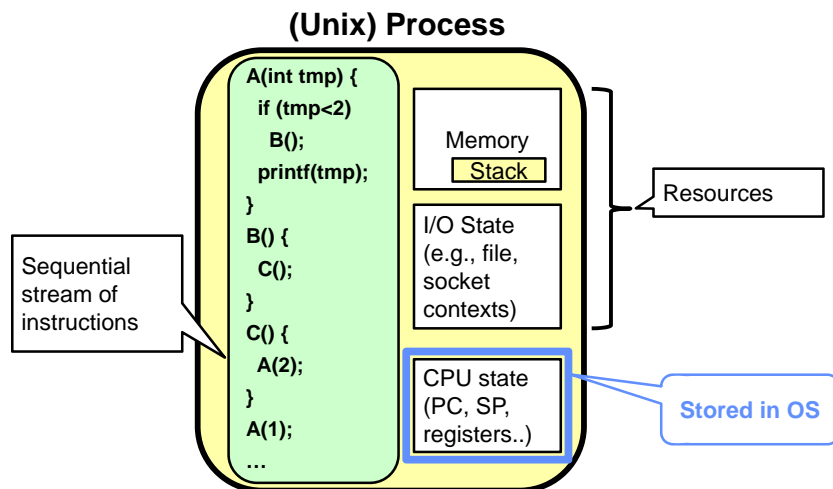
Included by Ali R. Butt in CS3204 from Virginia Tech

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.18

## Putting it together: Process

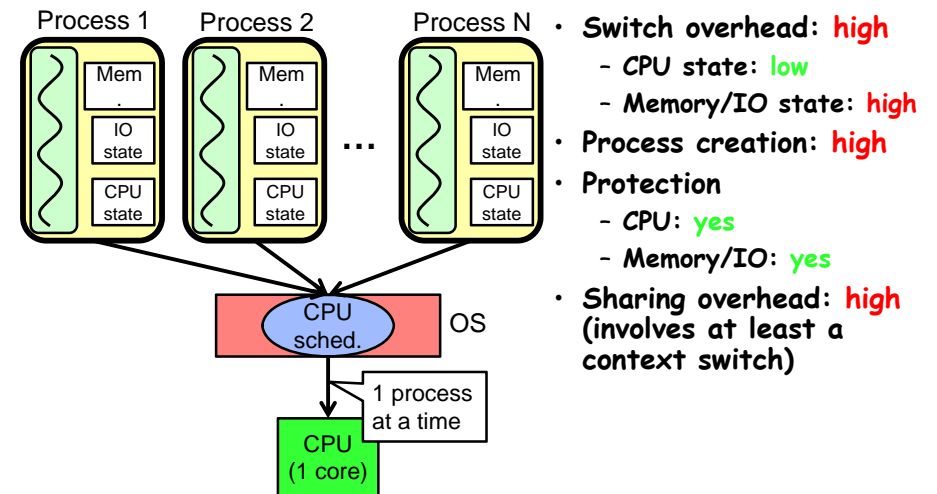


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.19

## Putting it together: Processes




2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.20

## Putting it together: Threads



- **Switch overhead:** **low** (only CPU state)
- **Thread creation:** **low**
- **Protection**
  - CPU: **yes**
  - Memory/IO: **No**
- **Sharing overhead:** **low** (thread switch overhead low)

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.21

## Kernel versus User-Mode threads

- We have been talking about Kernel threads
  - Native threads supported directly by the kernel
  - Every thread can run or block independently
  - One process may have several threads waiting on different things
- Downside of kernel threads: a bit expensive
  - Need to make a crossing into kernel mode to schedule
- Even lighter weight option: User Threads
  - User program provides scheduler and thread package
  - May have several user threads per kernel thread
  - User threads may be scheduled non-preemptively relative to each other (only switch on yield())
  - Cheap
- Downside of user threads:
  - When one thread blocks on I/O, all threads block
  - Kernel cannot adjust scheduling among all threads
  - Option: *Scheduler Activations*
    - » Have kernel inform user level when thread blocks...



2/9/15

Kubiatowicz CS162 ©UCB Spring 2015


Lec 6.22

## Some Threading Models

### Simple One-to-One Threading Model



Many-to-One



Many-to-Many

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.23

## Threads in a Process

- Threads are useful at user-level
  - Parallelism, hide I/O latency, interactivity
- Option A (early Java): user-level library, within a single-threaded process
  - Library does thread context switch
  - Kernel time slices between processes, e.g., on system call I/O
- Option B (SunOS, Linux/Unix variants): green Threads
  - User-level library does thread multiplexing
- Option C (Windows): scheduler activations
  - Kernel allocates processors to user-level library
  - Thread library implements context switch
  - System call I/O that blocks triggers upcall
- Option D (Linux, MacOS, Windows): use kernel threads
  - System calls for thread fork, join, exit (and lock, unlock,...)
  - Kernel does context switching
  - Simple, but a lot of transitions between user and kernel mode

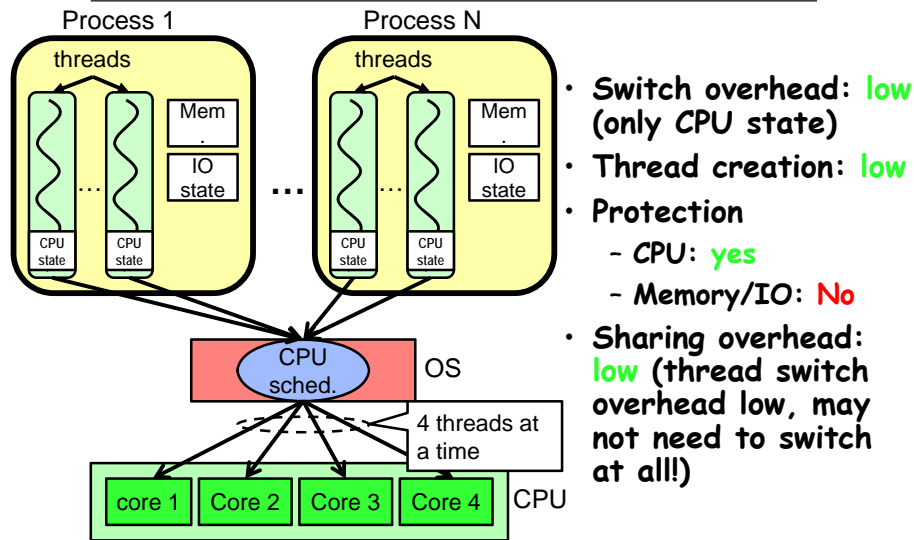
2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.24



## Putting it together: Multi-Cores



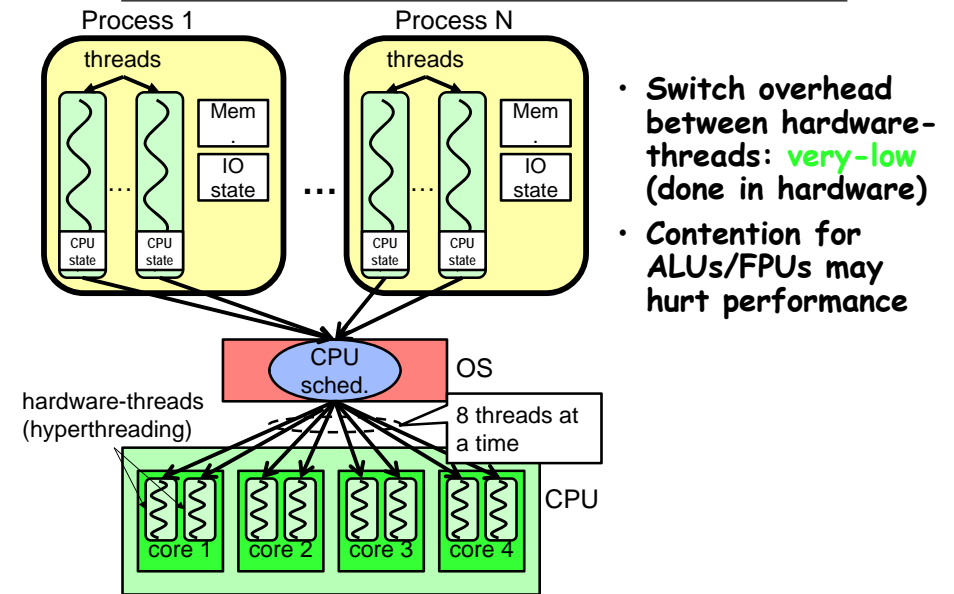
- Switch overhead: **low** (only CPU state)
- Thread creation: **low**
- Protection
  - CPU: **yes**
  - Memory/IO: **No**
- Sharing overhead: **low** (thread switch overhead low, may not need to switch at all!)

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.25

## Putting it together: Hyper-Threading



- Switch overhead between hardware-threads: **very-low** (done in hardware)
- Contention for ALUs/FPU's may hurt performance


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.26

## Multiprocessing vs Multiprogramming

- Remember Definitions:
  - Multiprocessing  $\equiv$  Multiple CPUs
  - Multiprogramming  $\equiv$  Multiple Jobs or Processes
  - Multithreading  $\equiv$  Multiple threads per Process
- What does it mean to run two threads "concurrently"?
  - Scheduler is free to run threads in any order and interleaving: FIFO, Random, ...
  - Dispatcher can choose to run each thread to completion or time-slice in big chunks or small chunks

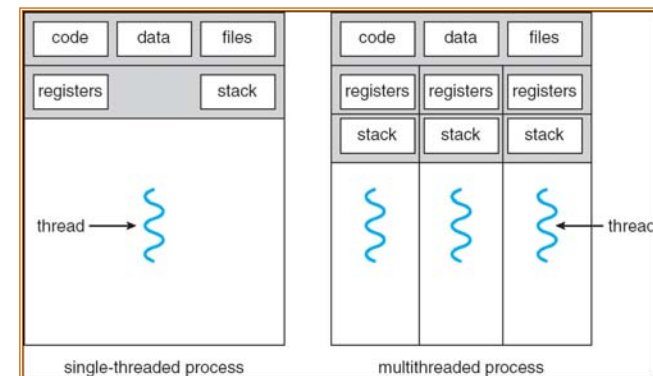


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.27

## Single and Multithreaded Processes

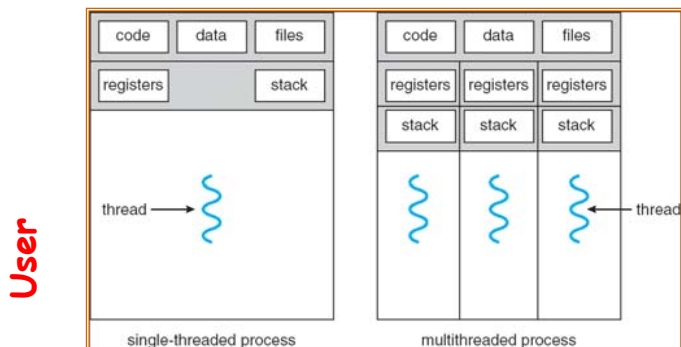


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.28

## Supporting 1T and MT Processes




User

System



## Supporting 1T and MT Processes



User

System



## 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 Windows 9x??? Win NT to XP, Solaris, HP-UX, OS X

- 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

## You are here... why?

- Processes
  - Thread(s) + address space
- Address Space
- Protection
- Dual Mode
- Interrupt handlers
  - Interrupts, exceptions, syscall
- File System
  - Integrates processes, users, cwd, protection
- Key Layers: OS Lib, Syscall, Subsystem, Driver
  - User handler on OS descriptors
- Process control
  - fork, wait, signal, exec
- Communication through sockets
  - Integrates processes, protection, file ops, concurrency
- Client-Server Protocol
- Concurrent Execution: Threads
- Scheduling

OS Concepts (6) intro

Address Space (4)

File Systems (6)

Distributed Systems (8)

Reliability, Security, Cloud (9)

Concurrency (6)



## Perspective on 'groking' 162

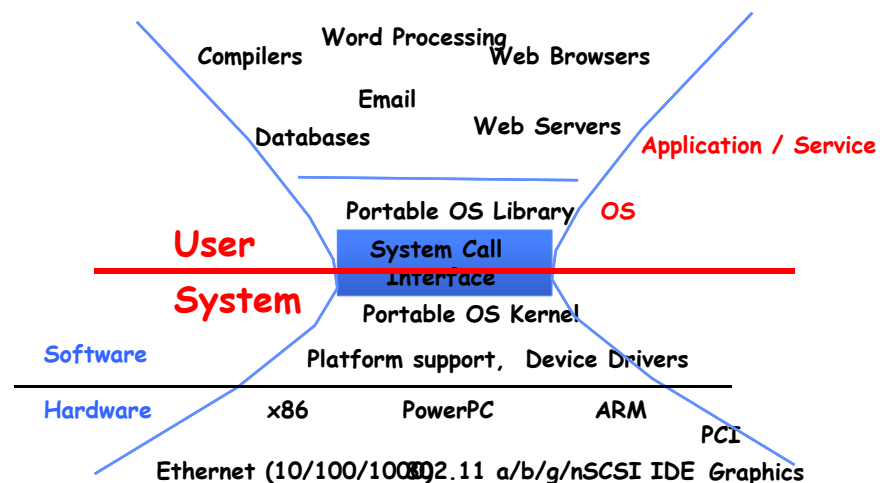
- Historically, OS was the most complex software
  - Concurrency, synchronization, processes, devices, communication, ...
  - Core systems concepts developed there
- Today, many "applications" are complex software systems too
  - These concepts appear there
  - But they are realized out of the capabilities provided by the operating system
- Seek to understand how these capabilities are implemented upon the basic hardware.
- See concepts multiple times from multiple perspectives
  - Lecture provides conceptual framework, integration, examples, ...
  - Book provides a reference with some additional detail
  - Lots of other resources that you need to learn to use
    - » man pages, google, reference manuals, includes (.h)
- Section, Homework and Project provides detail down to the actual code AND direct hands-on experience

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.33

## Operating System as Design

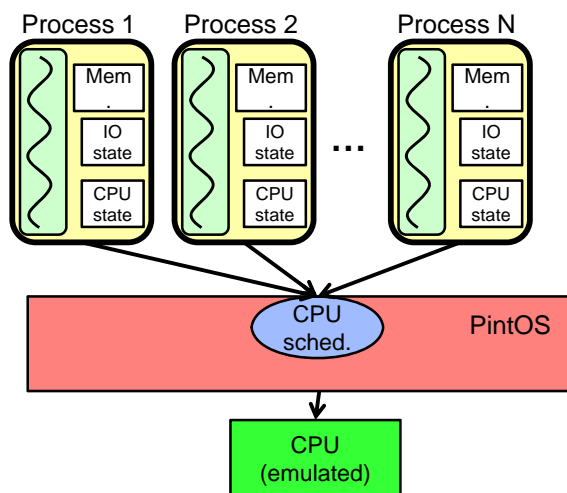


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.34

## Starting today: Pintos Projects



- Groups almost all formed
- Work as one!
- 10x homework
- P1: threads & scheduler
- P2: user process

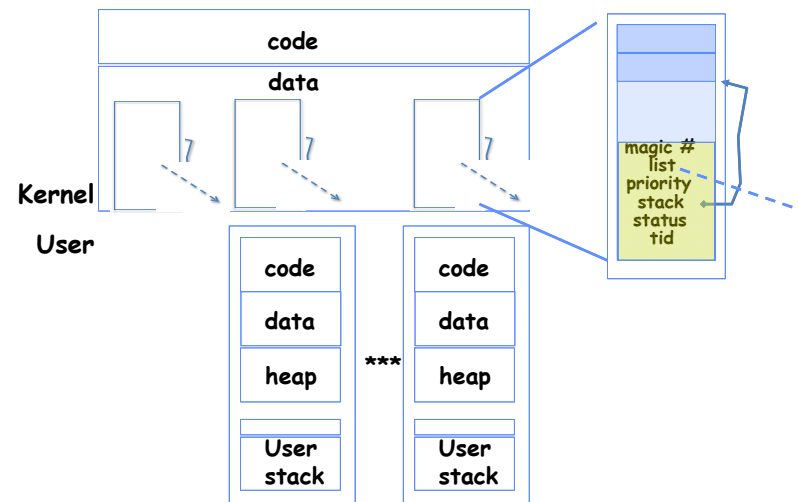
2/9/15

cs162 fa14 LZ Kubiatowicz CS162 ©UCB Spring 2015

35

Lec 6.35

## MT Kernel 1T Process ala Pintos/x86



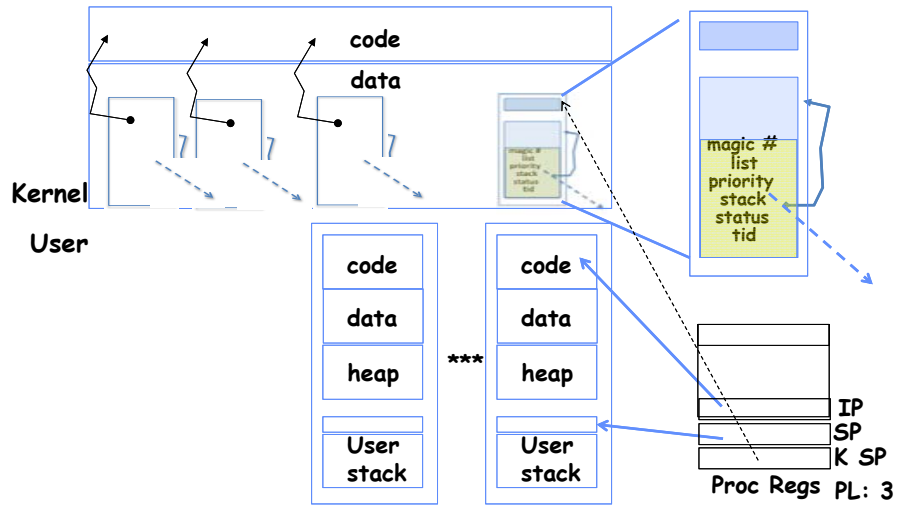
- Each user process/thread associated with a kernel thread, described by a 4kb Page object containing TCB and kernel stack for the kernel thread

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.36

### In User thread, w/ k-thread waiting



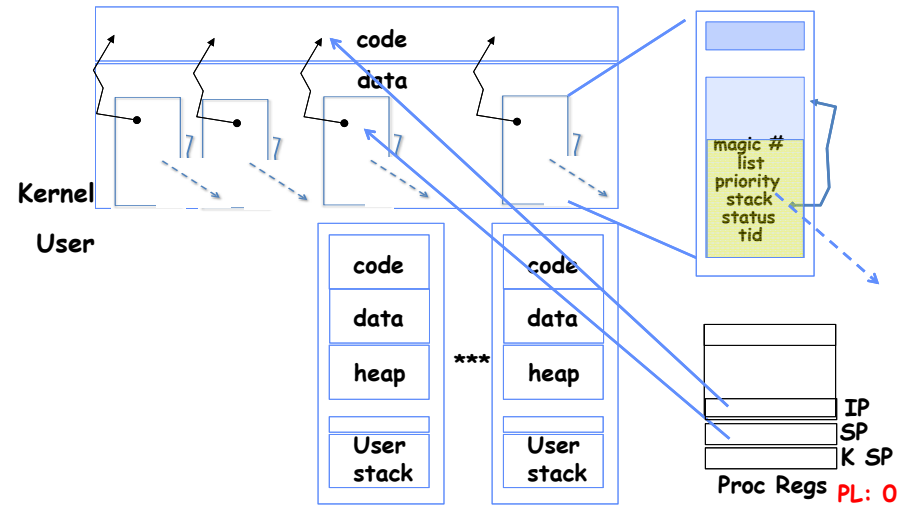
- x86 proc holds interrupt SP high system level
- During user thread exec, associate kernel thread is "standing by"

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.37

### In Kernel thread



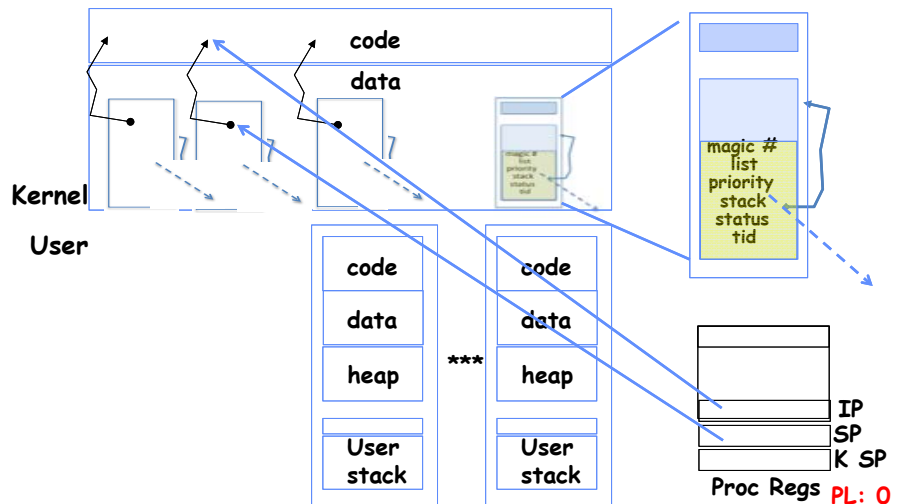
- Kernel threads execute with small stack in thread struct
- Scheduler selects among ready kernel and user threads

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.38

### Thread Switch (switch.S)



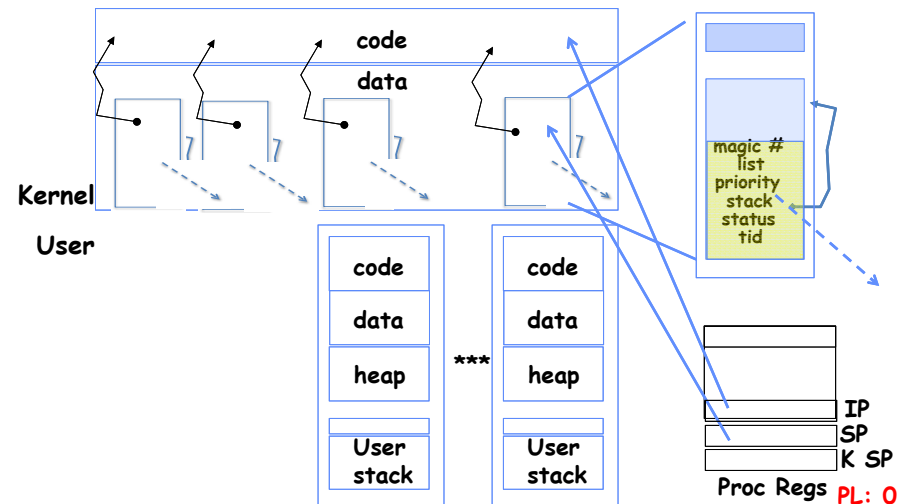
- `switch_threads`: save regs on current small stack, change SP, return from destination threads call to `switch_threads`

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.39

### Switch to Kernel Thread for Process

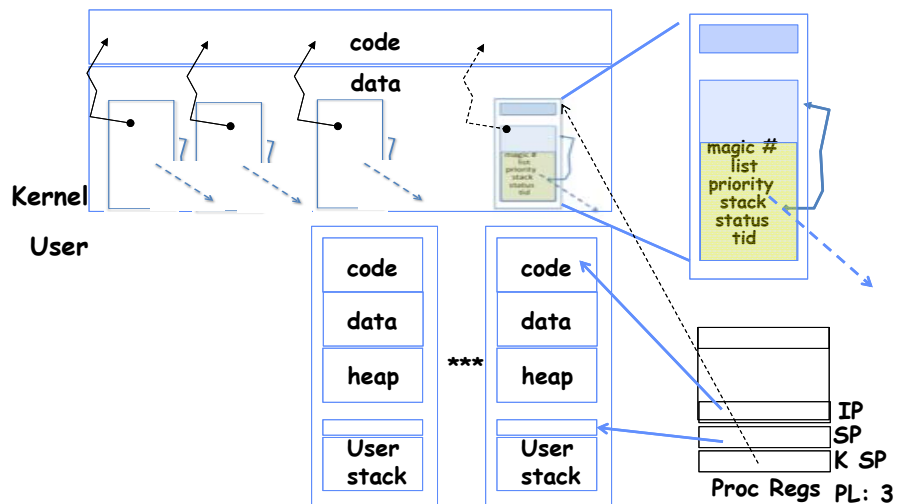


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.40

## Kernel->User



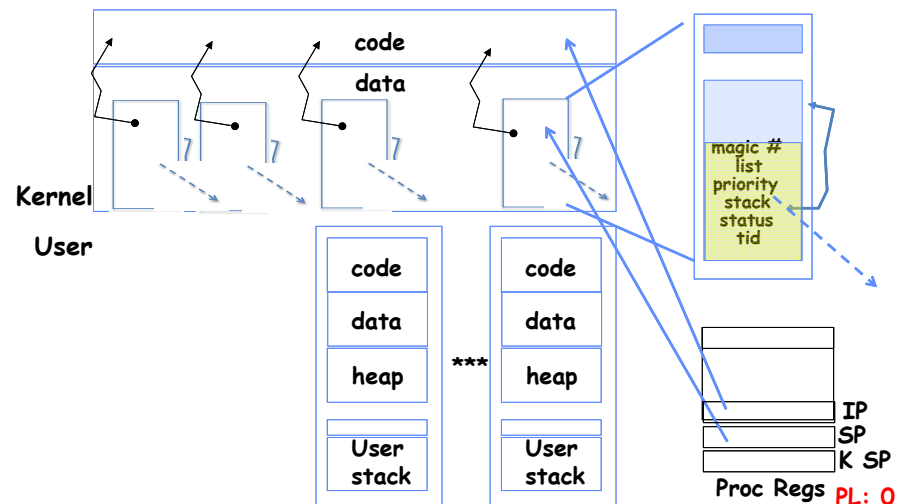
- `iret` restores user stack and PL

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.41

## User->Kernel



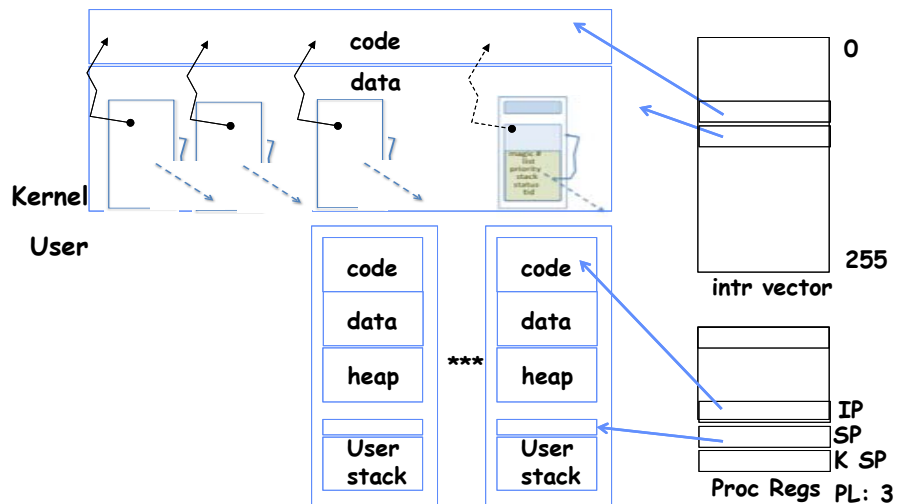
- Mechanism to resume k-thread goes through interrupt vector

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.42

## User->Kernel via interrupt vector



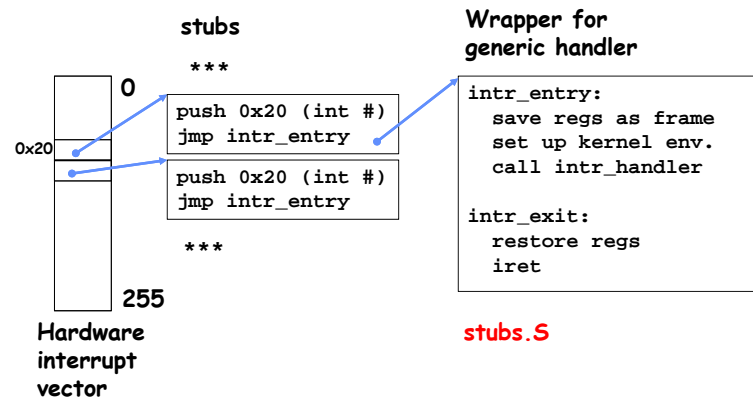
- Interrupt transfers control through the IV (IDT in x86)
- `iret` restores user stack and PL

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.43

## Pintos Interrupt Processing



**stubs.S**

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.44

## Recall: cs61C THE STACK FRAME

### Basic Structure of a Function

**Prologue**

```
entry_label:
addi $sp,$sp, -framesize
sw $ra, framesize-4($sp) # save $ra
save other regs if need be
```

**Body...** (call other functions...)

**Epilogue**

```
restore other regs if need be
lw $ra, framesize-4($sp) # restore $ra
addi $sp,$sp, framesize
jr $ra
```

### The Stack (review)

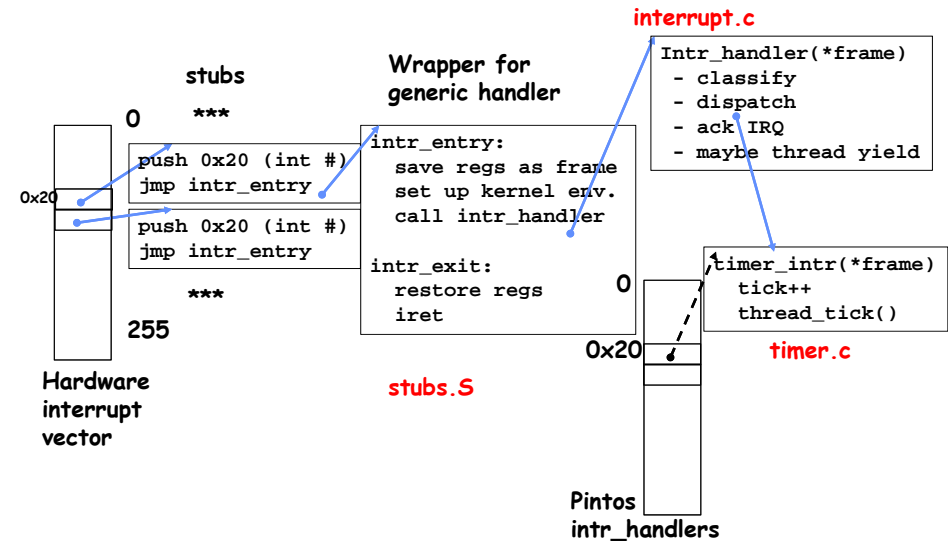
- Stack frame includes:
  - Return "instruction" address
  - Parameters
  - Space for other local variables
- Stack frames contiguous blocks of memory; stack pointer tells where bottom of stack frame is
- When procedure ends, stack frame is tossed off the stack, frees memory for future stack frames

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.45

## Pintos Interrupt Processing



2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.46

## Timer may trigger thread switch

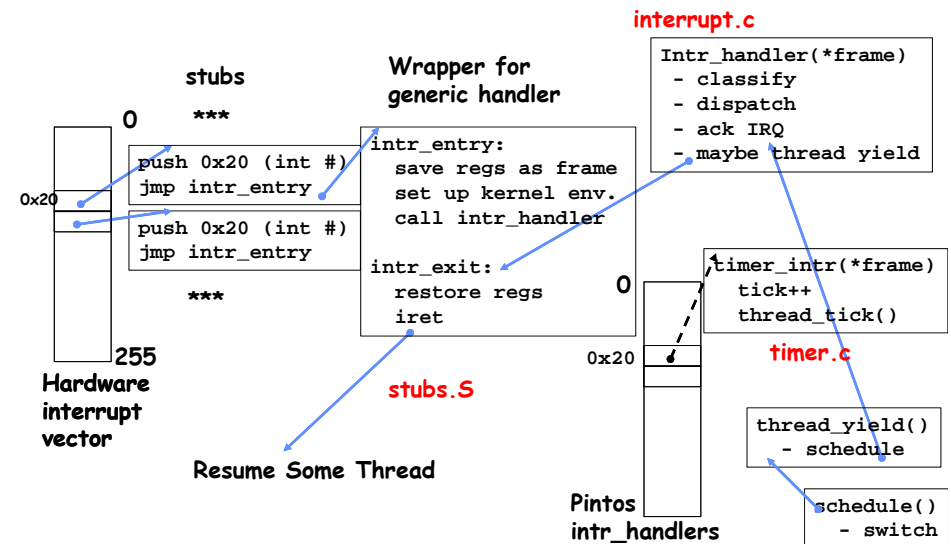
- thread\_tick**
  - Updates thread counters
  - If quanta exhausted, sets yield flag
- thread\_yield**
  - On path to rtn from interrupt
  - Sets current thread back to READY
  - Pushes it back on ready\_list
  - Calls schedule to select next thread to run upon iret
- Schedule**
  - Selects next thread to run
  - Calls switch\_threads to change regs to point to stack for thread to resume
  - Sets its status to RUNNING
  - If user thread, activates the process
  - Returns back to intr\_handler

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.47

## Pintos Return from Processing



2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.48

## Correctness for systems with concurrent threads

- If dispatcher can schedule threads in any way, programs must work under all circumstances
  - Can you test for this?
  - How can you know if your program works?
- **Independent Threads:**
  - No state shared with other threads
  - Deterministic  $\Rightarrow$  Input state determines results
  - Reproducible  $\Rightarrow$  Can recreate Starting Conditions, I/O
  - Scheduling order doesn't matter (if `switch()` works!!!)
- **Cooperating Threads:**
  - Shared State between multiple threads
  - Non-deterministic
  - Non-reproducible
- Non-deterministic and Non-reproducible means that bugs can be intermittent
  - Sometimes called "Heisenbugs"

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.49

## Interactions Complicate Debugging

- Is any program truly independent?
  - Every process shares the file system, OS resources, network, etc
  - Extreme example: buggy device driver causes thread A to crash "independent thread" B
- You probably don't realize how much you depend on reproducibility:
  - Example: Evil C compiler
    - » Modifies files behind your back by inserting errors into C program unless you insert debugging code
  - Example: Debugging statements can overrun stack
- Non-deterministic errors are really difficult to find
  - Example: Memory layout of kernel+user programs
    - » depends on scheduling, which depends on timer/other things
    - » Original UNIX had a bunch of non-deterministic errors
  - Example: Something which does interesting I/O
    - » User typing of letters used to help generate secure keys

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.50

## 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
    - » Many different file systems do read-ahead
  - Multiprocessors - chop up program into parallel pieces
- Advantage 3: Modularity
  - More important than you might think
  - Chop large problem up into simpler pieces
    - » To compile, for instance, gcc calls `cpp | cc1 | cc2 | as | ld`
    - » Makes system easier to extend

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.51

## High-level Example: Web Server



- Server must handle many requests
- Non-cooperating version:

```
serverLoop() {
    con = AcceptCon();
    ProcessFork(ServiceWebPage(), con);
}
```
- What are some disadvantages of this technique?

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.52

## Threaded Web Server

- Now, use a single process
- Multithreaded (cooperating) version:
 

```
serverLoop() {
    connection = AcceptCon();
    ThreadFork(ServiceWebPage(), connection);
}
```
- Looks almost the same, but has many advantages:
  - 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
- Question: would a user-level (say one-to-many) thread package make sense here?
  - When one request blocks on disk, all block...
- What about Denial of Service attacks or digg / Slash-dot effects?




2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.55

## Thread Pools

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



```

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

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


```

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.54

## ATM Bank Server



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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.55

## 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)

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.56



## 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

## 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(actId); /* 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	

## Review: Multiprocessing vs Multiprogramming

- What does it mean to run two threads “concurrently”?
  - Scheduler is free to run threads in any order and interleaving: FIFO, Random, ...
  - Dispatcher can choose to run each thread to completion or time-slice in big chunks or small chunks



### • Also recall: Hyperthreading

- Possible to interleave threads on a per-instruction basis
- Keep this in mind for our examples (like multiprocessing)

## 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!)
- Could even be 3 for serial processors:
  - » Thread A writes 0001, B writes 0010.
  - » Scheduling order ABABABBA yields 3!

## 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
  - Consequently - weird example that produces "3" on previous slide can't happen
- 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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.61

## 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/Xray 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


2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.62

## 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



2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.63

## 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?

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.64

## Hand Simulation Multiprocessor Example

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.65

## 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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.66

## Definitions

- **Synchronization**: using atomic operations to ensure cooperation between threads
  - For now, only loads and stores are atomic
  - We are going to show that 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. Only one thread at a time will get into this section of code.
  - Critical section is the result of mutual exclusion
  - Critical section and mutual exclusion are two ways of describing the same thing.

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.67

## 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 OJ



- Of Course - We don't know how to make a lock yet

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.68

## 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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.69

## 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?
  - Still too much milk **but only occasionally!**
  - 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 dispatcher does!

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.70

## 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) {  
        leave Note;  
        buy milk;  
    }  
}  
remove note;
```



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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.71

## Too Much Milk Solution #2

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

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


- Does this work?
- Possible for neither thread to buy milk
  - Context switches at exactly the wrong times can lead each to think that the other is going to buy
- Really insidious:
  - **Extremely unlikely** that this would happen, but will at worse possible time
  - Probably something like this in UNIX

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.72

## Too Much Milk Solution #2: problem!



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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.73

## Too Much Milk Solution #3

- Here is a possible two-note solution:

### Thread A

```
leave note A;
while (note B) { //X
    do nothing;
}
if (noMilk) {
    buy milk;
}
remove note A;
```

### Thread B

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

- 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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.74

## 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

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.75

## 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:

```
milklock.Acquire();
if (nomilk)
    buy milk;
milklock.Release();
```
- Once again, section of code between Acquire() and Release() called a "Critical Section"
- Of course, you can make this even simpler: suppose you are out of ice cream instead of milk
  - Skip the test since you always need more ice cream.

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.76



## Where are we going with synchronization?

Programs	Shared Programs
Higher-level API	Locks   Semaphores   Monitors   Send/Receive
Hardware	Load/Store   Disable Ints   Test&Set   Comp&Swap

- We are going to implement various higher-level synchronization primitives using atomic operations
  - Everything is pretty painful if only atomic primitives are load and store
  - Need to provide primitives useful at user-level

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.77

## Summary (1 of 2)

- Processes have two parts
  - Threads (Concurrency)
  - Address Spaces (Protection)
- Concurrency accomplished by multiplexing CPU Time:
  - Unloading current thread (PC, registers)
  - Loading new thread (PC, registers)
  - Such context switching may be voluntary (`yield()`, I/O operations) or involuntary (timer, other interrupts)
- Protection accomplished restricting access:
  - Memory mapping isolates processes from each other
  - Dual-mode for isolating I/O, other resources
- Various Textbooks talk about *processes*
  - When this concerns concurrency, really talking about thread portion of a process
  - When this concerns protection, talking about address space portion of a process

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.78

## Summary (2 of 2)

- 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
- Showed how to protect a critical section with only atomic load and store ⇒ pretty complex!

2/9/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 6.79