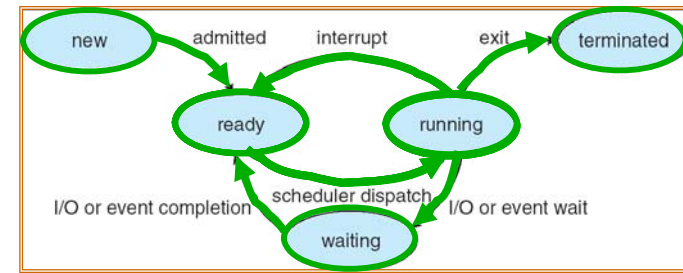


CS162 Operating Systems and Systems Programming Lecture 6

Concurrency (Continued), Synchronization (Start)

September 16th, 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.txt"));
}
  
```

- 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: Multithreaded stack switching

- Consider the following code blocks:

```


proc A() {
  B();
}

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



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

What happens when thread blocks on I/O?



- What happens when a thread requests a block of data from the file system?
 - User code invokes a system call
 - Read operation is initiated
 - Run new thread/switch
- Thread communication similar
 - Wait for Signal/Join
 - Networking

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.5

External Events

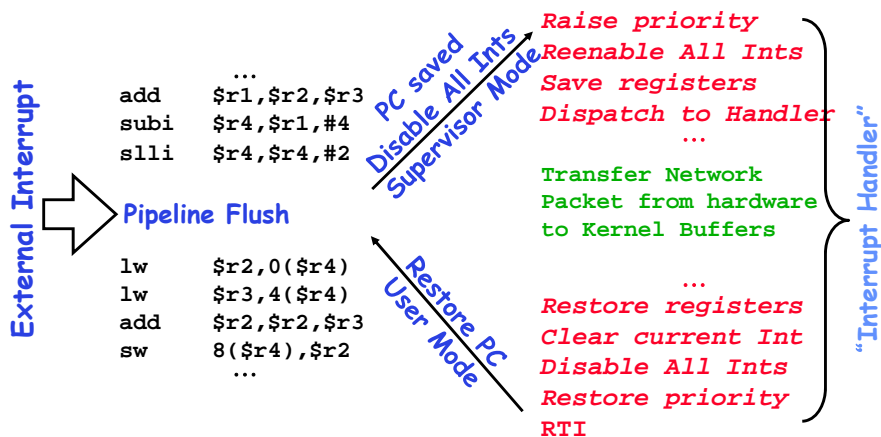
- What happens if thread never does any I/O, never waits, and never yields control?
 - Could the ComputePI program grab all resources and never release the processor?
 - » What if it didn't print to console?
 - Must find way that dispatcher can regain control!
- Answer: Utilize External Events
 - Interrupts: signals from hardware or software that stop the running code and jump to kernel
 - Timer: like an alarm clock that goes off every some many milliseconds
- If we make sure that external events occur frequently enough, can ensure dispatcher runs

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.6

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


9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.7

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().

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.8

ThreadFork(): Create a New Thread

- ThreadFork() is a user-level procedure that creates a new thread and places it on ready queue
 - We called this CreateThread() earlier
- Arguments to ThreadFork()
 - Pointer to application routine (fcnPtr)
 - Pointer to array of arguments (fcnArgPtr)
 - Size of stack to allocate
- Implementation
 - Sanity Check arguments
 - Enter Kernel-mode and Sanity Check arguments again
 - Allocate new Stack and TCB
 - Initialize TCB and place on ready list (Runnable).

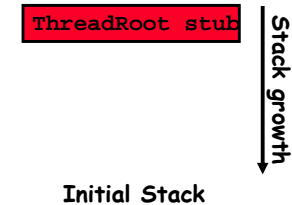
9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.9

How do we initialize TCB and Stack?

- Initialize Register fields of TCB
 - Stack pointer made to point at stack
 - PC return address \Rightarrow OS (asm) routine ThreadRoot()
 - Two arg registers (a0 and a1) initialized to fcnPtr and fcnArgPtr, respectively
- Initialize stack data?
 - No. Important part of stack frame is in registers (ra)
 - Think of stack frame as just before body of ThreadRoot() really gets started




9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.10

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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015


Lec 6.11

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



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.12

Administrivia

- **Group formation: should be completed by tonight!**
 - Will handle stragglers tonight
- **Section assignment**
 - Form due tonight by midnight!
 - We will try to do final section assignment tomorrow
- **Your section is your home for CS162**
 - The TA needs to get to know you to judge participation
 - All design reviews will be conducted by your TA
 - You can attend alternate section by same TA, but try to keep the amount of such cross-section movement to a minimum
- **Project #1: Released!**
 - Technically starts today
 - Autograder should be up by tomorrow.
- **HW1 due next Monday**
 - Must be submitted via the recommended "push" mechanism through git

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.13

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


9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.14

Multithreaded Processes

- **Process Control Block (PCBs) points to multiple Thread Control Blocks (TCBs):**



- **Switching threads within a block is a simple thread switch**
- **Switching threads across blocks requires changes to memory and I/O address tables.**

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.15

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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.16

Example multithreaded programs (con't)

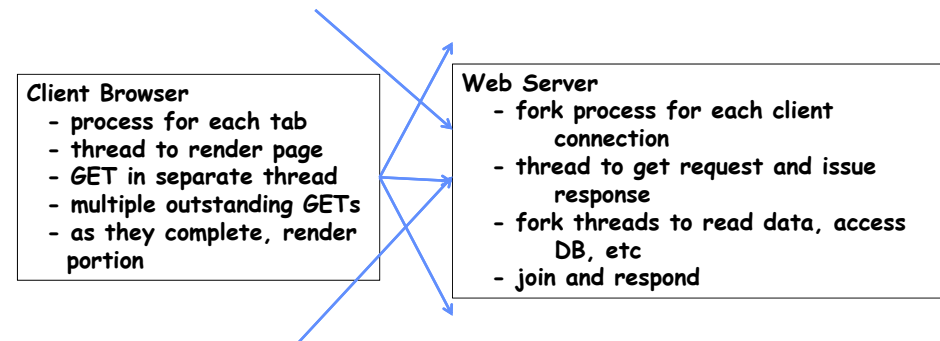
- **Network Servers**
 - Concurrent requests from network
 - Again, single program, multiple concurrent operations
 - File server, Web server, and airline reservation systems
- **Parallel Programming (More than one physical CPU)**
 - Split program into multiple threads for parallelism
 - This is called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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.17

A typical use case



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.18

Some Numbers

- Frequency of performing context switches: 10-100ms
- **Context switch time in Linux: 3-4 μ secs (Current Intel i7 & E5).**
 - Thread switching faster than process switching (100 ns).
 - But switching across cores about 2x more expensive than within-core switching.
- Context switch time increases sharply with the size of the working set*, and can increase 100x or more.

* The working set is the subset of memory used by the process in a time window.

- **Moral:** Context switching depends mostly on cache limits and the process or thread's hunger for memory.

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.19

Some Numbers

- Many process 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	jfc	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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.20

Kernel Use Cases

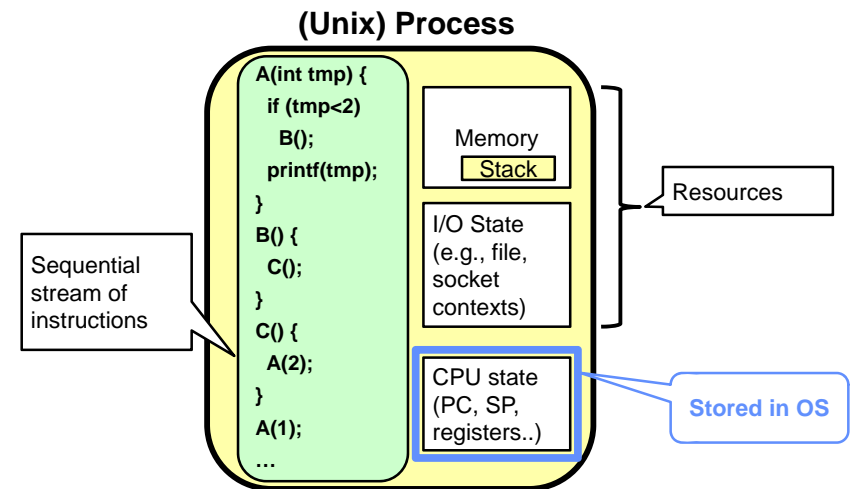
- Thread for each user process
- Thread for sequence of steps in processing I/O
- Threads for device drivers
- ...

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.21

Putting it together: Process

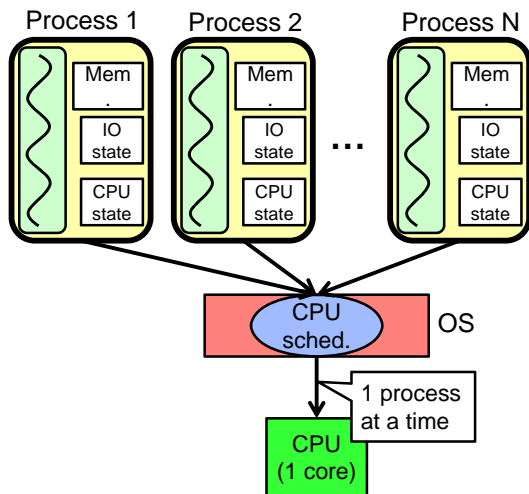


9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.22

Putting it together: Processes




- **Switch overhead: high**
 - Kernel entry: **low (ish)**
 - CPU state: **low**
 - Memory/IO state: **high**
- **Process creation: high**
- **Protection**
 - CPU: **yes**
 - Memory/IO: **yes**
- **Sharing overhead: high** (involves at least a context switch)

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.23

Putting it together: Threads



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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.24

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
- 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...



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015


Lec 6.25

Some Threading Models

Simple One-to-One Threading Model



Many-to-One



Many-to-Many

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.26

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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.27

Putting it together: Multi-Cores



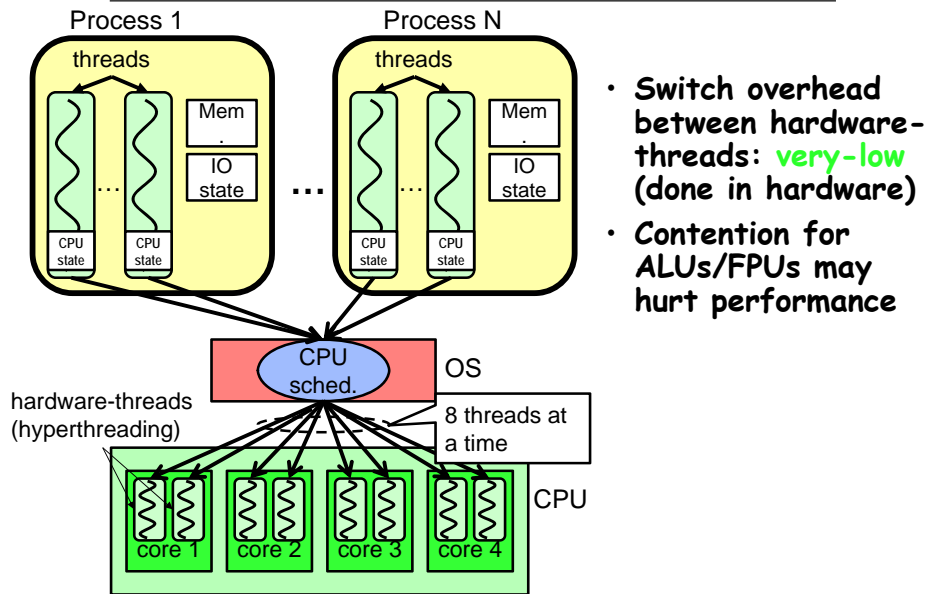
- Switch overhead: **low** (only CPU state)
- Thread creation: **low**
- Protection
 - CPU: **yes**
 - Memory/IO: **No**
- Sharing overhead: **low** (thread switch overhead low, may not need to switch at all!)

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.28

Putting it together: Hyper-Threading



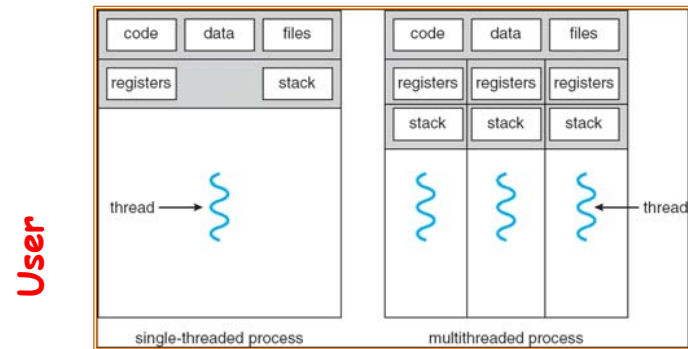
- Switch overhead between hardware-threads: **very-low** (done in hardware)
- Contention for ALUs/FPUs may hurt performance

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.29

Supporting 1T and MT Processes



User

System

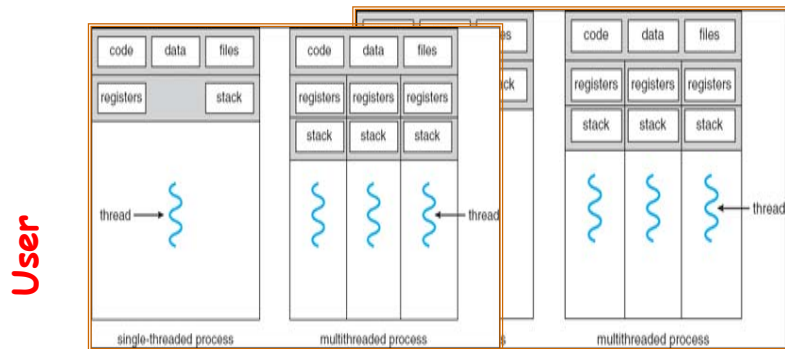


9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.30

Supporting 1T and MT Processes



User

System



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.31

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

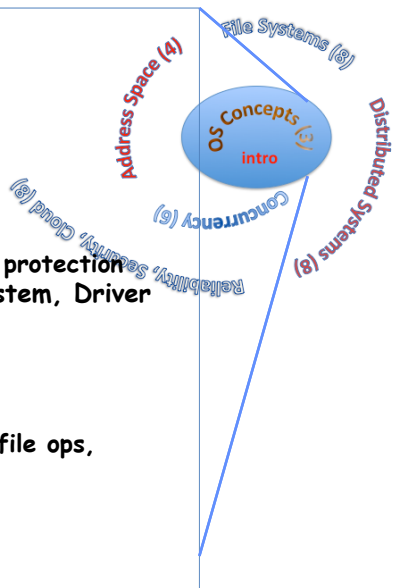
9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.32

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



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.33

Perspective on 'groking' 162

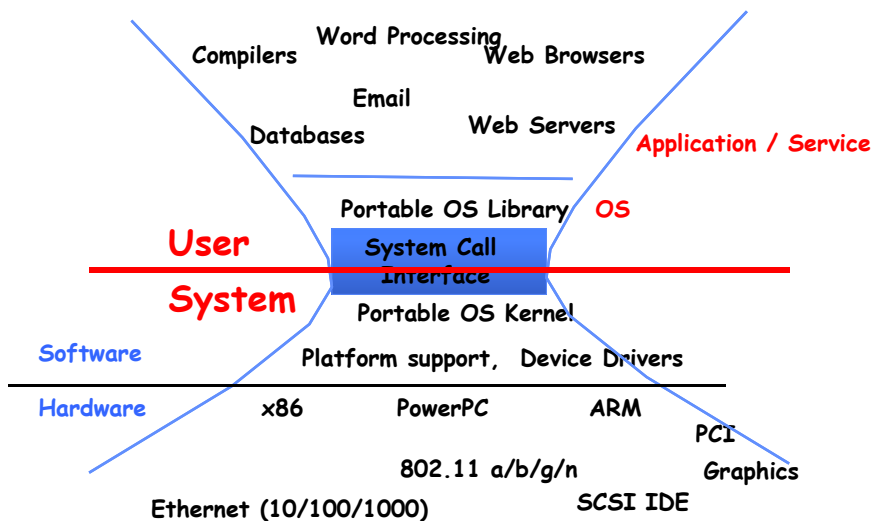
- Historically, OS was the most complex software
 - Concurrency, synchronization, processes, devices, communication, ...
 - Core systems concepts developed there
- Today, many "applications" are complex software systems too
 - These concepts appear there
 - But they are realized out of the capabilities 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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.34

Operating System as Design

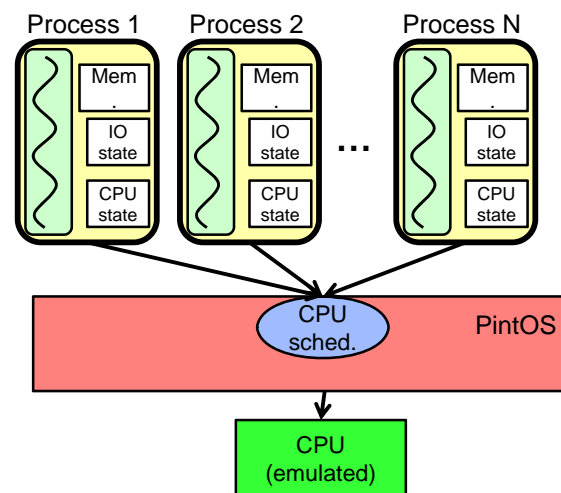


9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.35

Starting today: Pintos Projects



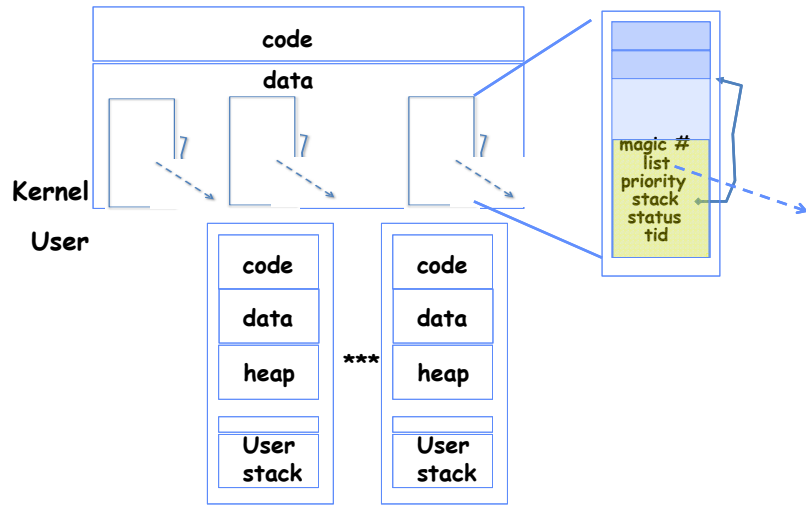
- Groups almost all formed
- Work as one!
- 10x homework
- P1: threads & scheduler
- P2: user process
- P3: file system

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.36

MT Kernel 1T Process ala Pintos/x86



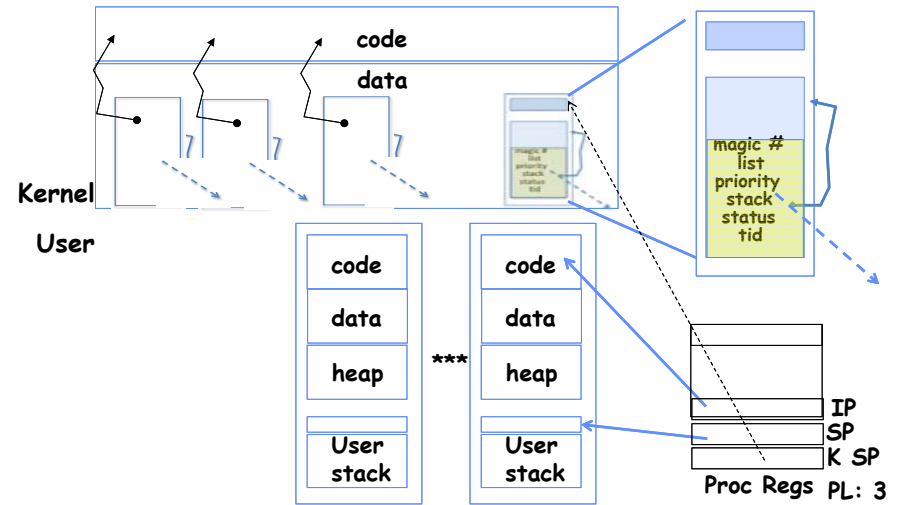
- Each user process/thread associated with a kernel thread, described by a 4kb Page object containing TCB and kernel stack for the kernel thread

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.37

In User thread, w/ k-thread waiting



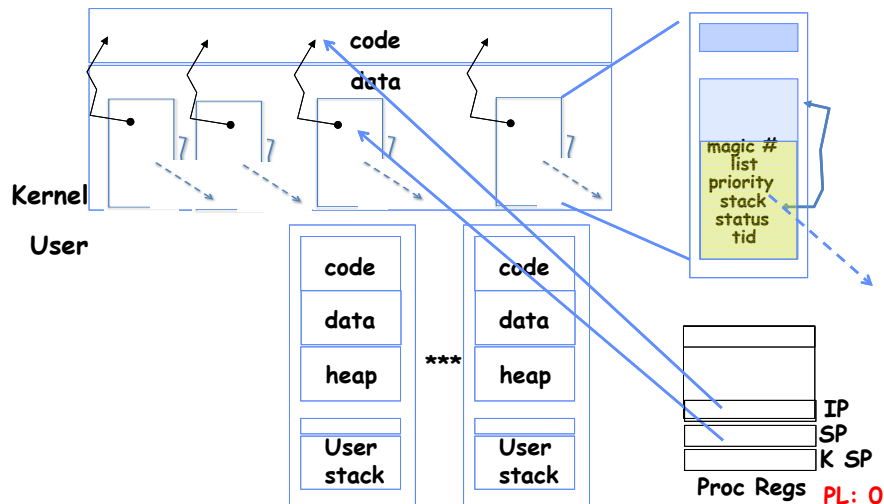
- x86 proc holds interrupt SP high system level
- During user thread exec, associate kernel thread is "standing by"

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.38

In Kernel thread



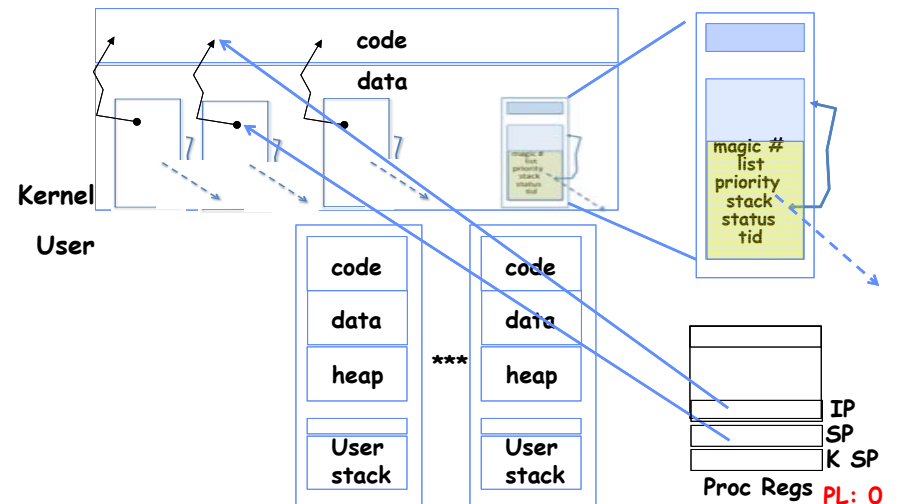
- Kernel threads execute with small stack in thread struct
- Scheduler selects among ready kernel and user threads

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.39

Thread Switch (switch.S)



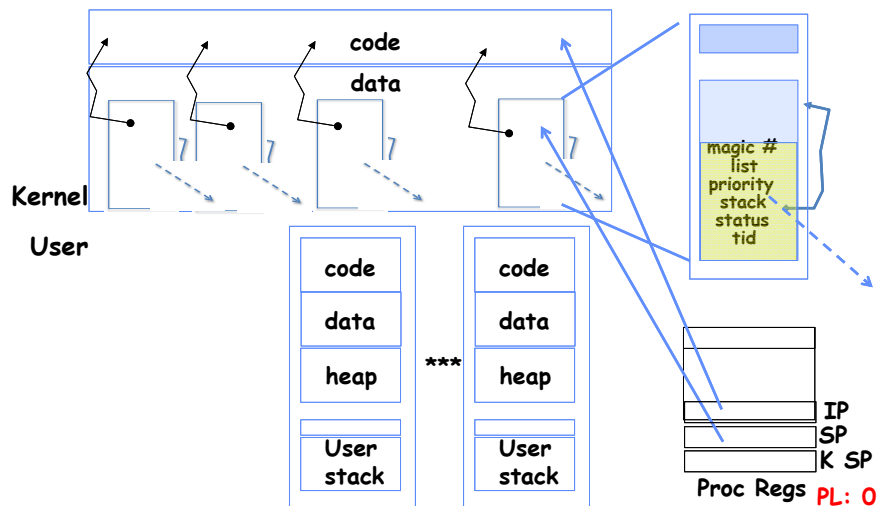
- switch_threads: save regs on current small stack, change SP, return from destination threads call to switch_threads

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.40

Switch to Kernel Thread for Process

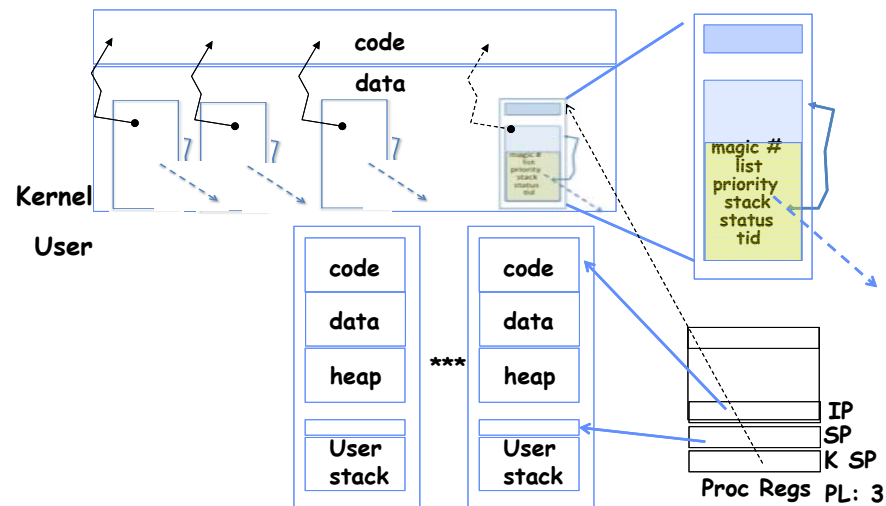


9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.41

Kernel->User



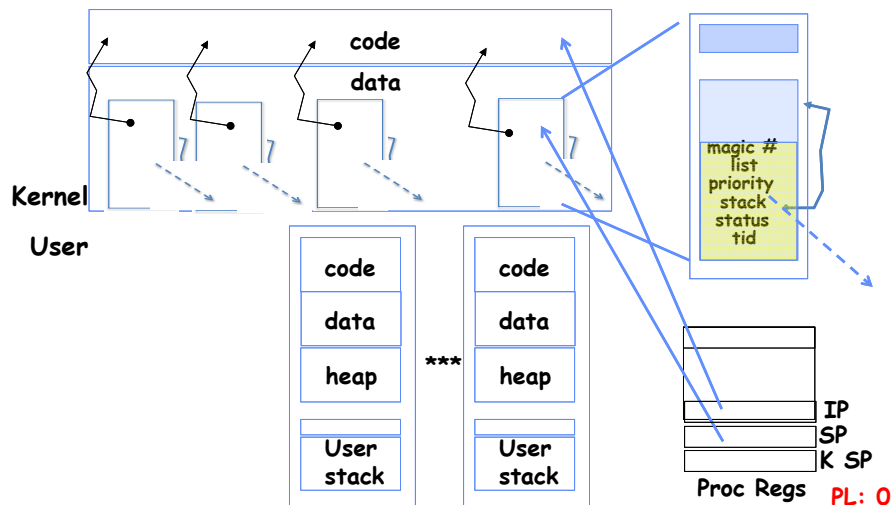
- iret restores user stack and PL

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.42

User->Kernel



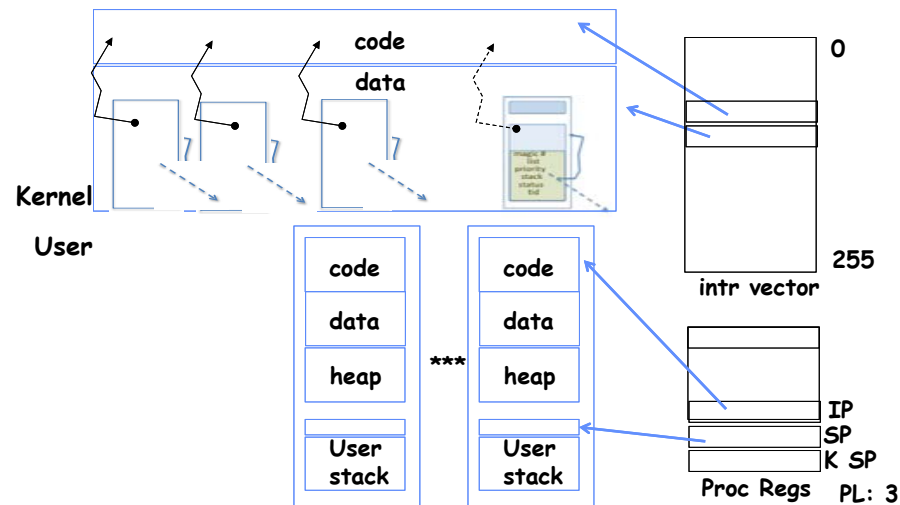
- Mechanism to resume k-thread goes through interrupt vector

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.43

User->Kernel via interrupt vector



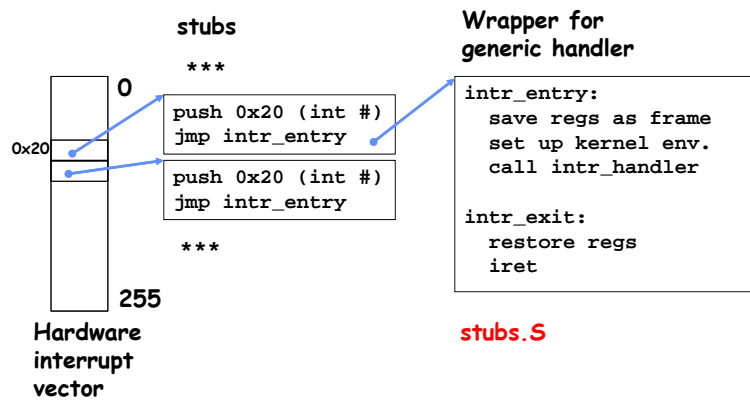
- Interrupt transfers control through the IV (IDT in x86)
- iret restores user stack and PL

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.44

Pintos Interrupt Processing



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.45

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)

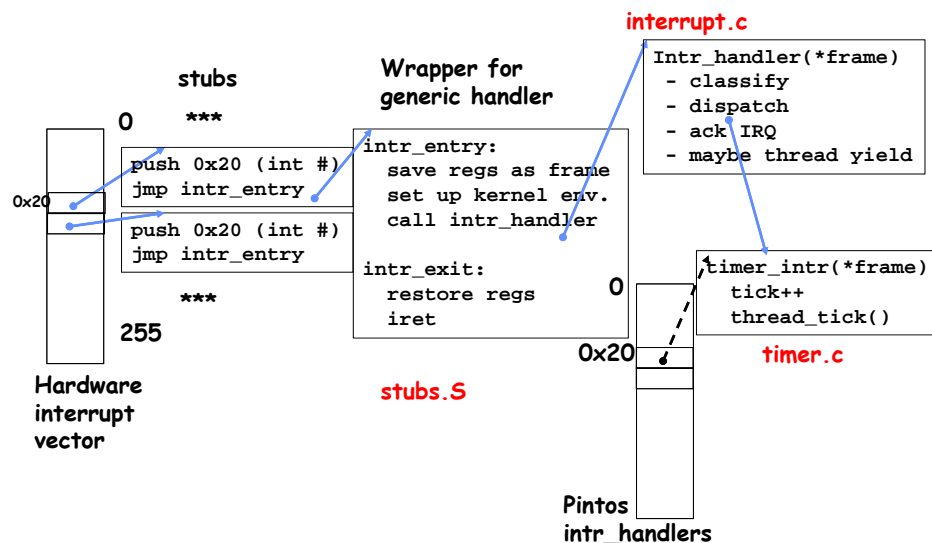
- Stack frame includes:
 - Return "instruction" address
 - Parameters
 - Space for other local variables
- Stack frames contiguous blocks of memory; stack pointer tells where bottom of stack frame is
- When procedure ends, stack frame is tossed off the stack, frees memory for future stack frames

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.46

Pintos Interrupt Processing



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.47

Timer may trigger thread switch

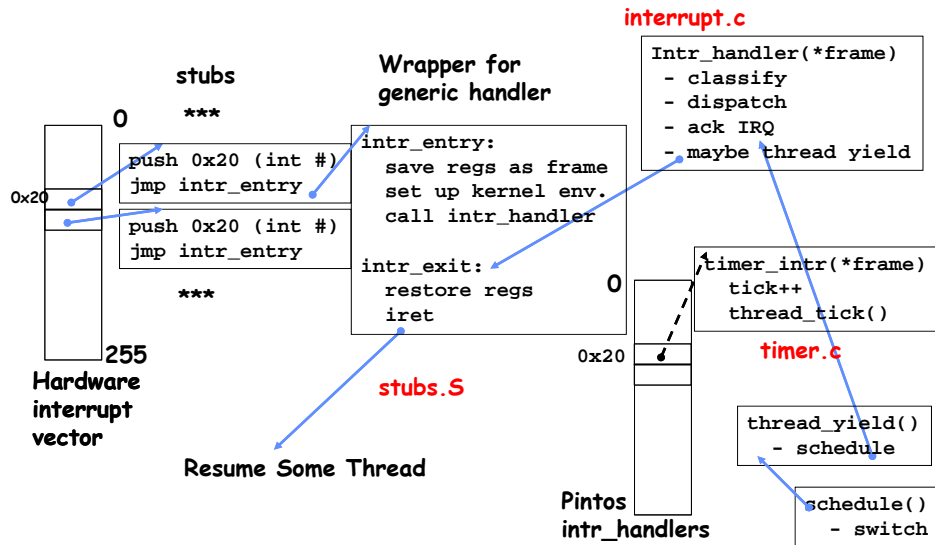
- **thread_tick**
 - Updates thread counters
 - If quanta exhausted, sets yield flag
- **thread_yield**
 - On path to rtn from interrupt
 - Sets current thread back to READY
 - Pushes it back on ready_list
 - Calls schedule to select next thread to run upon iret
- **Schedule**
 - Selects next thread 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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.48

Pintos Return from Processing




9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.49

Thread Abstraction



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


9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.50

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



9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.51

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"

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.52

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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.53

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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.54

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?

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.55

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?

9/16/15


Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.56



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);
    }
}
```

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.57

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

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.58

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!

9/16/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 6.59