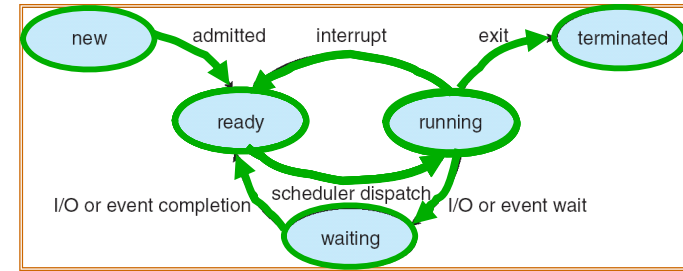


CS162 Operating Systems and Systems Programming Lecture 6

Concurrency (Continued), Thread and Processes

February 7th, 2019
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, "classlist.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

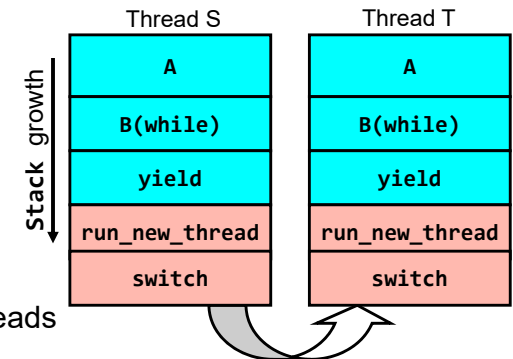


Recall: Multithreaded Stack Switching

- Consider the following code blocks:


```

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



- Suppose we have 2 threads running same code:
 - Threads S and T
 - Assume S and T have been running for a while

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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.5

External Events

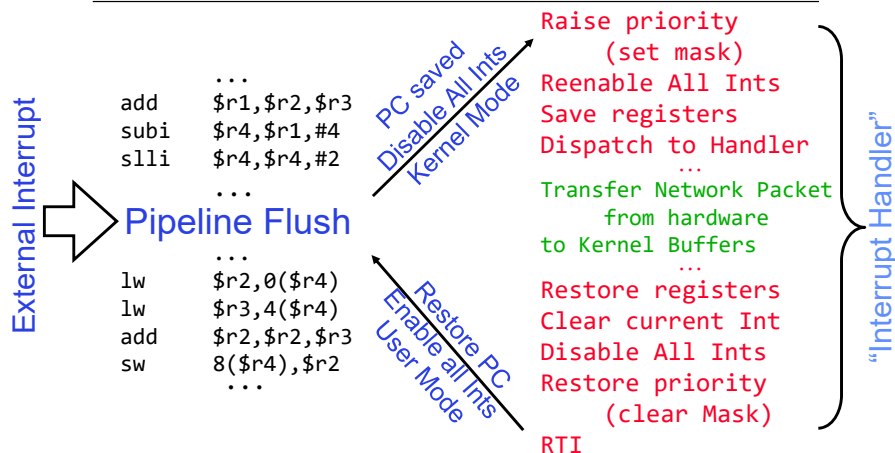
- What happens if thread never does any I/O, never waits, and never yields control?
 - Could the ComputePI program grab all resources and never release the processor?
 - » What if it didn't print to console?
 - Must find way that dispatcher can regain control!
- Answer: utilize external events
 - Interrupts: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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

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


2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

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();
}
```

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

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
- 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)

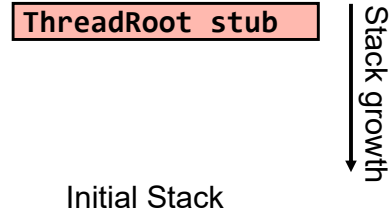
2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

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




2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019


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



2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.12

Administrivia


- Waitlist was closed last Friday/Early Drop passed Friday
- Recommendation: Read assigned readings *before* lecture
- Group sign up this week
 - Get finding groups ASAP – deadline Friday 2/8 at 11:59PM
 - 4 people in a group!
- Continue to attend whichever section is convenient
 - Next week, we start official section attendance!
- TA *preference* signup form due Tuesday 2/12 at 11:59PM
 - Everyone in a group must have the same TA!
 - » Preference given to same section
 - Participation: Get to know your TA!

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.13

Thread Abstraction




- Illusion: Infinite number of processors

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.14

Thread Abstraction



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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.15

Programmer vs. Processor View

Programmer's View	Possible Execution
	#1
	.
	.
	.
$x = x + 1;$	$x = x + 1;$
$y = y + x;$	$y = y + x;$
$z = x + 5y;$	$z = x + 5y;$
	.
	.
	.

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.16


Programmer vs. Processor View

Programmer's View	Possible Execution #1	Possible Execution #2
.	.	.
.	.	.
.	.	.
$x = x + 1;$	$x = x + 1;$	$x = x + 1$
$y = y + x;$	$y = y + x;$
$z = x + 5y;$	$z = x + 5y;$	thread is suspended
.	.	other thread(s) run
.	.	thread is resumed
.
		$y = y + x$
		$z = x + 5y$

Programmer vs. Processor View


Programmer's View	Possible Execution #1	Possible Execution #2	Possible Execution #3
.	.	.	.
.	.	.	.
.	.	.	.
$x = x + 1;$	$x = x + 1;$	$x = x + 1$	$x = x + 1$
$y = y + x;$	$y = y + x;$	$y = y + x$
$z = x + 5y;$	$z = x + 5y;$	thread is suspended
.	.	other thread(s) run	thread is suspended
.	.	thread is resumed	other thread(s) run
.	thread is resumed
		$y = y + x$
		$z = x + 5y$	$z = x + 5y$

Possible Executions



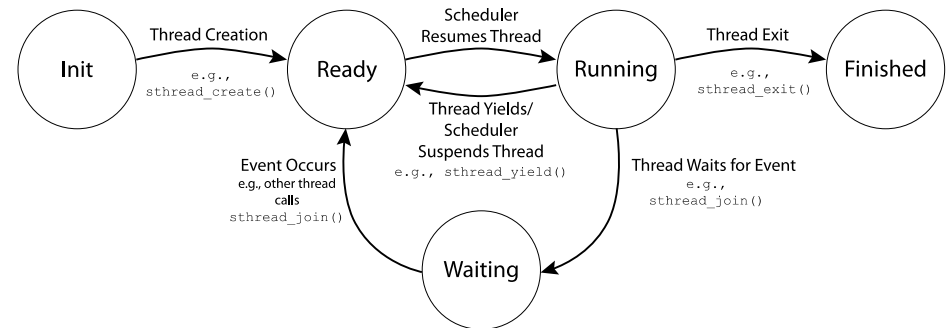
a) One execution

b) Another execution



c) Another execution

Thread Lifecycle



Per Thread Descriptor (Kernel Supported Threads)

- Each Thread has a **Thread Control Block (TCB)**
 - Execution State: CPU registers, program counter (PC), pointer to stack (SP)
 - Scheduling info: state, priority, CPU time
 - Various Pointers (for implementing scheduling queues)
 - Pointer to enclosing process (PCB) – user threads
 - ... (add stuff as you find a need)
- OS Keeps track of TCBs in “kernel memory”
 - In Array, or Linked List, or ...
 - I/O state (file descriptors, network connections, etc)


2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.21

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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.22

Examples multithreaded programs

- Embedded systems
 - Elevators, planes, medical systems, smart watches
 - 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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.23

Example multithreaded programs (con't)

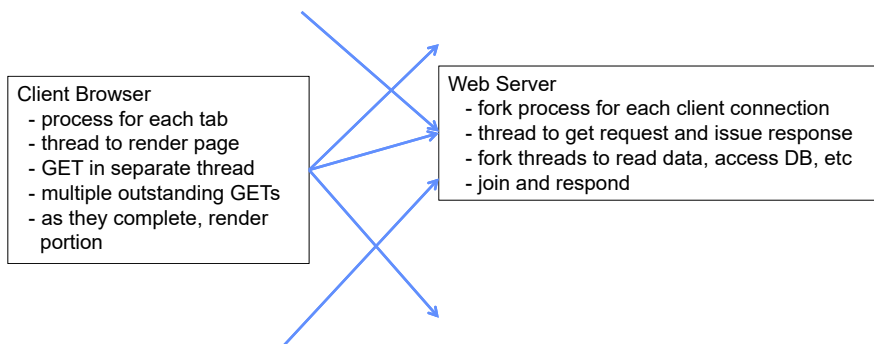
- Network servers
 - Concurrent requests from network
 - Again, single program, multiple concurrent operations
 - File server, Web server, and airline reservation systems
- Parallel programming (more than one physical CPU)
 - Split program into multiple threads for parallelism
 - This is called Multiprocessing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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.24

A Typical Use Case



2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.25

Some Numbers

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

*The working set is subset of memory used by process in a time window
- **Moral:** context switching depends mostly on cache limits and the process or thread's hunger for memory

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.26

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	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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.27

Kernel Use Cases

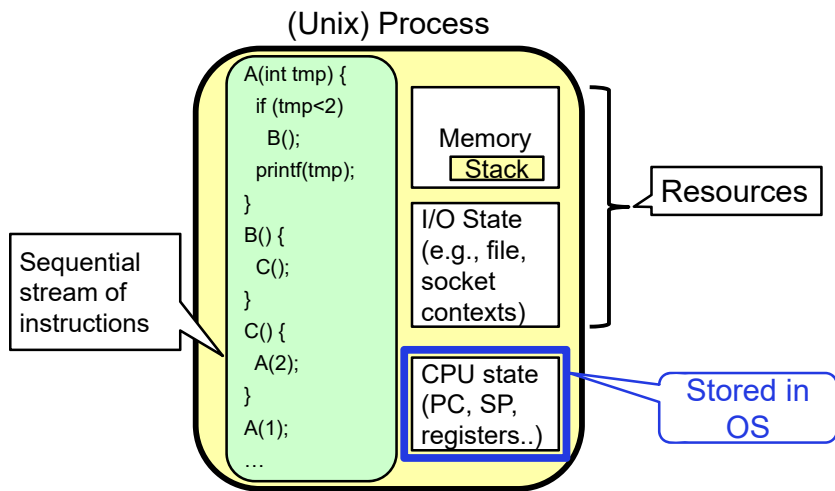
- Thread for each user process
- Thread for sequence of steps in processing I/O
- Threads for device drivers
- ...

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.28

Putting it Together: Process

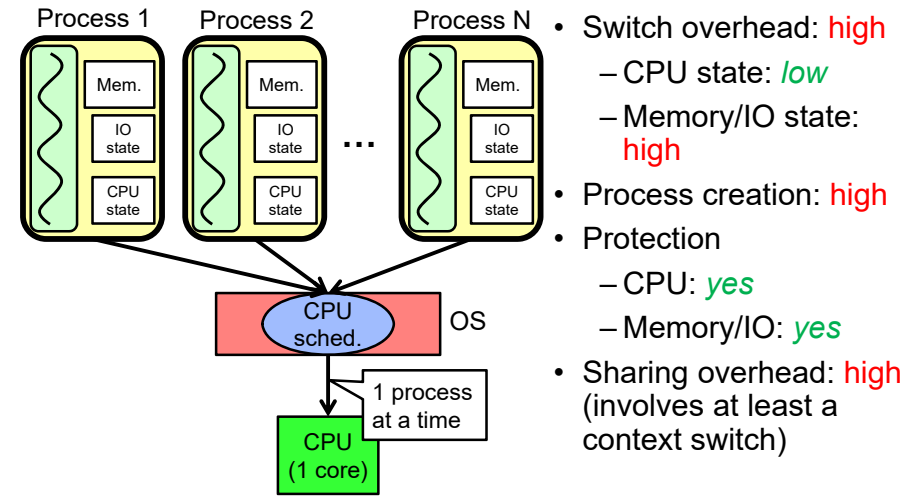


2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.29

Putting it Together: Processes

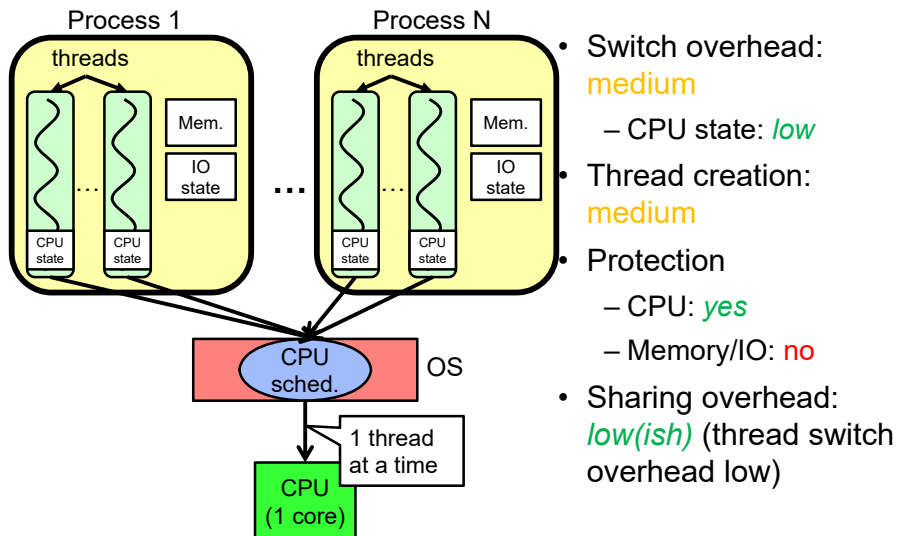


2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.30

Putting it Together: Threads



2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.31

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 level Threads

2/7/19


Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.32

User-Mode Threads

- Lighter weight option:

- 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...



2/7/19

Kubiatowicz CS162 ©UCB Fall 2019


Lec 6.33

Some Threading Models

Simple One-to-One Threading Model



Many-to-One



Many-to-Many

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.34

Threads in a Process


- Threads are useful at user-level: parallelism, hide I/O latency, interactivity
- Option A (early Java): user-level library, one multi-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): many single-threaded processes
 - User-level library does thread multiplexing
- Option C (Windows): scheduler activations
 - Kernel allocates processes 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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.35

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!)

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.36

Simultaneous MultiThreading/Hyperthreading

Hardware technique

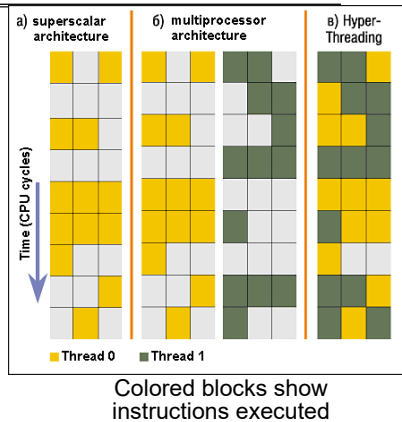
- Superscalar processors can execute multiple instructions that are independent
- Hyperthreading duplicates register state to make a second “thread,” allowing more instructions to run

Can schedule each thread as if were separate CPU

- But, sub-linear speedup!

Original called “Simultaneous Multithreading”

- <http://www.cs.washington.edu/research/smt/index.html>
- Intel, SPARC, Power (IBM)
- A virtual core on AWS’ EC2 is basically a hyperthread




2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.37

Putting it Together: Hyper-Threading



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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.38

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 10 Win NT to XP, Solaris, HP-UX, OS X

- Most operating systems have either
 - One or many address spaces
 - One or many threads per address space


2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.39

Multiprocessing vs Multiprogramming

- Remember Definitions:
 - Multiprocessing ≡ Multiple CPUs
 - Multiprogramming ≡ Multiple Jobs or Processes
 - Multithreading ≡ 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



2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.40

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"

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.41

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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.42

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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.43

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?

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.44

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?




2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.45

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);
    }
}


```

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.46

ATM Bank Server



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

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.47

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)

2/7/19

Kubiatowicz CS162 ©UCB Fall 2019

Lec 6.48

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	

Summary

- Processes have two parts
 - Threads (Concurrency)
 - Address Spaces (Protection)
- 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
- 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