


# CS162 Operating Systems and Systems Programming Lecture 2

## Introduction to the Process

January 26<sup>th</sup>, 2015  
Prof. John Kubiawicz  
<http://cs162.eecs.Berkeley.edu>

## Recall: What is an operating system?

- Special layer of software that provides application software access to hardware resources
  - Convenient abstraction of complex hardware devices
  - Protected access to shared resources
  - Security and authentication
  - Communication amongst logical entities



1/26/15

Kubiawicz CS162 ©UCB Spring 2015

Lec 2.2

## Review: What is an Operating System?



- Referee
  - Manage sharing of resources, Protection, Isolation
    - » Resource allocation, isolation, communication



- Illusionist
  - Provide clean, easy to use abstractions of physical resources
    - » Infinite memory, dedicated machine
    - » Higher level objects: files, users, messages
    - » Masking limitations, virtualization



- Glue
  - Common services
    - » Storage, Window system, Networking
    - » Sharing, Authorization
    - » Look and feel


1/26/15

Kubiawicz CS162 ©UCB Spring 2015

Lec 2.3

## Review: Increasing Software Complexity

Millions of lines of source code




From MIT's 6.033 course

1/26/15

Kubiawicz CS162 ©UCB Spring 2015

Lec 2.4

## Recall: Loading



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.5

## Very Brief History of OS

### • Several Distinct Phases:

- Hardware Expensive, Humans Cheap
  - » Eniac, ... Multics
- Hardware Cheaper, Humans Expensive
  - » PCs, Workstations, Rise of GUIs
- Hardware Really Cheap, Humans Really Expensive
  - » Ubiquitous devices, Widespread networking



"I think there is a world market for maybe five computers." -- Thomas Watson, chairman of IBM, 1943

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.6

## Very Brief History of OS

### • Several Distinct Phases:

- Hardware Expensive, Humans Cheap
  - » Eniac, ... Multics
- Hardware Cheaper, Humans Expensive
  - » PCs, Workstations, Rise of GUIs
- Hardware Really Cheap, Humans Really Expensive
  - » Ubiquitous devices, Widespread networking



Thomas Watson was often called "the worlds greatest salesman" by the time of his death in 1956

1/26/15


Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.7

## Very Brief History of OS

### • Several Distinct Phases:

- Hardware Expensive, Humans Cheap
  - » Eniac, ... Multics
- Hardware Cheaper, Humans Expensive
  - » PCs, Workstations, Rise of GUIs
- Hardware Really Cheap, Humans Really Expensive
  - » Ubiquitous devices, Widespread networking



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.8

## Very Brief History of OS

- Several Distinct Phases:
  - Hardware Expensive, Humans Cheap
    - » Eniac, ... Multics
  - Hardware Cheaper, Humans Expensive
    - » PCs, Workstations, Rise of GUIs
  - Hardware Really Cheap, Humans Really Expensive
    - » Ubiquitous devices, Widespread networking
- Rapid Change in Hardware Leads to changing OS
  - Batch ⇒ Multiprogramming ⇒ Timesharing ⇒ Graphical UI ⇒ Ubiquitous Devices
  - Gradual Migration of Features into Smaller Machines
- Situation today is much like the late 60s
  - Small OS: 100K lines/Large: 10M lines (5M browser!)
  - 100-1000 people-years

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.9

## OS Archaeology


- Because of the cost of developing an OS from scratch, most modern OSes have a long lineage:
- Multics → AT&T Unix → BSD Unix → Ultrix, SunOS, NetBSD,...
- Mach (micro-kernel) + BSD → NextStep → XNU → Apple OSX, iPhone iOS
- Linux → Android OS
- CP/M → QDOS → MS-DOS → Windows 3.1 → NT → 95 → 98 → 2000 → XP → Vista → 7 → 8 → phone → ...
- Linux → RedHat, Ubuntu, Fedora, Debian, Suse,...

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.10

## Migration of OS Concepts and Features



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.11

## Today: Four fundamental OS concepts

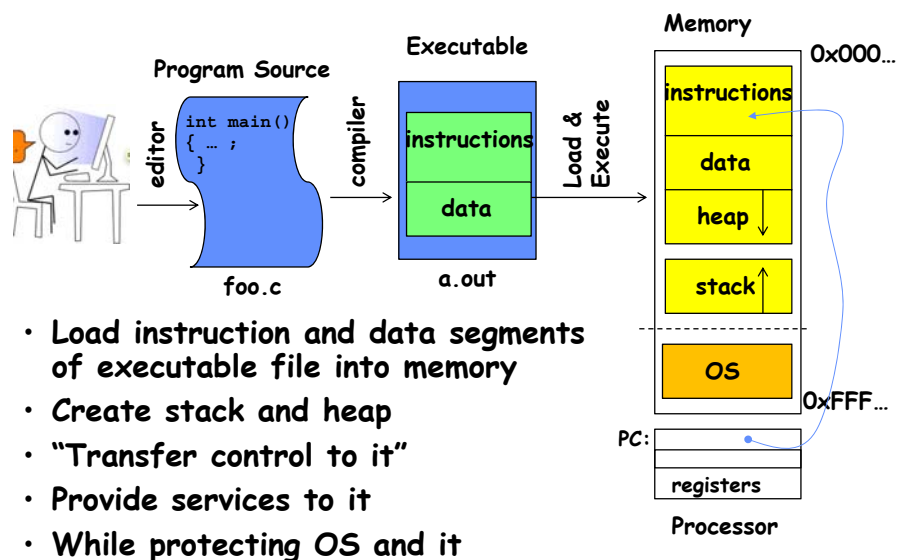
- Thread
  - Single unique execution context
  - Program Counter, Registers, Execution Flags, Stack
- Address Space w/ Translation
  - Programs execute in an *address space* that is distinct from the memory space of the physical machine
- Process
  - An instance of an executing program is a *process consisting of an address space and one or more threads of control*
- Dual Mode operation/Protection
  - Only the "system" has the ability to access certain resources
  - The OS and the hardware are protected from user programs and user programs are isolated from one another by *controlling the translation* from program virtual addresses to machine physical addresses

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.12

## OS Bottom Line: Run Programs



- Load instruction and data segments of executable file into memory
- Create stack and heap
- "Transfer control to it"
- Provide services to it
- While protecting OS and it


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.13

## Today we need one key 61B concept

### The instruction cycle

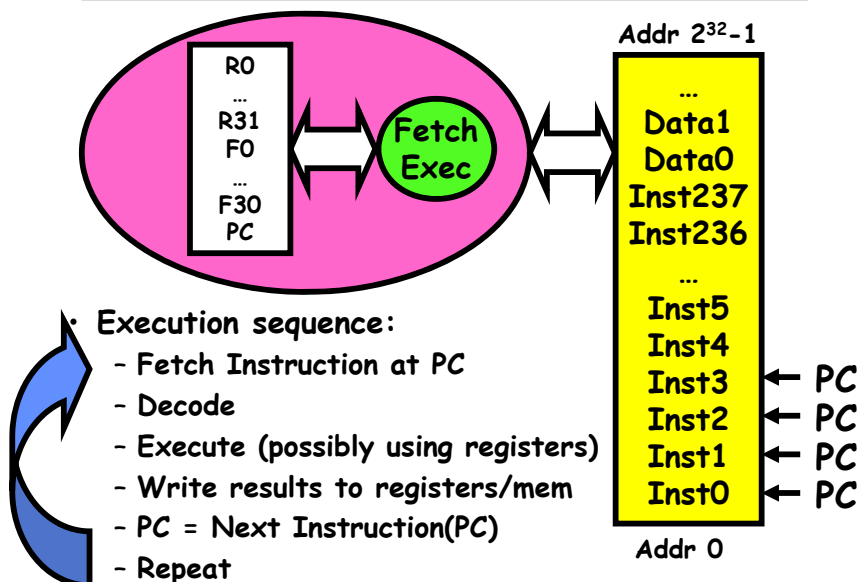


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.14

### Recall (61C): What happens during program execution?



#### Execution sequence:

- Fetch Instruction at PC
- Decode
- Execute (possibly using registers)
- Write results to registers/mem
- PC = Next Instruction(PC)
- Repeat

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.15

### First OS Concept: Thread of Control

- **Thread: Single unique execution context**
  - Program Counter, Registers, Execution Flags, Stack
- A thread is executing on a processor when it is resident in the processor registers.
- PC register holds the address of executing instruction in the thread.
- Certain registers hold the *context* of thread
  - Stack pointer holds the address of the top of stack
    - » Other conventions: Frame Pointer, Heap Pointer, Data
  - May be defined by the instruction set architecture or by compiler conventions
- Registers hold the root state of the thread.
  - The rest is "in memory"

1/26/15


Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.16

## Second OS Concept: Program's Address Space

• Address space  $\Rightarrow$  the set of accessible addresses + state associated with them:

- For a 32-bit processor there are  $2^{32} = 4$  billion addresses
- What happens when you read or write to an address?
  - Perhaps Nothing
  - Perhaps acts like regular memory
  - Perhaps ignores writes
  - Perhaps causes I/O operation
    - » (Memory-mapped I/O)
  - Perhaps causes exception (fault)

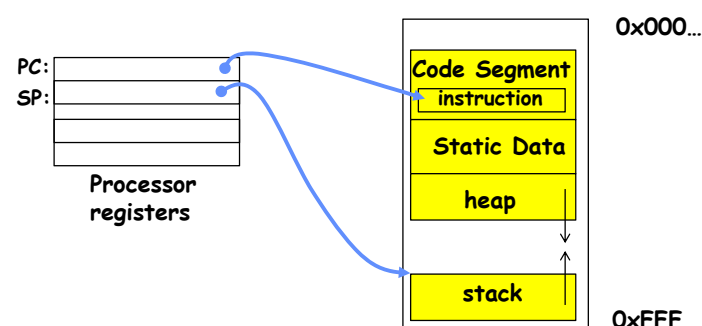


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.17

## Address Space: In a Picture



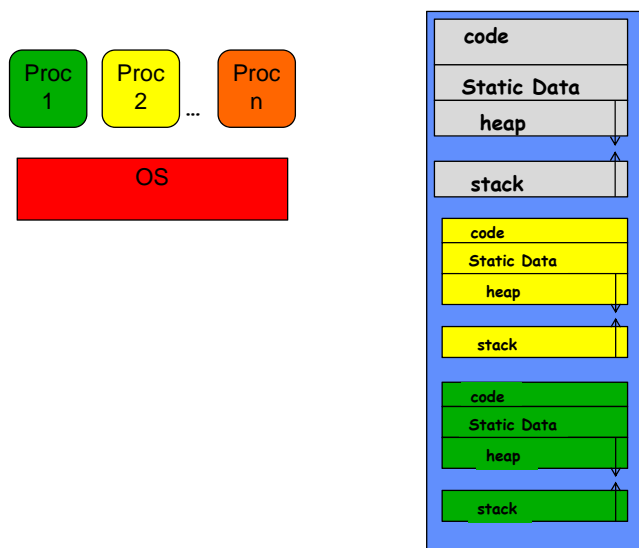
- What's in the code segment? Data?
- What's in the stack segment?
  - How is it allocated? How big is it?
- What's in the heap segment?
  - How is it allocated? How big?

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.18

## Multiprogramming - Multiple Threads of Control



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.19

## Administrivia: Getting started

- Start homework 0 immediately  $\Rightarrow$  **Due on Friday!**
  - Gets [cs162-xx@cory.eecs.berkeley.edu](mailto:cs162-xx@cory.eecs.berkeley.edu) (and other inst m/c)
  - Github account
  - Registration survey
  - Vagrant virtualbox - VM environment for the course
    - » Consistent, managed environment on your machine
  - icluster24.eecs.berkeley.edu is same
  - Get familiar with all the cs162 tools
  - Submit to autograder via git
- **Should be going to section already!**
- **Group sign up form out next week (after drop deadline)**
  - Get finding groups ASAP
  - 4 people in a group!

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.20

## Administrivia (Con't)

- Upcoming Workshops on Git: From Hackers@Berkeley
  - Introductory and advanced
  - Details on Piazza (link to facebook announcement)
- Kubiawicz Office Hours:
  - 2pm-3pm, Monday/Wednesday
  - May change as need arises (still have a bit of fluidity here as well)
- Online Textbooks:
  - Click on "Projects" link, under "Resources", there is a pointer to "Online Textbooks"
  - Can read these for free as long as on campus
  - First ones: Book on Git, two books on C
- Webcast:
  - We are webcasting this class
  - Will put link up off main page, but for now, go to:
    - » [webcast.Berkeley.edu](http://webcast.Berkeley.edu), click on "computer science" department
  - Webcast is **\*NOT\*** a replacement for coming to class!

1/26/15

Kubiawicz CS162 @UCB Spring 2015

Lec 2.21

## CS 162 Collaboration Policy

- ✓ Explaining a concept to someone in another group
- ✓ Discussing algorithms/testing strategies with other groups
- ✓ Helping debug someone else's code (in another group)
- ✓ Searching online for generic algorithms (e.g., hash table)
- ✗ Sharing code or test cases with another group
- ✗ Copying OR reading another group's code or test cases
- ✗ Copying OR reading online code or test cases from prior years


We compare all project submissions against prior year submissions and online solutions and will take actions (described on the course overview page) against offenders

1/26/15

Kubiawicz CS162 @UCB Spring 2015

Lec 2.22

## How can we give the illusion of multiple processors?



- Assume a single processor. How do we provide the illusion of multiple processors?
  - Multiplex in time!
- Each virtual "CPU" needs a structure to hold:
  - Program Counter (PC), Stack Pointer (SP)
  - Registers (Integer, Floating point, others...?)
- How switch from one virtual CPU to the next?
  - Save PC, SP, and registers in current state block
  - Load PC, SP, and registers from new state block
- What triggers switch?
  - Timer, voluntary yield, I/O, other things

1/26/15

Kubiawicz CS162 @UCB Spring 2015

Lec 2.23

## The Basic Problem of Concurrency

- The basic problem of concurrency involves resources:
  - Hardware: single CPU, single DRAM, single I/O devices
  - Multiprogramming API: processes think they have exclusive access to shared resources
- OS has to coordinate all activity
  - Multiple processes, I/O interrupts, ...
  - How can it keep all these things straight?
- Basic Idea: Use Virtual Machine abstraction
  - Simple machine abstraction for processes
  - Multiplex these abstract machines
- Dijkstra did this for the "THE system"
  - Few thousand lines vs 1 million lines in OS 360 (1K bugs)

1/26/15

Kubiawicz CS162 @UCB Spring 2015

Lec 2.24

## Properties of this simple multiprogramming technique

- All virtual CPUs share same non-CPU resources
  - I/O devices the same
  - Memory the same
- Consequence of sharing:
  - Each thread can access the data of every other thread (good for sharing, bad for protection)
  - Threads can share instructions (good for sharing, bad for protection)
  - Can threads overwrite OS functions?
- This (unprotected) model is common in:
  - Embedded applications
  - Windows 3.1/Early Macintosh (switch only with yield)
  - Windows 95—ME (switch with both yield and timer)

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.25

## Third OS Concept: Process

- **Process: execution environment with Restricted Rights**
  - **Address Space with One or More Threads**
  - Owns memory (address space)
  - Owns file descriptors, file system context, ...
  - Encapsulate one or more threads sharing process resources
- Why processes?
  - **Protected from each other!**
  - **OS Protected from them**
  - Navigate fundamental tradeoff between protection and efficiency
  - Processes provides memory protection
  - Threads more efficient than processes (later)
- Application instance consists of one or more processes

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.26

## Protection

- Operating System must protect itself from user programs
  - Reliability: compromising the operating system generally causes it to crash
  - Security: limit the scope of what processes can do
  - Privacy: limit each process to the data it is permitted to access
  - Fairness: each should be limited to its appropriate share
- It must protect User programs from one another
- Primary Mechanism: limit the translation from program address space to physical memory space
  - Can only touch what is mapped in
- Additional Mechanisms:
  - Privileged instructions, in/out instructions, special registers
  - syscall processing, subsystem implementation
    - » (e.g., file access rights, etc)

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.27

## Fourth OS Concept: Dual Mode Operation

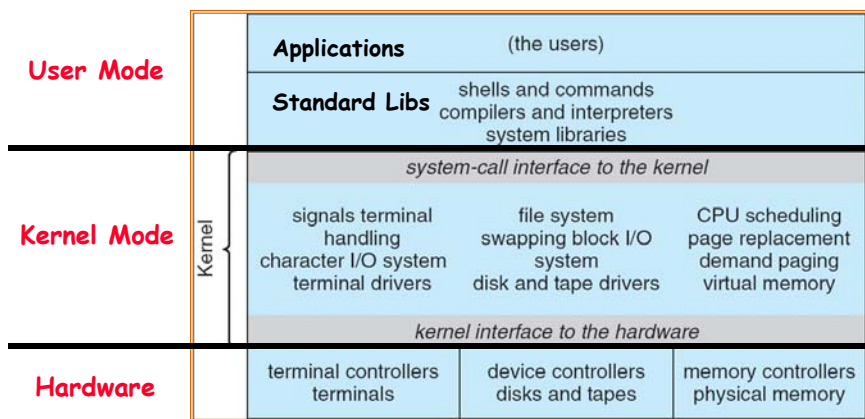
- **Hardware** provides at least two modes:
  - "Kernel" mode (or "supervisor" or "protected")
  - "User" mode: Normal programs executed
- What is needed in the hardware to support "dual mode" operation?
  - a bit of state (user/system mode bit)
  - Certain operations / actions only permitted in system/kernel mode
    - » In user mode they fail or trap
  - User- > Kernel transition *sets* system mode AND saves the user PC
    - » Operating system code carefully puts aside user state then performs the necessary operations
  - Kernel- > User transition clears system mode AND restores appropriate user PC
    - » return-from-interrupt

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.28

## For example: UNIX System Structure




1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.29

## User/Kernel(Privileged) Mode

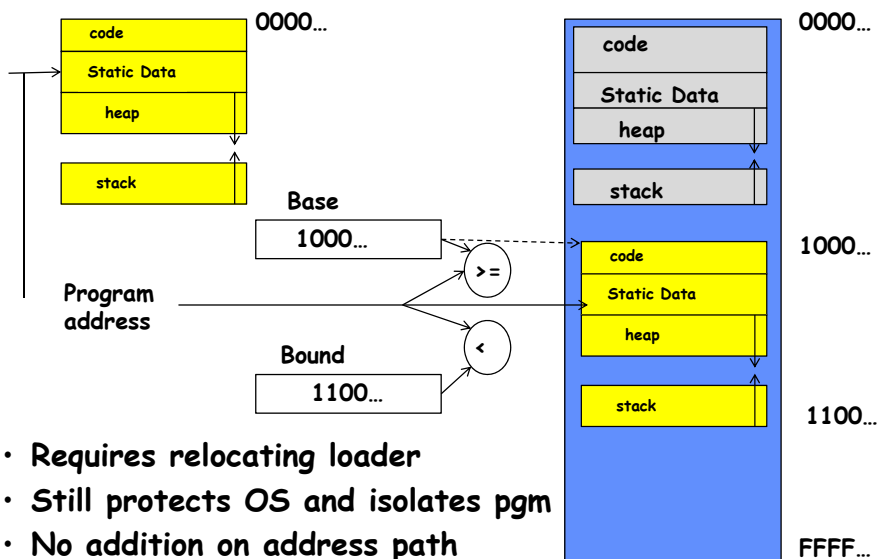


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.30

## Simple Protection: Base and Bound (B&B)



- Requires relocating loader
- Still protects OS and isolates pgm
- No addition on address path


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.31

## Another idea: Address Space Translation

- Program operates in an address space that is distinct from the physical memory space of the machine




1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.32



## A simple address translation with Base and Bound




- Can the program touch OS?
- Can it touch other programs?

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.33

## Tying it together: Simple B&B: OS loads process

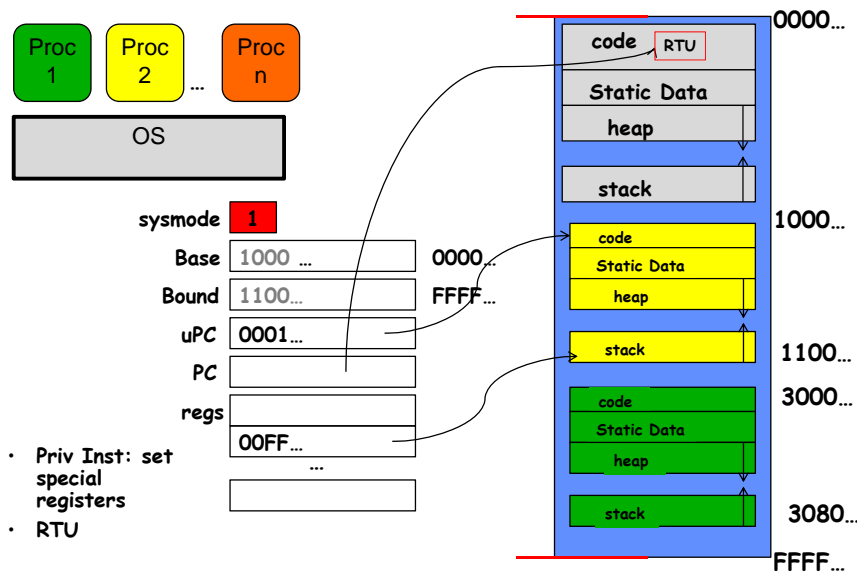


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.34

## Simple B&B: OS gets ready to switch



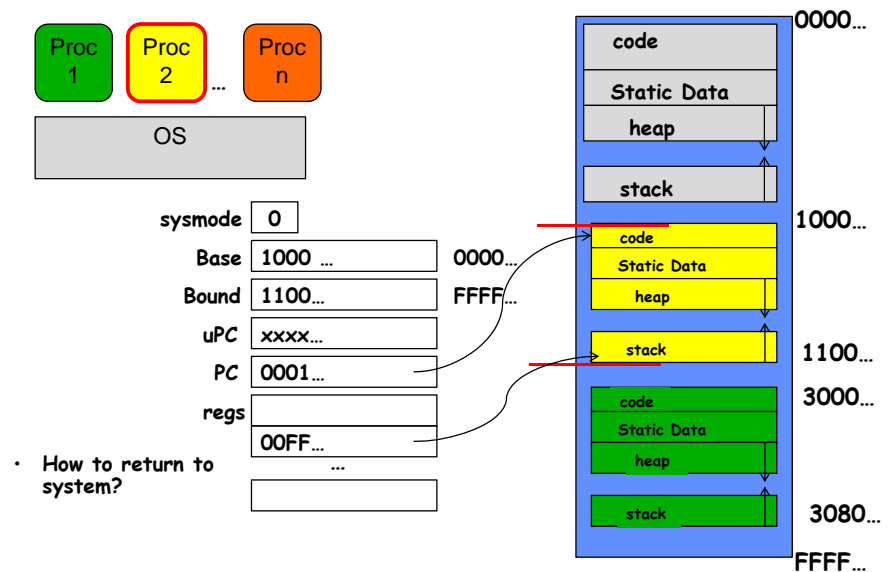
- Priv Inst: set special registers
- RTU

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.35

## Simple B&B: "Return" to User



- How to return to system?

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

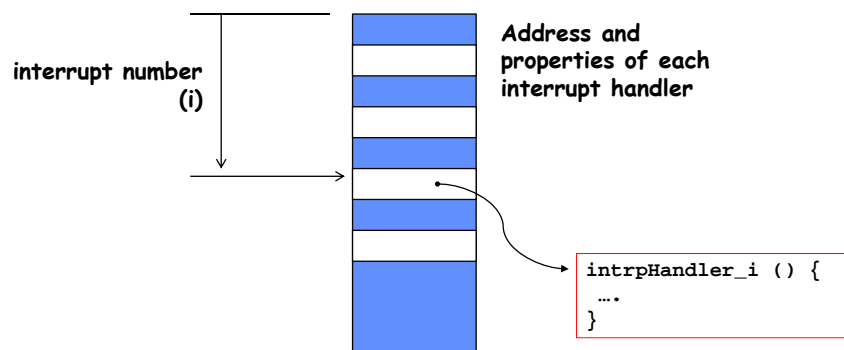
Lec 2.36

### 3 types of Mode Transfer

- **Syscall**
  - Process requests a system service, e.g., exit
  - Like a function call, but "outside" the process
  - Does not have the address of the system function to call
  - Like a Remote Procedure Call (RPC) - for later
  - Marshall the syscall id and args in registers and exec syscall
- **Interrupt**
  - External asynchronous event triggers context switch
  - eg. Timer, I/O device
  - Independent of user process
- **Trap or Exception**
  - Internal synchronous event in process triggers context switch
  - e.g., Protection violation (segmentation fault), Divide by zero, ...
- **All 3 are an UNPROGRAMMED CONTROL TRANSFER**
  - Where does it go?

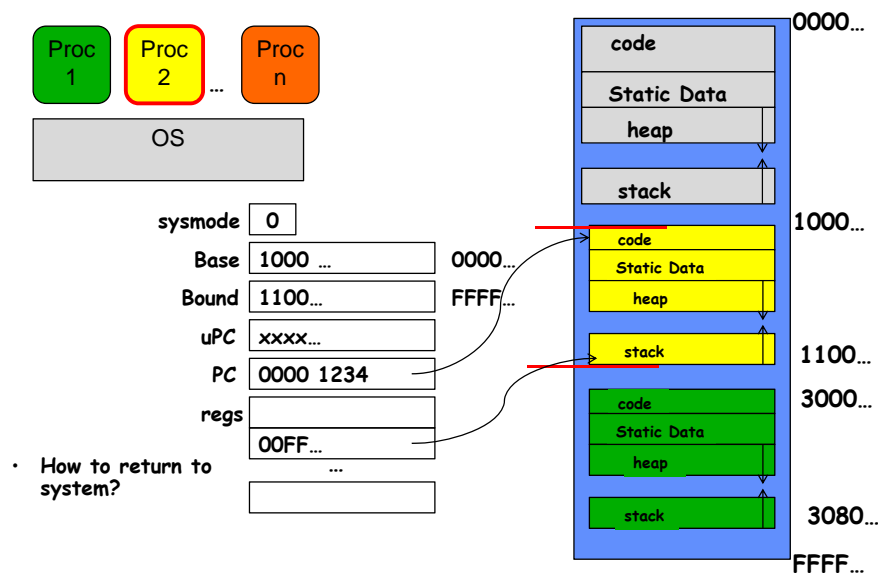
How do we get the system target address of the "unprogrammed control transfer?"

### Interrupt Vector

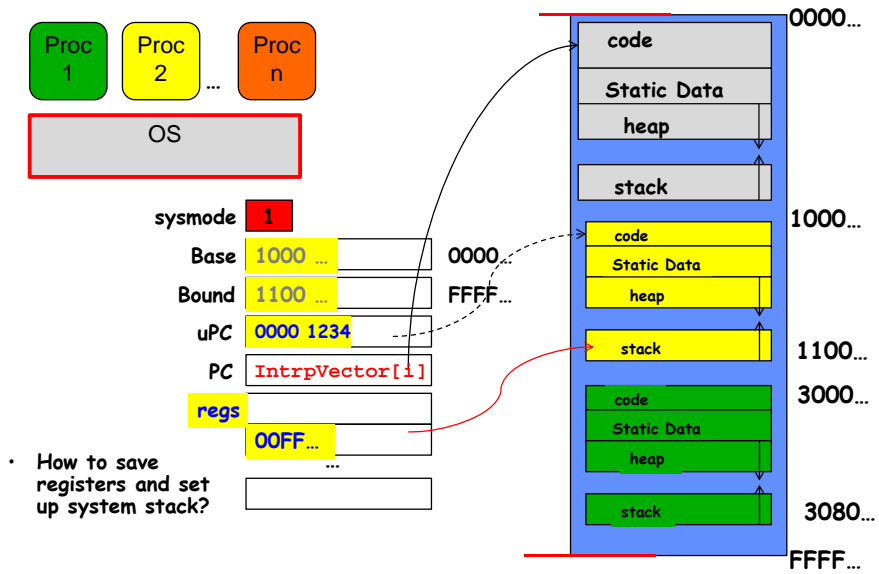


- Where else do you see this dispatch pattern?

### Simple B&B: User => Kernel



## Simple B&B: Interrupt

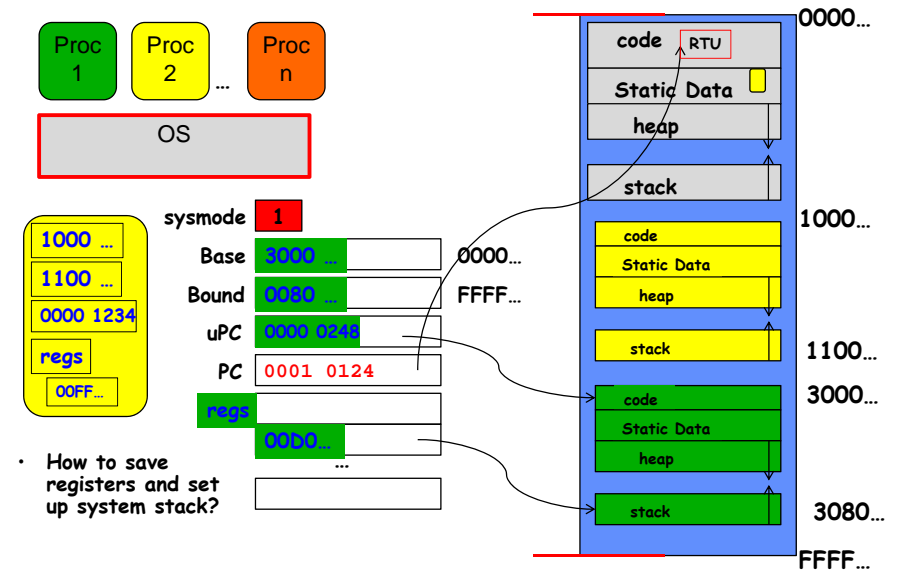


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.41

## Simple B&B: Switch User Process

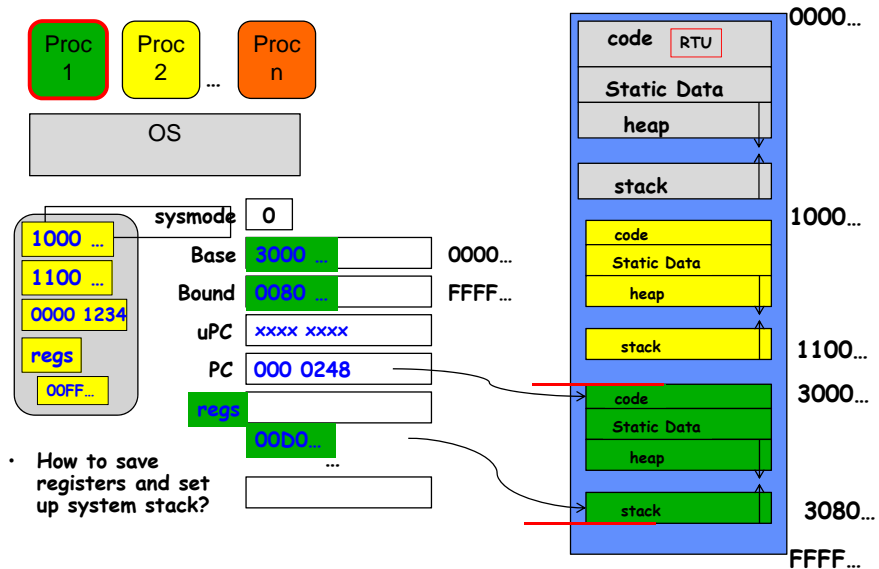


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.42

## Simple B&B: "resume"



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.43

What's wrong with this simplistic address translation mechanism?

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

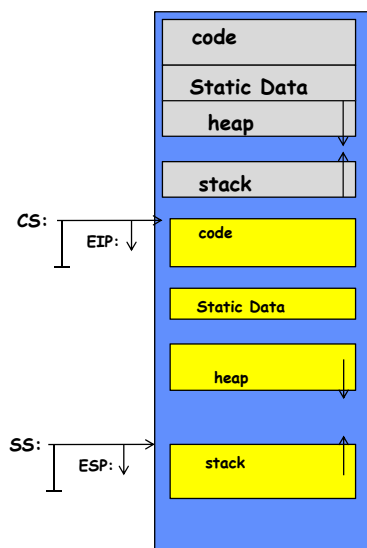
Lec 2.44

## x86 - segments and stacks

### Processor Registers

CS	EIP
SS	ESP
DS	EAX
ES	EBX
	ECX
	EDX
	ESI
	EDI

Start address, length and access rights associated with each segment



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.45

## Virtual Address Translation

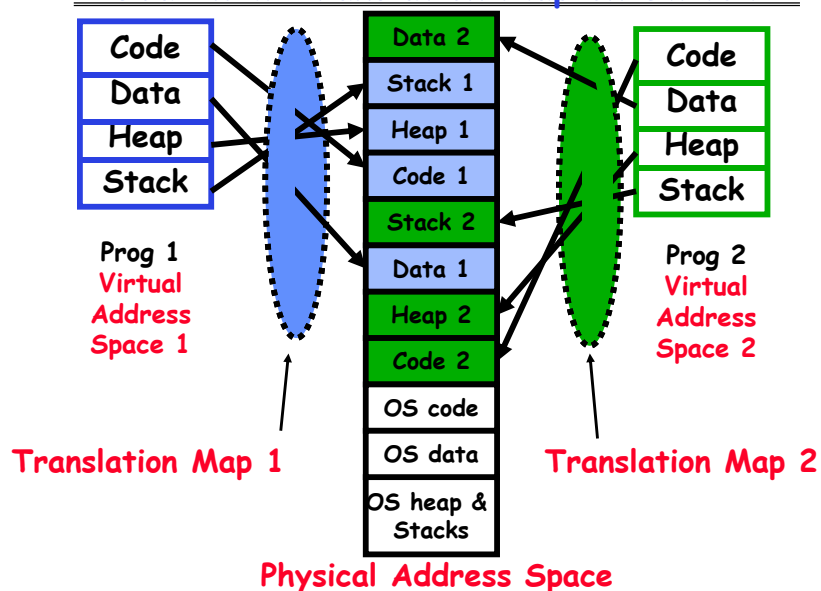
- Simpler, more useful schemes too!
- Give every process the illusion of its own BIG FLAT ADDRESS SPACE
  - Break it into pages
  - More on this later

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.46

## Providing Illusion of Separate Address Space: Load new Translation Map on Switch



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.47

## Running Many Programs ???

- We have the basic mechanism to
  - switch between user processes and the kernel,
  - the kernel can switch among user processes,
  - Protect OS from user processes and processes from each other
- Questions ???
- How do we decide which user process to run?
- How do we represent user processes in the OS?
- How do we pack up the process and set it aside?
- How do we get a stack and heap for the kernel?
- Aren't we wasting are lot of memory?
- ...

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.48

## Process Control Block

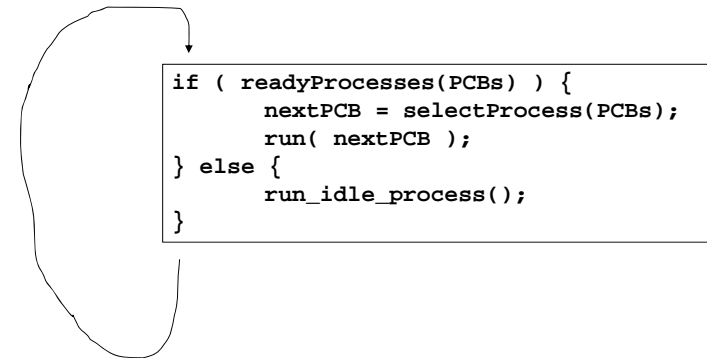
- Kernel represents each process as a process control block (PCB)
  - Status (running, ready, blocked, ...)
  - Register state (when not ready)
  - Process ID (PID), User, Executable, Priority, ...
  - Execution time, ...
  - Memory space, translation, ...
- Kernel Scheduler maintains a data structure containing the PCBs
- Scheduling algorithm selects the next one to run

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.49

## Scheduler

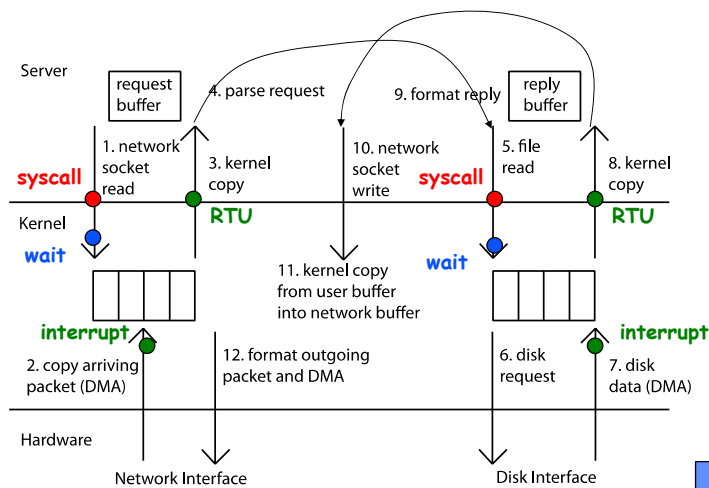


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.50

## Putting it together: web server



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.51

## Digging Deeper: Discussion & Questions

1/26/15


Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.52

## 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 technique called "Simultaneous Multithreading"
- <http://www.cs.washington.edu/research/smt/index.html>
  - SPARC, Pentium 4/Xeon ("Hyperthreading"), Power 5



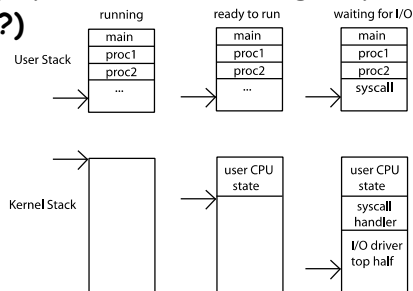
Colored blocks show instructions executed

## Implementing Safe Mode Transfers

- Carefully constructed kernel code packs up the user process state and sets it aside.
  - Details depend on the machine architecture
- Should be impossible for buggy or malicious user program to cause the kernel to corrupt itself.
- Interrupt processing not be visible to the user process:
  - Occurs between instructions, restarted transparently
  - No change to process state
  - What can be observed even with perfect interrupt processing?

## Kernel Stack Challenge

- Kernel needs space to work
- Cannot put anything on the user stack (Why?)
- Two-stack model
  - OS thread has interrupt stack (located in kernel memory) plus User stack (located in user memory)
  - Syscall handler copies user args to kernel space before invoking specific function (e.g., open)
  - Interrupts (???)



## Hardware support: Interrupt Control

- Interrupt Handler invoked with interrupts 'disabled'
  - Re-enabled upon completion
  - Non-blocking (run to completion, no waits)
  - Pack it up in a queue and pass off to an OS thread to do the hard work
    - » wake up an existing OS thread
- OS kernel may enable/disable interrupts
  - On x86: CLI (disable interrupts), STI (enable)
  - Atomic section when select next process/thread to run
  - Atomic return from interrupt or syscall
- HW may have multiple levels of interrupt
  - Mask off (disable) certain interrupts, eg., lower priority
  - Certain non-maskable-interrupts (nmi)
    - » e.g., kernel segmentation fault

## How do we take interrupts safely?

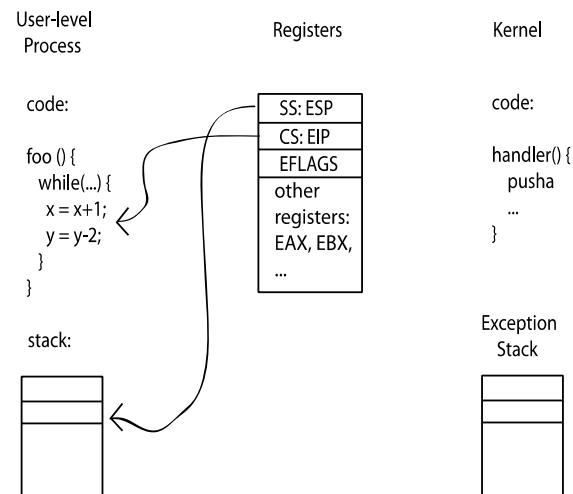
- **Interrupt vector**
  - Limited number of entry points into kernel
- **Kernel interrupt stack**
  - Handler works regardless of state of user code
- **Interrupt masking**
  - Handler is non-blocking
- **Atomic transfer of control**
  - "Single instruction"-like to change:
    - » Program counter
    - » Stack pointer
    - » Memory protection
    - » Kernel/user mode
- **Transparent restartable execution**
  - User program does not know interrupt occurred

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.57

## Before

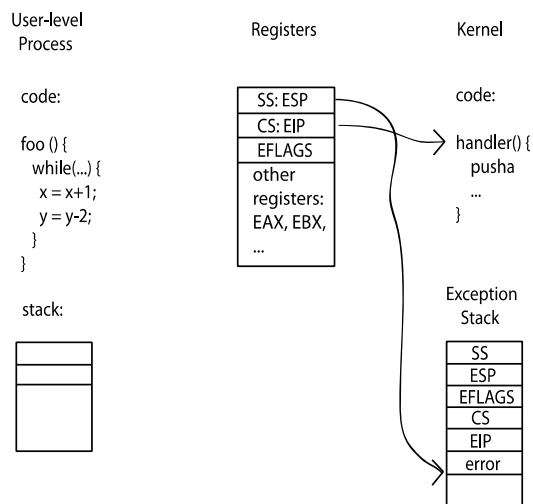


1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.58

## During



1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.59

## Kernel System Call Handler

- **Locate arguments**
  - In registers or on user(!) stack
- **Copy arguments**
  - From user memory into kernel memory
  - Protect kernel from malicious code evading checks
- **Validate arguments**
  - Protect kernel from errors in user code
- **Copy results back**
  - into user memory

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.60

## Multiprocessors - Multicores - Multiple Threads

---

- What do we need to support Multiple Threads
  - Multiple kernel threads?
  - Multiple user threads in a process?
- What if we have multiple Processors / Cores

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.61

## Idle Loop & Power

---

- Measly do-nothing unappreciated trivial piece of code that is central to low-power

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.62

## Performance

---

- Performance = Operations / Time
- How can the OS ruin application performance?
- What can the OS do to increase application performance?

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.63

## Conclusion: Four fundamental OS concepts

---

- Thread
  - Single unique execution context
  - Program Counter, Registers, Execution Flags, Stack
- Address Space w/ Translation
  - Programs execute in an *address space* that is distinct from the memory space of the physical machine
- Process
  - An instance of an executing program is a *process consisting of an address space and one or more threads of control*
- Dual Mode operation/Protection
  - Only the "system" has the ability to access certain resources
  - The OS and the hardware are protected from user programs and user programs are isolated from one another by *controlling the translation* from program virtual addresses to machine physical addresses

1/26/15

Kubiatowicz CS162 ©UCB Spring 2015

Lec 2.64