

# **CSE 153**

# **Design of Operating Systems**

**Fall 18**

Lecture 2: OS model and Architectural  
Support

# Last time/Today

---

- | Historic evolution of Operating Systems (and computing!)
- | Today:
  - ▣ We start our journey in exploring Operating Systems
  - ▣ Try to answer questions such as:
    - » What is the OS?
    - » What does it need to do?
    - » How/When does the OS run?
    - » How do programs interact with it?
    - » How is this supported by CPUs?

# Some questions to get you thinking

---

- What is the OS? Software?
- Is the OS always executing?
  - ◆ If not, how do we make sure it gets to run?
- How do we prevent user programs from directly manipulating hardware?

# Sleeping Beauty Model

---

- Answer: Sleeping beauty model
  - ◆ Technically known as *controlled direct execution*
  - ◆ OS runs in response to “events”; we support the switch in hardware
  - ◆ Only the OS can manipulate hardware or critical system state
  
- Most of the time the OS is sleeping
  - ◆ Good! Less overhead
  - ◆ Good! Applications are running directly on the hardware

# What do we need from the architecture/CPU?

---

- Manipulating privileged machine state
  - ◆ Protected instructions
  - ◆ Manipulate device registers, TLB entries, etc.
  - ◆ Controlling access
- Generating and handling “events”
  - ◆ Interrupts, exceptions, system calls, etc.
  - ◆ Respond to external events
  - ◆ CPU requires software intervention to handle fault or trap
- Other stuff
  - ◆ Mechanisms to handle concurrency, Isolation, virtualization ...

# Types of Arch Support

---

- Manipulating privileged machine state
  - ◆ Protected instructions
  - ◆ Manipulate device registers, TLB entries, etc.
  - ◆ Controlling access
  
- Generating and handling “events”
  - ◆ Interrupts, exceptions, system calls, etc.
  - ◆ Respond to external events
  - ◆ CPU requires software intervention to handle fault or trap
  
- Other stuff
  - ◆ Interrupts, atomic instructions, isolation

# Protected Instructions

---

- OS must have exclusive access to hardware and critical data structures
- Only the operating system can
  - ◆ Directly access I/O devices (disks, printers, etc.)
    - » Security, fairness (why?)
  - ◆ Manipulate memory management state
    - » Page table pointers, page protection, TLB management, etc.
  - ◆ Manipulate protected control registers
    - » Kernel mode, interrupt level
  - ◆ Halt instruction (why?)

# Privilege mode

---

- Hardware restricts privileged instructions to OS
- Q: How does the HW know if the executed program is OS?
  - ◆ HW must support (at least) two execution modes: OS (kernel) mode and user mode
- Mode kept in a status bit in a protected control register
  - ◆ User programs execute in user mode
  - ◆ OS executes in kernel mode (OS == “kernel”)
  - ◆ CPU checks mode bit when protected instruction executes
  - ◆ Attempts to execute in user mode trap to OS



# Switching back and forth

---

- Going from higher privilege to lower privilege
  - ◆ Easy: can directly modify the mode register to drop privilege
- But how do we escalate privilege?
  - ◆ Special instructions to change mode
    - » System calls (`int 0x80, syscall, svc`)
    - » Saves context and invokes designated handler
      - You jump to the privileged code; you cannot execute your own
    - » OS checks your syscall request and honors it only if safe
  - ◆ Or, some kind of event happens in the system

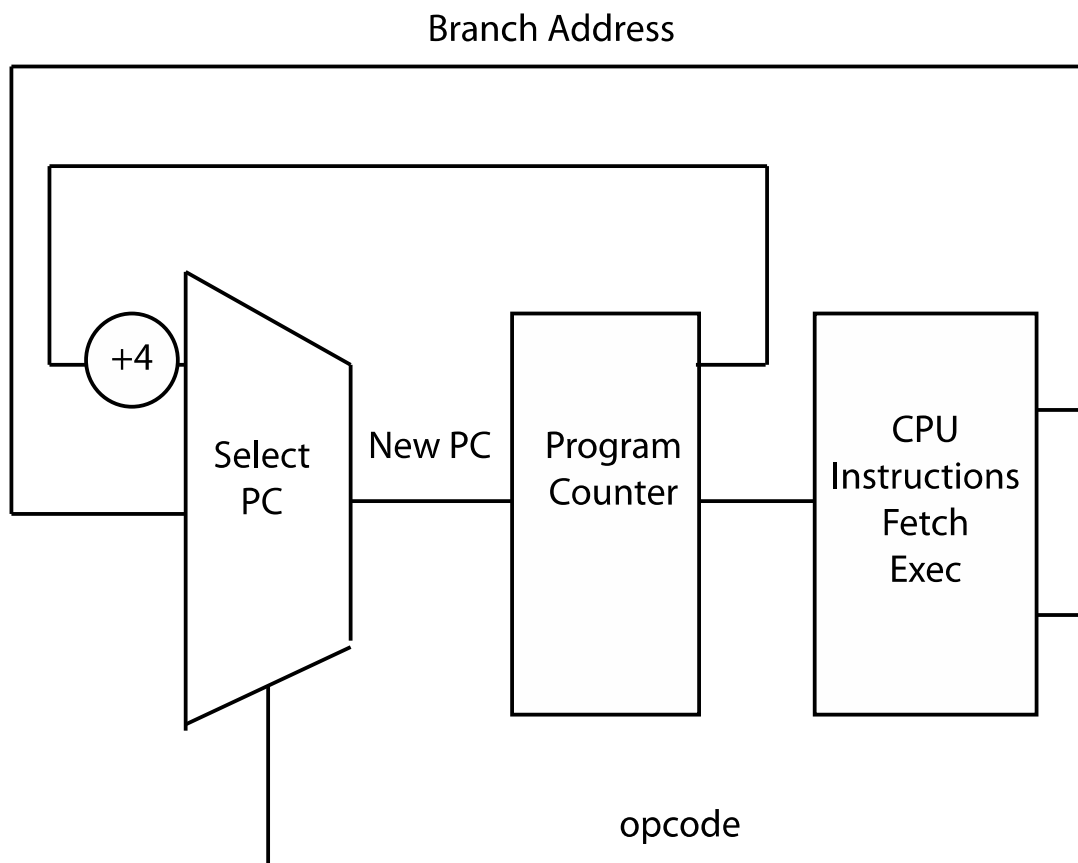
# Types of Arch Support

---

- Manipulating privileged machine state
  - ◆ Protected instructions
  - ◆ Manipulate device registers, TLB entries, etc.
  - ◆ Controlling access
- **Generating and handling “events”**
  - ◆ Interrupts, exceptions, system calls, etc.
  - ◆ Respond to external events
  - ◆ CPU requires software intervention to handle fault or trap
- Other stuff

# Review: Computer Organization

---

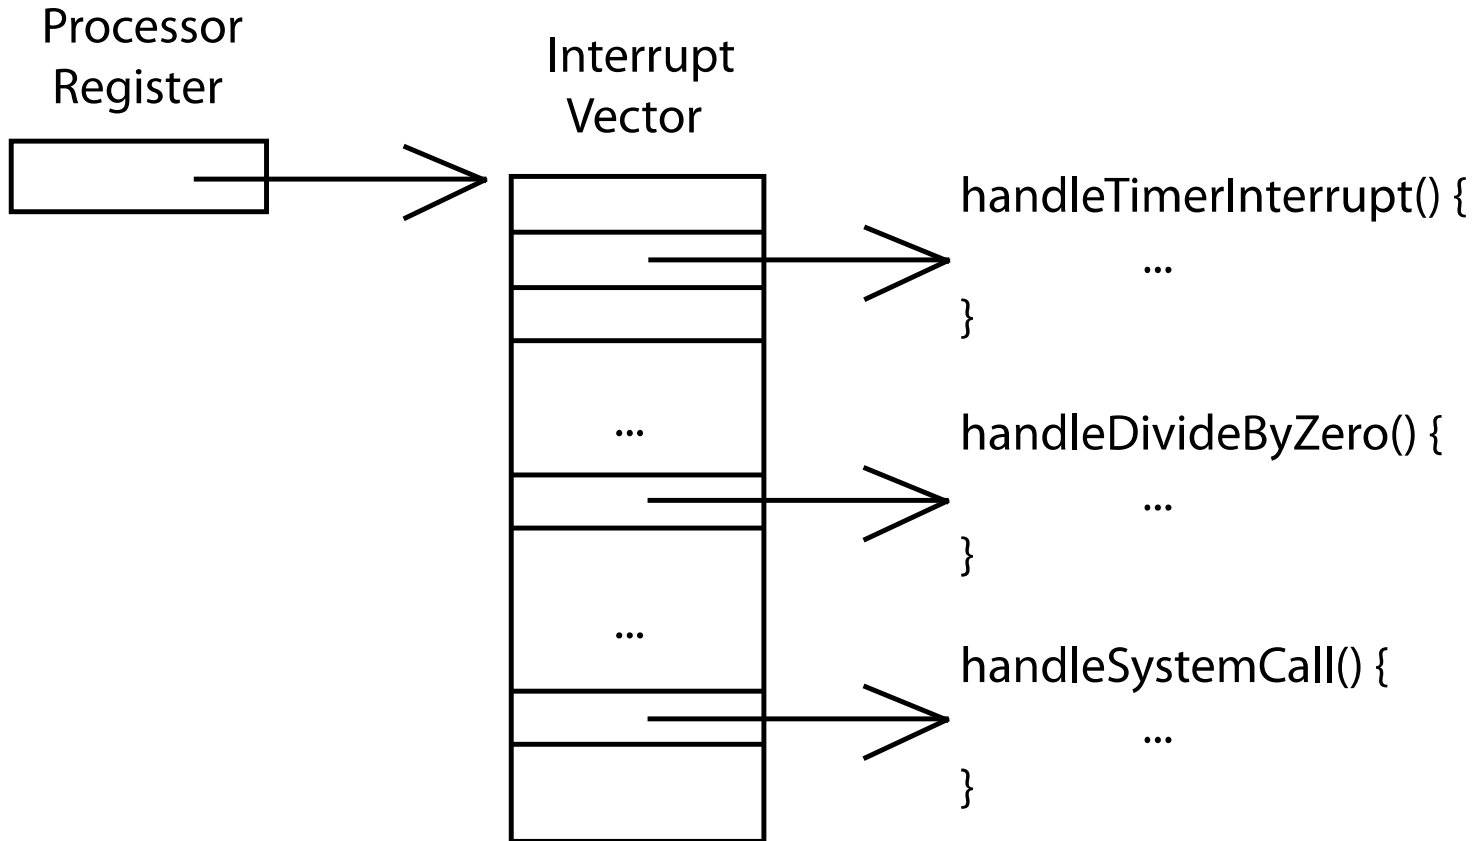


# Events

---

- An event is an “unnatural” change in control flow
  - ◆ Events immediately stop current execution
  - ◆ Changes mode, context (machine state), or both
- The kernel defines a handler for each event type
  - ◆ Event handlers always execute in kernel mode
  - ◆ The specific types of events are defined by the machine
- Once the system is booted, OS is one big event handler
  - ◆ all entry to the kernel occurs as the result of an event


# Handling events – Interrupt vector table



# Categorizing Events

---

- Two *kinds* of events: **synchronous** and **asynchronous**
- Sync events are caused by executing instructions
  - ◆ Example?
- Async events are caused by an external event
  - ◆ Example?



# Categorizing Events

---

- Two *kinds* of events: **synchronous** and **asynchronous**
  - ◆ Sync events are caused by executing instructions
  - ◆ Async events are caused by an external event
- Two *reasons* for events: **unexpected** and **deliberate**
- Unexpected events are, well, unexpected
  - ◆ Example?
- Deliberate events are scheduled by OS or application
  - ◆ Why would this be useful?

# Categorizing Events

---

- This gives us a convenient table:

	Unexpected	Deliberate
Synchronous	fault	syscall trap
Asynchronous	interrupt	signal

- Terms may be slightly different by OS and architecture
  - ◆ E.g., POSIX signals, asynch system traps, async or deferred procedure calls



# Faults

---

- Hardware detects and reports “exceptional” conditions
  - ◆ Page fault, memory access violation (unaligned, permission, not mapped, bounds...), illegal instruction, divide by zero
  
- Upon exception, hardware “faults” (verb)
  - ◆ Must save state (PC, regs, mode, etc.) so that the faulting process can be restarted
  - ◆ Invokes registered handler

# Handling Faults

---

- Some faults are handled by “fixing” the exceptional condition and returning to the faulting context
  - ◆ Page faults cause the OS to place the missing page into memory
  - ◆ Fault handler resets PC of faulting context to re-execute instruction that caused the page fault

# Handling Faults

---

- The kernel may handle unrecoverable faults by killing the user process
  - ◆ Program fault with no registered handler
  - ◆ Halt process, write process state to file, destroy process
  - ◆ In Unix, the default action for many signals (e.g., SIGSEGV)
  
- What about faults in the kernel?
  - ◆ Dereference NULL, divide by zero, undefined instruction
  - ◆ These faults considered fatal, operating system crashes
  - ◆ **Unix panic**, **Windows “Blue screen of death”**
    - » Kernel is halted, state dumped to a core file, machine locked up

# Categorizing Events

---

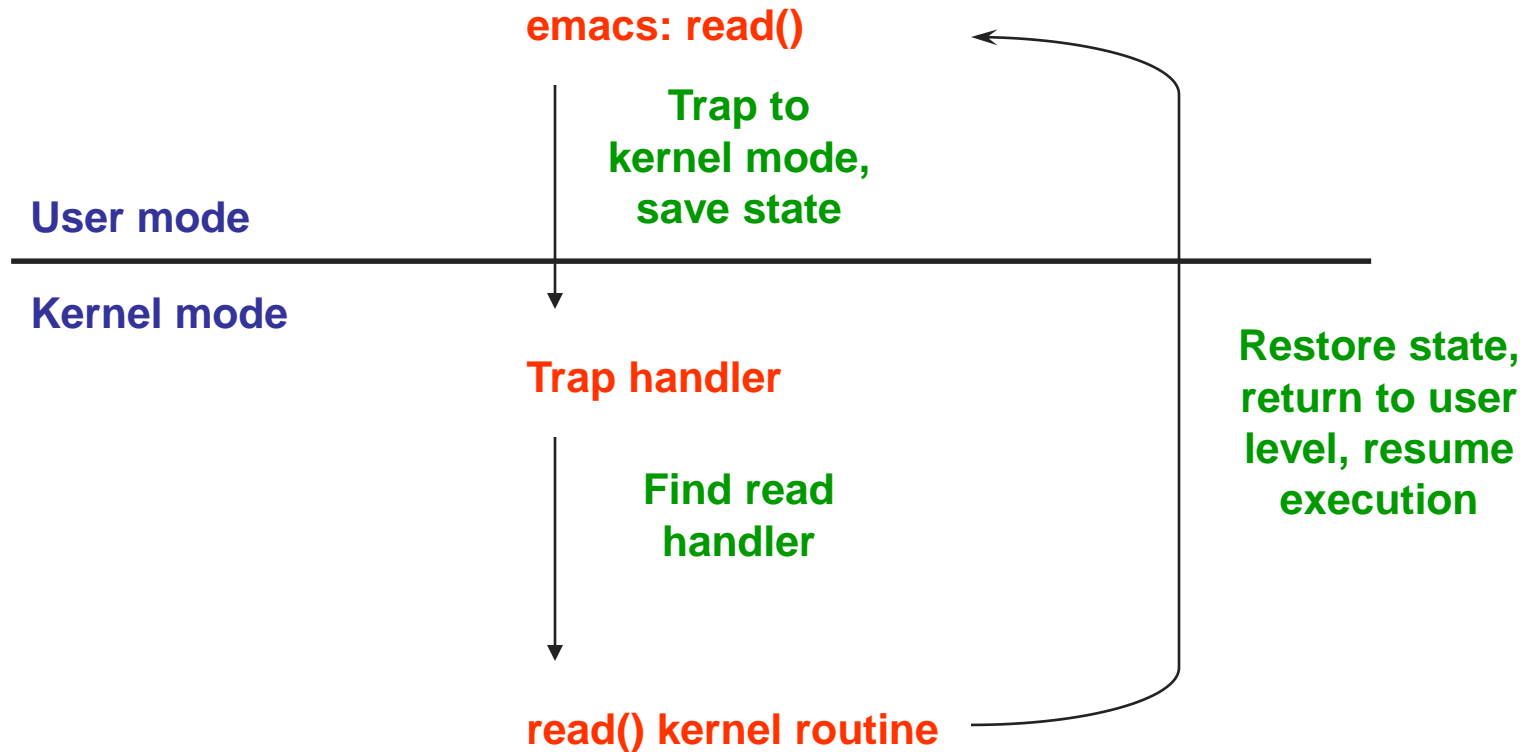
	Unexpected	Deliberate
Synchronous	fault	<b>syscall trap</b>
Asynchronous	interrupt	signal

# System Calls

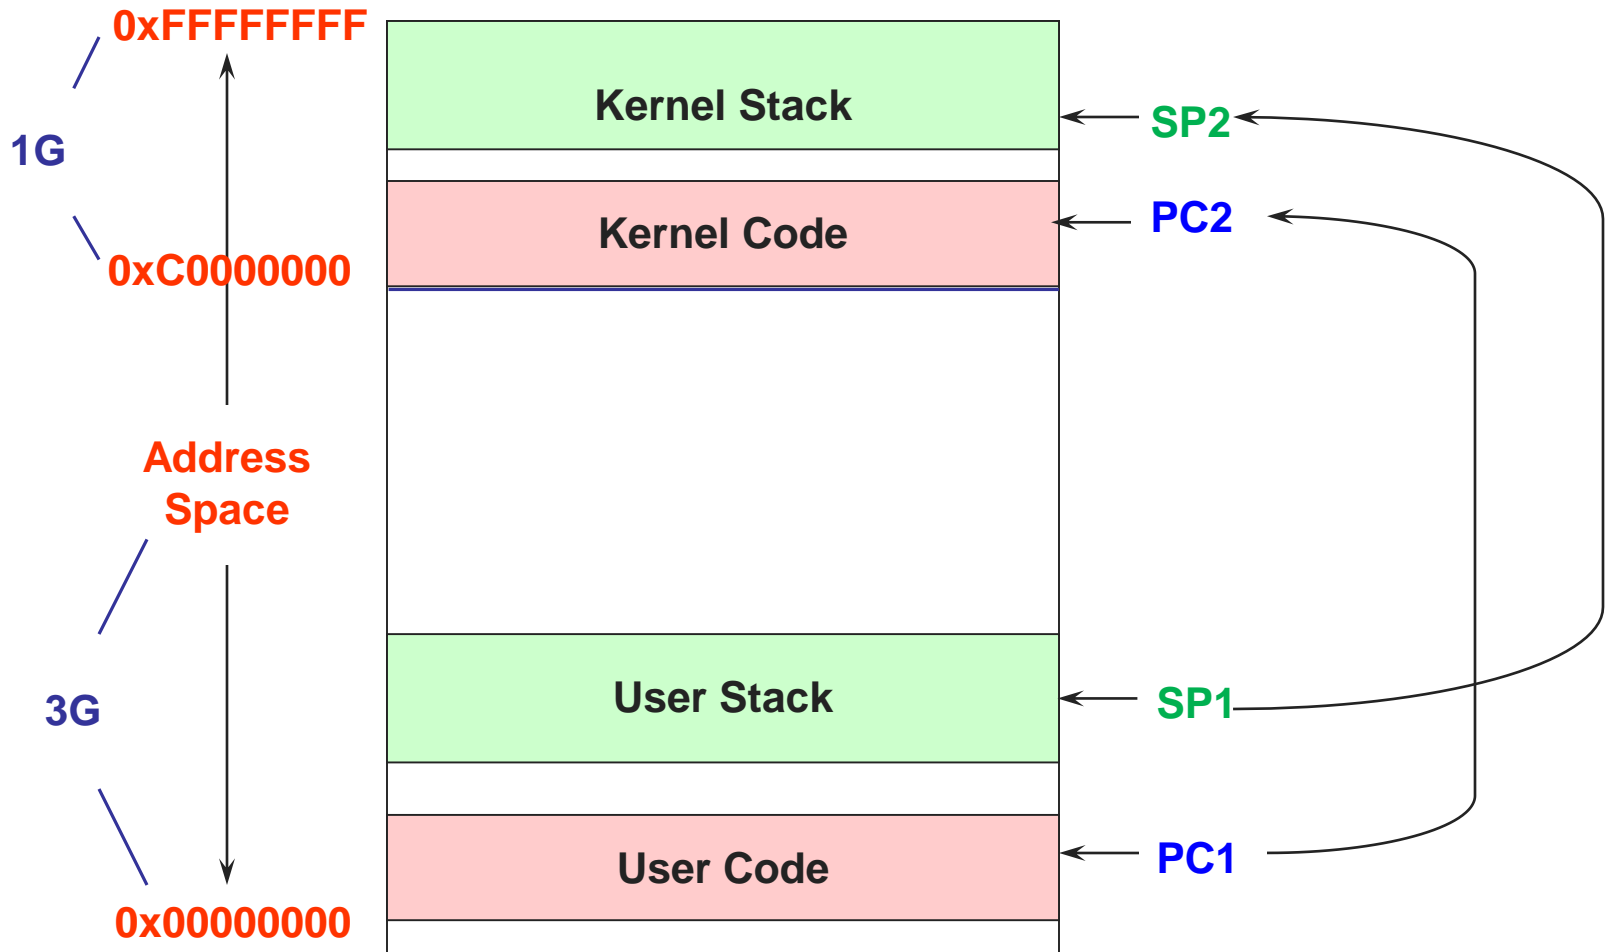
---

- For a user program to do something “privileged” (e.g., I/O) it must call an OS procedure
  - ◆ Known as **crossing the protection boundary**, or a **protected procedure call**
- Hardware provides a **system call** instruction that:
  - ◆ Causes an exception, which invokes a kernel handler
    - » Passes a parameter determining the system routine to call
  - ◆ Saves caller state (PC, regs, mode) so it can be restored
    - » **Why save mode?**
  - ◆ Returning from system call restores this state

# System Call



# Another view



# System Call Questions

---

- There are hundreds of syscalls. How do we let the kernel know which one we intend to invoke?
  - ◆ Before issuing **int \$0x80** or **sysenter**, set **%eax/%rax** with the syscall number
  
- System calls are like function calls, but how to pass parameters?
  - ◆ Just like calling convention in syscalls, typically passed through **%ebx, %ecx, %edx, %esi, %edi, %ebp**



# More questions

---

- How to reference kernel objects (e.g., files, sockets)?
  - ◆ Naming problem – an integer mapped to a unique object
    - » `int fd = open("file"); read(fd, buffer);`
  - ◆ Why can't we reference the kernel objects by memory address?

# System calls in xv6

---

- Look at trap.h and trap.c
  - ◆ Interrupt handlers are initialized in two arrays (idt and vectors)
    - » Tvinit() function does the initialization
  - ◆ Syscalls have a single trap handler (T\_SYSCALL, 64)
  - ◆ Trap() handles all exceptions, including system calls
    - » If the exception is a system call, it calls syscall()
- Keep digging from there to understand how system calls are supported
  - ◆ You will be adding a new system call in Lab 1

# Categorizing Events

---

	Unexpected	Deliberate
Synchronous	fault	syscall trap
Asynchronous	<b>interrupt</b>	software interrupt

- **Interrupts signal asynchronous events**
  - ◆ I/O hardware interrupts
  - ◆ Software and hardware timers