CSE 120
Principles of Operating Systems

Fall 2004

Lecture 3: Architectural Support for Operating Systems

Geoffrey M. Voelker

Administrivia

- Mailing list
  - You should be getting mail on the list. If not, let me know.
- Homework #1
  - Due 10/5
- Project groups
  - Send your group info to Charles (cplucas@cs.ucsd.edu)
- Office hours
  - Mon 3-4pm, Thu 4-5pm

Why Start With Architecture?

- Operating system functionality fundamentally depends upon the architectural features of the computer
- Architectural support can greatly simplify – or complicate – OS tasks
  - Early PC operating systems (DOS, MacOS) lacked virtual memory in part because the architecture did not support it
  - Early Sun 1 computers used two M68000 CPUs to implement virtual memory (M68000 did not have VM hardware support)

Architectural Features for OS

- Features that directly support the OS include
  - Protection (kernel/user mode)
  - Protected instructions
  - Memory protection
  - System calls
  - Interrupts and exceptions
  - Timer (clock)
  - I/O control and operation
  - Synchronization (atomic instructions)

Types of Arch Support

- Manipulating privileged machine state
  - Protected instructions
  - Manipulate device registers, TLB entries, etc.
- Generating and handling "events"
  - Interrupts, exceptions, system calls, etc.
  - Respond to external events
  - CPU requires software intervention to handle fault or trap

Protected Instructions

- A subset of instructions of every CPU is restricted to use only by the OS
  - Known as protected (privileged) instructions
- 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?)

OS Protection

- How do we know if we can execute a protected instruction?
    - Architecture must support (at least) two modes of operation: kernel mode and user mode
        - » VAX, x86 support four modes; earlier archs (Multics) even more
        - » Why? Protect the OS from itself (software engineering)
    - Mode is indicated by a status bit in a protected control register
    - User programs execute in user mode
    - OS executes in kernel mode (OS == "kernel")
- Protected instructions only execute in kernel mode
    - CPU checks mode bit when protected instruction executes
    - Setting mode bit must be a protected instruction

Memory Protection

- OS must be able to protect programs from each other
- OS must protect itself from user programs
- May or may not protect user programs from OS
- Memory management hardware provides memory protection mechanisms
    - Base and limit registers
    - Page table pointers, page protection, TLB
    - Virtual memory
    - Segmentation
- Manipulating memory management hardware uses protected (privileged) operations

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, all entry to the kernel occurs as the result of an event
  - In effect, the operating system is one big event handler

Categorizing Events

- Two kinds of events, interrupts and exceptions
- Exceptions are caused by executing instructions
  - CPU requires software intervention to handle a fault or trap
- Interrupts are caused by an external event
  - Device finishes I/O, timer expires, etc.

- Two reasons for events, unexpected and deliberate
- Unexpected events are, well, unexpected
  - What is an example?
- Deliberate events are scheduled by OS or application
  - Why would this be useful?

Categorizing Events (2)

- This gives us a convenient table:

	Unexpected	Deliberate
Exceptions (sync)	fault	syscall trap
Interrupts (async)	interrupt	software interrupt

- Terms may be used slightly differently by various OSes, CPU architectures…
- Software interrupt – a.k.a. async system trap (AST), async or deferred procedure call (APC or DPC)
- Will cover faults, system calls, and interrupts next
  - Does anyone remember from CSE 141 what a software interrupt is?

Faults

- Hardware detects and reports "exceptional" conditions
  - Page fault, unaligned access, divide by zero
- Upon exception, hardware "faults" (verb)
  - Must save state (PC, regs, mode, etc.) so that the faulting process can be restarted
- Modern OSes use VM faults for many functions
  - Debugging, distributed VM, GC, copy-on-write
- Fault exceptions are a performance optimization
  - Could detect faults by inserting extra instructions into code (at a significant performance penalty)

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
- Some faults are handled by notifying the process
  - Fault handler changes the saved context to transfer control to a user-mode handler on return from fault
  - Handler must be registered with OS
  - Unix signals or NT user-mode Async Procedure Calls (APCs)
    » SIGALRM, SIGHUP, SIGTERM, SIGSEGV, etc.

Handling Faults (2)

- 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"

System Calls
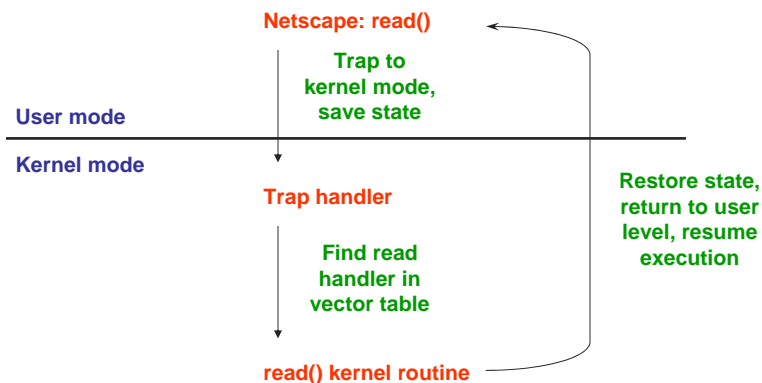
- For a user program to do something "privileged" (e.g., I/O) it must call an OS procedure
  - Known as crossing the protection boundary, or a protected procedure call
- Arch provides a system call instruction that:
  - Causes an exception, which vectors to a kernel handler
  - Passes a parameter determining 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
- Arch must permit OS to:
  - Verify input parameters (e.g., valid addresses for buffers)
  - Restore saved state, reset mode, resume execution

System Call

Netscape: read()

Trap to kernel mode, save state

User mode

Kernel mode

Trap handler

Find read handler in vector table

read() kernel routine

Restore state, return to user level, resume execution

8

System Call Questions

- What would happen if kernel did not save state?
- What if the kernel executes a system call?
- What if a user program returns from a system call?
- How to reference kernel objects as arguments or results to/from system calls?
  - A naming issue
  - Use integer object handles or descriptors
    - » E.g., Unix file descriptors
  - Also called capabilities (more later)
  - Why not use kernel addresses to name kernel objects?

Interrupts

- Interrupts
  - Precise interrupts – CPU transfers control only on instruction boundaries
  - Imprecise interrupts – CPU transfers control in the middle of instruction execution
  - OS designers like precise interrupts, CPU designers like imprecise interrupts
    - » Why?
- Interrupts signal asynchronous events
  - Timer, I/O, etc.

Timer

- The timer is critical for an operating system
- It is the fallback mechanism by which the OS reclaims control over the machine
  - Timer is set to generate an interrupt after a period of time
    - » Setting timer is a privileged instruction
  - When timer expires, generates an interrupt
- Prevents infinite loops
  - OS can always regain control from erroneous or malicious programs that try to hog CPU
- Also used for time-based functions (e.g., *sleep()*)

I/O Control

- I/O issues
  - Initiating an I/O
  - Completing an I/O
- Initiating an I/O
  - Special instructions
  - Memory-mapped I/O
    - » Device registers mapped into address space
    - » Writing to address sends data to I/O device

I/O Completion

- Interrupts are the basis for asynchronous I/O
  - OS initiates I/O
  - Device operates independently of rest of machine
  - Device sends an interrupt signal to CPU when done
  - OS maintains a vector table containing a list of addresses of kernel routines to handle various events
  - CPU looks up kernel address indexed by interrupt number, context switches to routine
- If you have ever installed Windows, you now know what IRQs are for

I/O Example

1. Ethernet receives packet, writes packet into memory
2. Ethernet signals an interrupt
3. CPU stops current operation, switches to kernel mode, saves machine state (PC, mode, etc.) on kernel stack
4. CPU reads address from vector table indexed by interrupt number, branches to address (Ethernet device driver)
5. Ethernet device driver processes packet (reads device registers to find packet in memory)
6. Upon completion, restores saved state from stack

Interrupt Questions

- Interrupts halt the execution of a process and transfer control (execution) to the operating system
  - Can the OS be interrupted? (Consider why there might be different IRQ levels)
- Interrupts are used by devices to have the OS do stuff
  - What is an alternative approach to using interrupts?
  - What are the drawbacks of that approach?

Synchronization

- Interrupts cause difficult problems
  - An interrupt can occur at any time
  - A handler can execute that interferes with code that was interrupted
- OS must be able to synchronize concurrent execution
- Need to guarantee that short instruction sequences execute atomically
  - Disable interrupts – turn off interrupts before sequence, execute sequence, turn interrupts back on
  - Special atomic instructions – read/modify/write a memory address, test and conditionally set a bit based upon previous value

Summary

- Protection
  - User/kernel modes
  - Protected instructions
- System calls
  - Used by user-level processes to access OS functions
  - Access what is "in" the OS
- Exceptions
  - Unexpected event during execution (e.g., divide by zero)
- Interrupts
  - Timer, I/O

Next Time...

- Read Chapter 4 (Processes)
- Homework #1 assigned
  - Due Tue 10/5