

CSE 451: Operating Systems Winter 2007

Module 2 Architectural Support for Operating Systems

Brian Bershad
Bershad@cs.washington.edu
562 Allen Center

Architectural features affecting OS's

- These features were built primarily to support OS's:
 - timer (clock) operation
 - synchronization instructions (e.g., atomic test-and-set)
 - memory protection
 - I/O control operations
 - interrupts and exceptions
 - protected modes of execution (kernel vs. user)
 - privileged instructions
 - system calls (and software interrupts)
- [2006] virtualization architectures (aka Intel discovers the early 1970s)
 - Intel: <http://www.intel.com/technology/itj/2006/v10i3/1-hardware/1-abstract.htm>
 - AMD: <http://enterprise.amd.com/us-en/AMD-Business/Business-Solutions/Consolidation/Virtualization.aspx>

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

2

Privileged instructions

- some instructions are restricted to the OS
 - known as **protected** or **privileged** instructions
- e.g., only the OS can:
 - directly access I/O devices (disks, network cards)
 - why?
 - manipulate memory state management
 - page table pointers, TLB loads, etc.
 - why?
 - manipulate special 'mode bits'
 - interrupt priority level
 - why?
 - halt instruction
 - why?

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

3

OS protection

- So how does the processor know if a privileged instruction should be executed?
 - the architecture must support at least two modes of operation: **kernel mode** and **user mode**
 - VAX, x86 support 4 protection modes
 - mode is set by status bit in a protected processor register
 - user programs execute in user mode
 - OS executes in kernel mode (OS == kernel)
- Privileged instructions can only be executed in kernel mode
 - what happens if user mode attempts to execute a privileged instruction?

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

4

Crossing protection boundaries

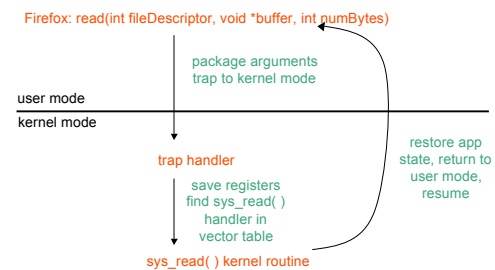
- So how do user programs do something privileged?
 - e.g., how can you write to a disk if you can't execute I/O instructions?
- User programs must call an OS procedure
 - OS defines a sequence of **system calls**
 - how does the user-mode to kernel-mode transition happen?
- There must be a system call instruction, which:
 - causes an exception (throws a **software interrupt**), which vectors to a kernel handler
 - passes a parameter indicating which system call to invoke
 - saves caller's state (registers, mode bit) so they can be restored
 - OS must verify caller's parameters (e.g., pointers)
 - must be a way to return to user mode once done

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

5

A kernel crossing illustrated



10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

6

System call issues

- What would happen if kernel didn't save state?
- Why must the kernel verify arguments?
- How can you reference kernel objects as arguments or results to/from system calls?


10/1/07

© 2007 Bershad, Gribble, Lazowska, Levy, Zahorjan

7

Memory protection

- OS must protect user programs from each other
 - maliciousness, ineptitude
- OS must also protect itself from user programs
 - integrity and security
 - what about protecting user programs from OS?
- Simplest scheme: **base** and **limit** registers
 - are these protected?



base and limit registers are loaded by OS before starting program

10/1/07

© 2007 Bershad, Gribble, Lazowska, Levy, Zahorjan

8

More sophisticated memory protection

- coming later in the course
- paging, segmentation, virtual memory
 - page tables, page table pointers
 - translation lookaside buffers (TLBs)
 - page fault handling

10/1/07

© 2007 Bershad, Gribble, Lazowska, Levy, Zahorjan

9

OS control flow

- After the OS has booted, all entry to the kernel happens as the result of an **event**
 - event immediately stops current execution
 - changes mode to kernel mode, event handler is called
- Kernel defines handlers for each event type
 - specific types are defined by the architecture
 - e.g.: timer event, I/O interrupt, system call trap
 - when the processor receives an event of a given type, it
 - transfers control to handler within the OS
 - handler saves program state (PC, regs, etc.)
 - handler functionality is invoked
 - handler restores program state, returns to program

10/1/07

© 2007 Bershad, Gribble, Lazowska, Levy, Zahorjan

10

Interrupts and exceptions

- Two main types of events: **interrupts** and **exceptions**
 - exceptions are caused by software executing instructions
 - e.g., the x86 'int' instruction
 - e.g., a page fault, or an attempted write to a read-only page
 - an expected exception is a "trap", unexpected is a "fault"
 - interrupts are caused by hardware devices
 - e.g., device finishes I/O
 - e.g., timer fires

10/1/07

© 2007 Bershad, Gribble, Lazowska, Levy, Zahorjan

11

I/O control

- Issues:
 - how does the kernel start an I/O?
 - special I/O instructions
 - memory-mapped I/O
 - how does the kernel notice an I/O has finished?
 - polling
 - interrupts
- Interrupts are basis for asynchronous I/O
 - device performs an operation asynchronously to CPU
 - device sends an interrupt signal on bus when done
 - in memory, a **vector table** contains list of addresses of kernel routines to handle various interrupt types
 - who populates the vector table, and when?
 - CPU switches to address indicated by vector index specified by interrupt signal

10/1/07

© 2007 Bershad, Gribble, Lazowska, Levy, Zahorjan

12

Timers

- How can the OS prevent runaway user programs from hogging the CPU (infinite loops?)
 - use a hardware timer that generates a periodic interrupt
 - before it transfers to a user program, the OS loads the timer with a time to interrupt
 - “quantum” – how big should it be set?
 - when timer fires, an interrupt transfers control back to OS
 - at which point OS must decide which program to schedule next
 - very interesting policy question: we’ll dedicate a class to it
- Should the timer be privileged?
 - for reading or for writing?

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

13

Synchronization

- Interrupts cause a wrinkle:
 - may occur any time, causing code to execute that interferes with code that was interrupted
 - OS must be able to **synchronize** concurrent processes
- Synchronization:
 - guarantee that short instruction sequences (e.g., read-modify-write) execute atomically
 - one method: turn off interrupts before the sequence, execute it, then re-enable interrupts
 - architecture must support disabling interrupts
 - another method: have special complex atomic instructions
 - read-modify-write
 - test-and-set
 - load-linked store-conditional

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

14

“Concurrent programming”

- Management of concurrency and asynchronous events is biggest difference between “systems programming” and “traditional application programming”
 - modern “event-oriented” application programming is a middle ground
- Arises from the architecture
- Can be sugar-coated, but cannot be totally abstracted away
- Huge intellectual challenge
 - Unlike vulnerabilities due to buffer overruns, which are just sloppy programming

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

15

Some questions

- Why wouldn’t you want a user program to be able to access an I/O device (e.g., the disk) directly?
- OK, so what keeps this from happening? What prevents user programs from directly accessing the disk?
- So, how does a user program cause disk I/O to occur?
- What prevents a user program from scribbling on the memory of another user program?
- What prevents a user program from scribbling on the memory of the operating system?
- What prevents a user program from running away with the CPU?

10/1/07

© 2007 Bershad, Gibble, Lazowska, Levy, Zahorjan

16