





- Jun Li (junl.li@mail.utoronto.ca)
- Ali Shariat (shariat@gmail.com)

1|14|13

- Xin Tong (<u>xtong@eecg.toronto.edu</u>)
- TAs will be at the lab sessions (3 TAs on Thursday and 3 TAs on Friday)

Ding Yuan, ECE344 Operating System



























Why hardware has to support User/Kernel mode?











- 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

21

Ding Yuan. ECE344 Operating Syste

- Mechanisms to handle concurrency
 - Interrupts, atomic instructions

1|14|13



OS Protection

- Hardware must support (at least) two modes of operation: kernel mode and user mode
 - 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
 - Attempts to execute in user mode are detected and prevented

23

Ding Yuan. ECE344 Operating System

• x86: General Protection Fault

1|14|13





- Memory management hardware provides memory protection mechanisms
 - Base and limit registers
 - Page table pointers, page protection, TLB
 - Virtual memory
 - Segmentation

1|14|13

• Manipulating memory management hardware uses protected (privileged) operations

25

Ding Yuan. ECE344 Operating System









- 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, unexpectedWhat is an example?
- Deliberate events are scheduled by OS or application
 - Why would this be useful?

1/14/13

Categorizing Events			
• This gives us a con	nvenient table:		
	Unexpected	Deliberate	
Exceptions (sync)	fault	syscall trap	
Interrupts (async)	interrupt	software interrupt	
• Terms may be use architectures	ed slightly differer	ntly by various OSes, CPU	
• No need to "men	norize" all the term	S	
Software interrup or deferred proce	t – a.k.a. async sy dure call (APC or	ystem trap (AST), async DPC)	
• Will cover faults,	system calls, and	l interrupts next	
	20		

29

Ding Yuan, ECE344 Operating System







- 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 usermode handler on return from fault

33

Ding Yuan, ECE344 Operating System

- Handler must be registered with OS
- Unix signals

1|14|13

• SIGSEGV, SIGALRM, SIGTERM, etc.









System call		
open (path, flags, mode);	<pre>open: ;Linux convention: ;parameters via registers. mov eax, 5 ; syscall number for open mov ebx, path ; ebx: first parameter mov ecx, flags ; ecx: 2nd parameter mov edx, mode ; edx: 3rd parameter int 80h</pre>	
More information: http://www.int80h.org	<pre>open: ; FreeBSD convention: ; parameters via stacks. push dword mode push dword flags push dword path mov eax, 5 push dword eax ; syscall number int 80h add esp, byte 16</pre>	























- 1. Ethernet receives packet, writes packet into memory
- 2. Ethernet signals an interrupt

1/14/13

- 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)

49

Ding Yuan, ECE344 Operating System

6. Upon completion, restores saved state from stack







Summary (2)

- 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
 - When the processor receives an event of a given type, it
 - transfers control to handler within the OS
 - handler saves program state (PC, registers, etc.)
 - handler functionality is invoked

1|14|13

• handler restores program state, returns to program

53

Ding Yuan, ECE344 Operating System



May you live in Interesting Times

