



CS 423

Operating System Design:

Far too much information

about interrupts

Professor Adam Bates
Fall 2018

Goals for Today



- Learning Objectives:
 - Understand the role and types of of Interrupts
- Announcements:
 - C4 Week 2 Summaries due today! **Week 3 is out.**
 - MP0 is available on Compass! **Due Jan 29**
 - HW0 is available on Compass! **Due Jan 29**

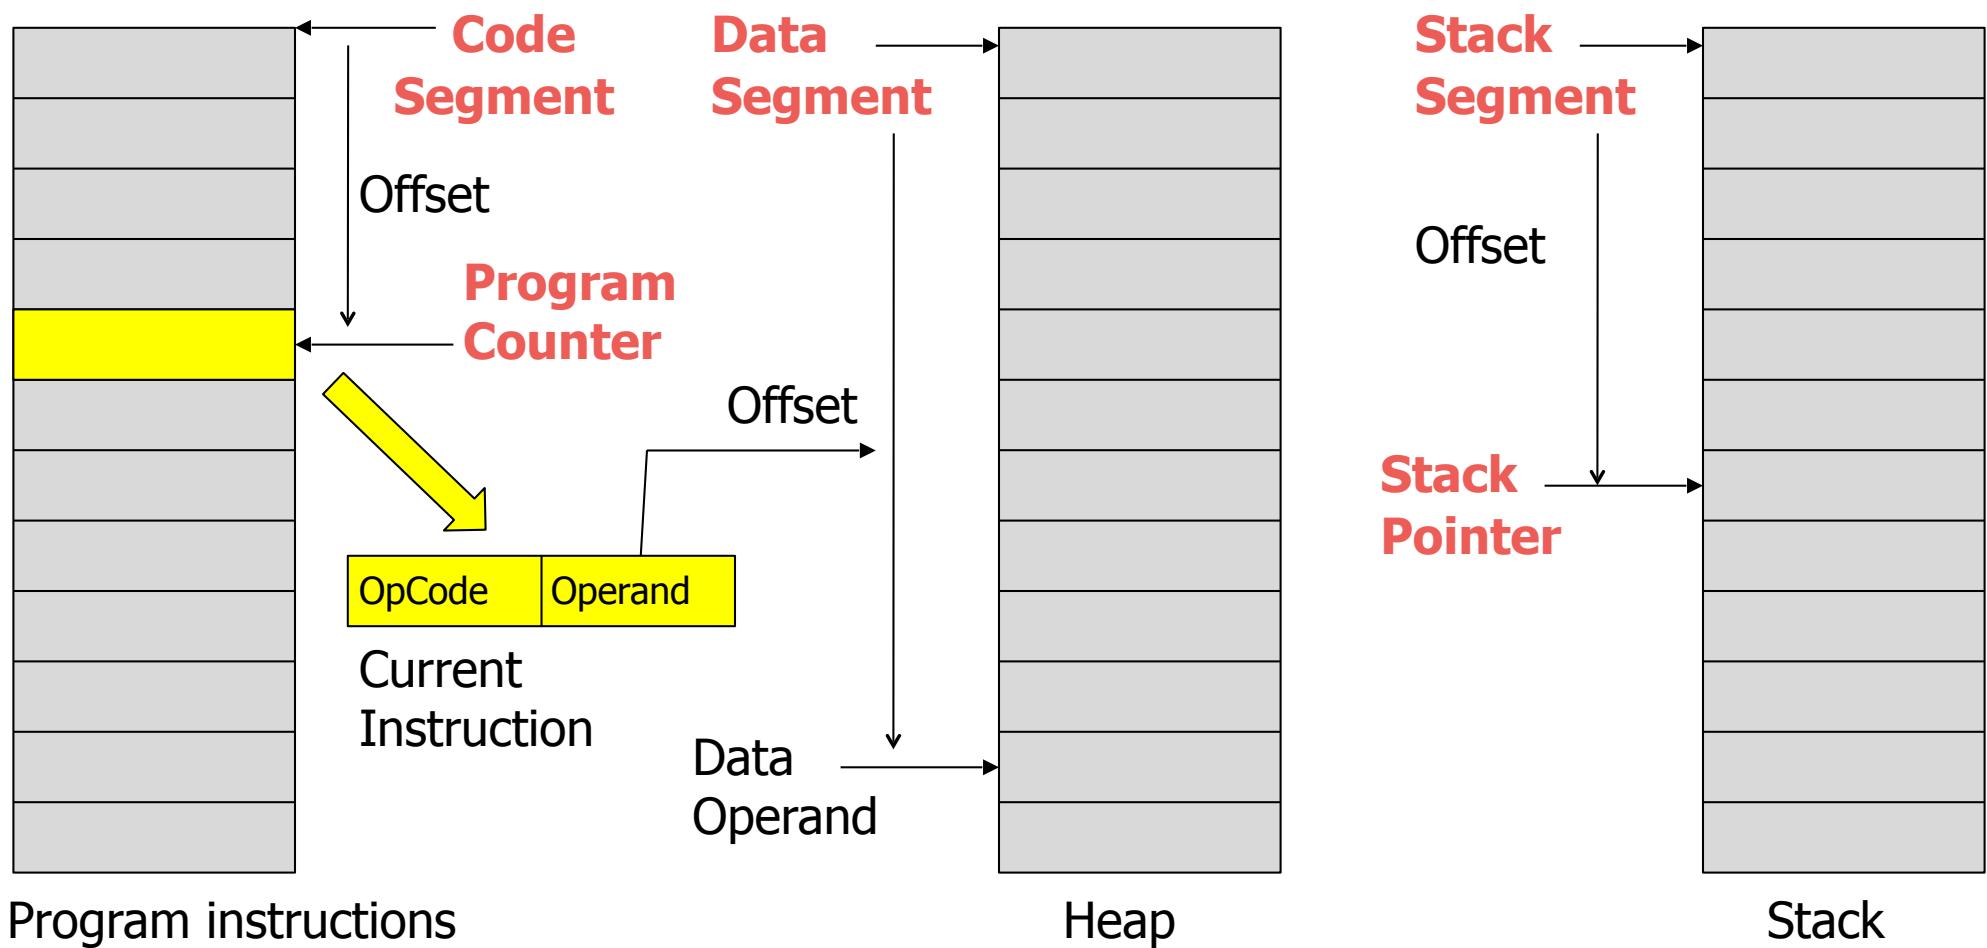


Reminder: Please put away devices at the start of class

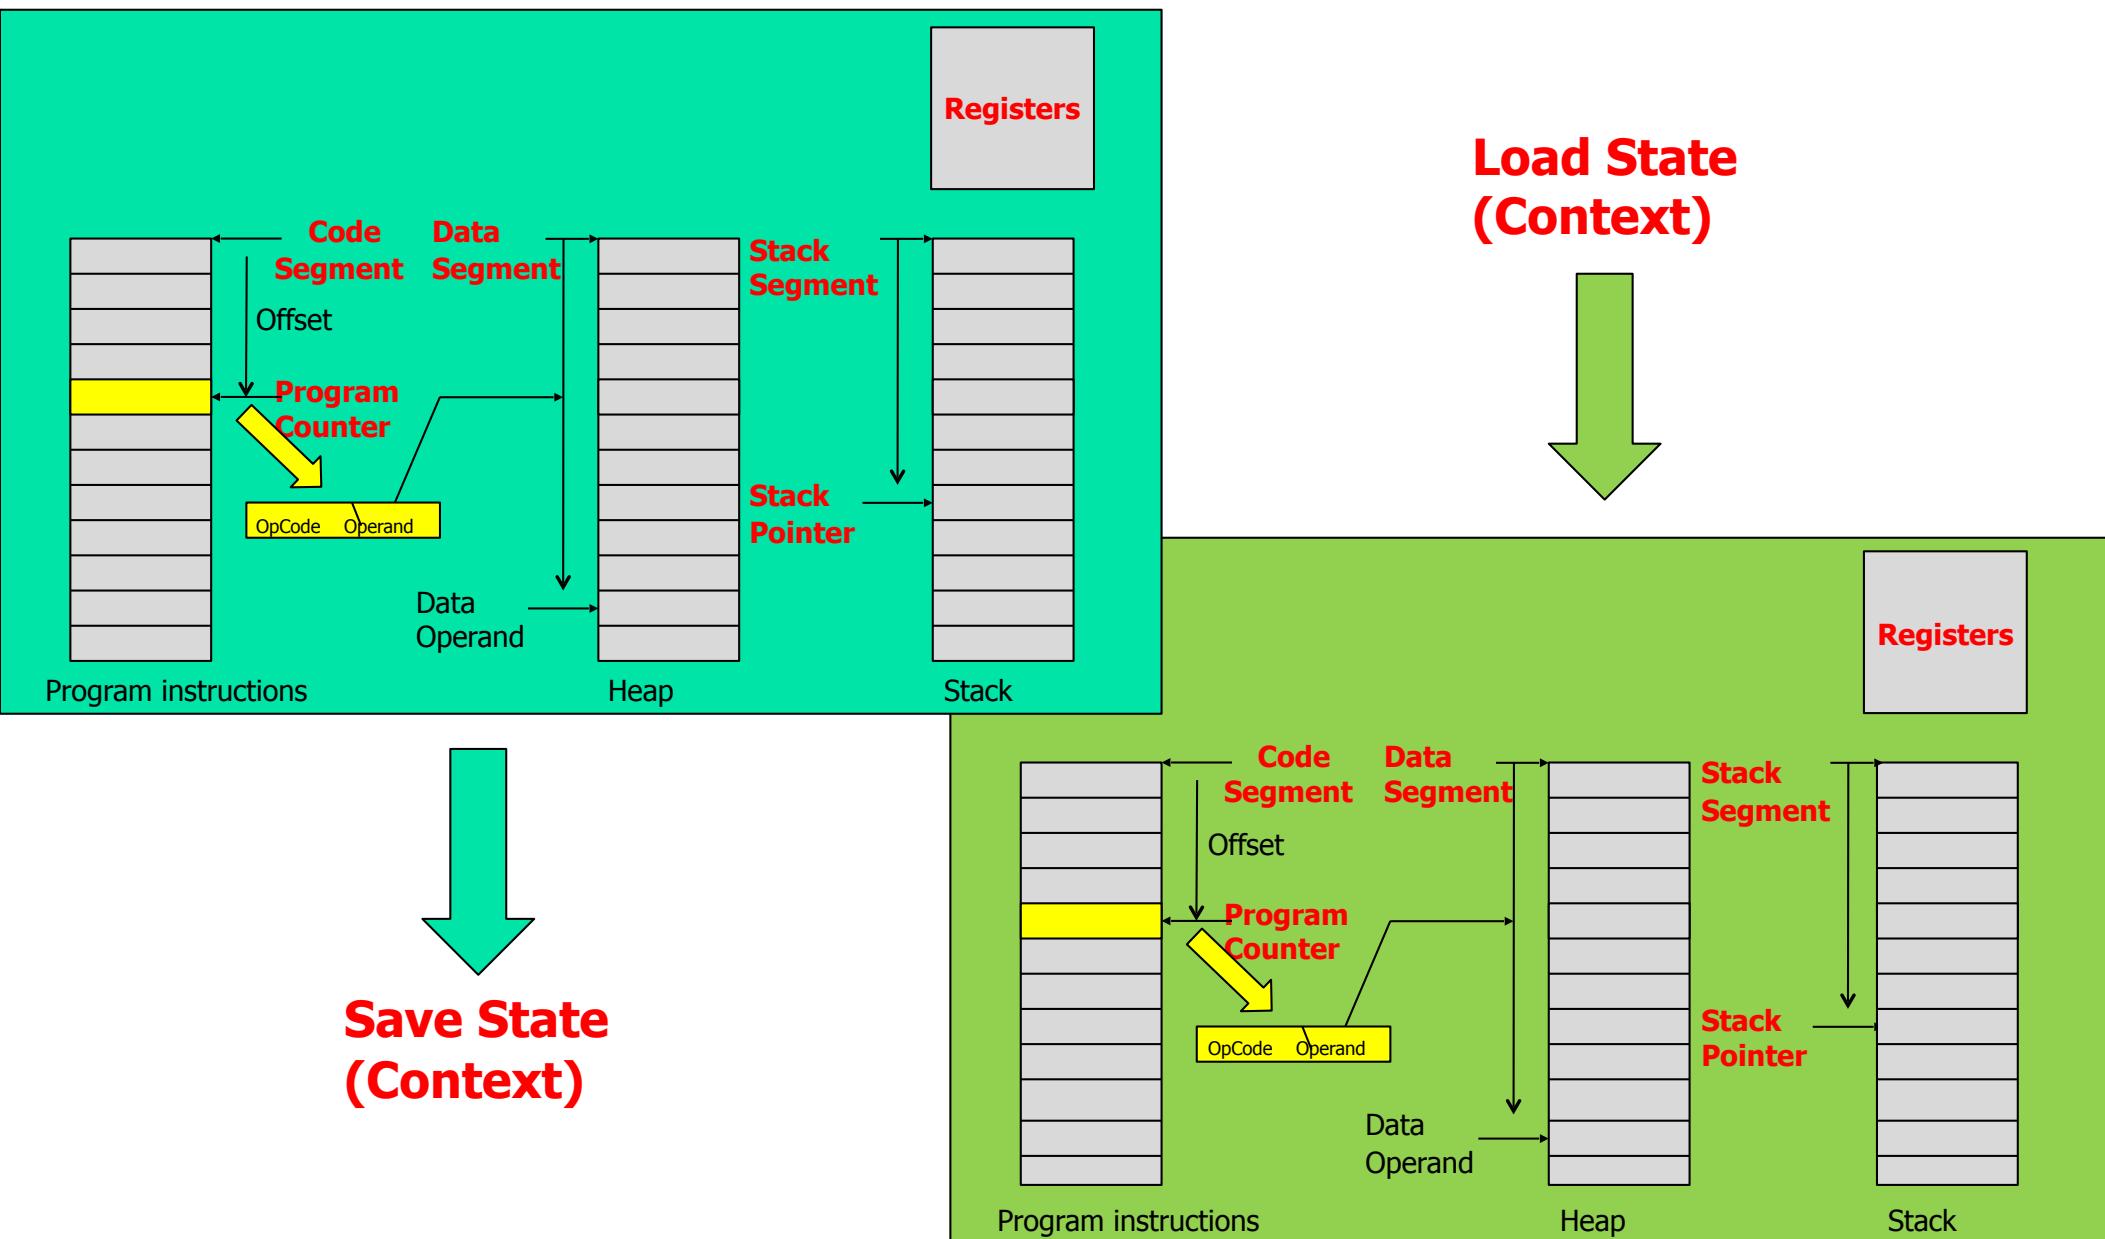
What's a 'real' CPU?



What's the **STATE** of a real CPU?

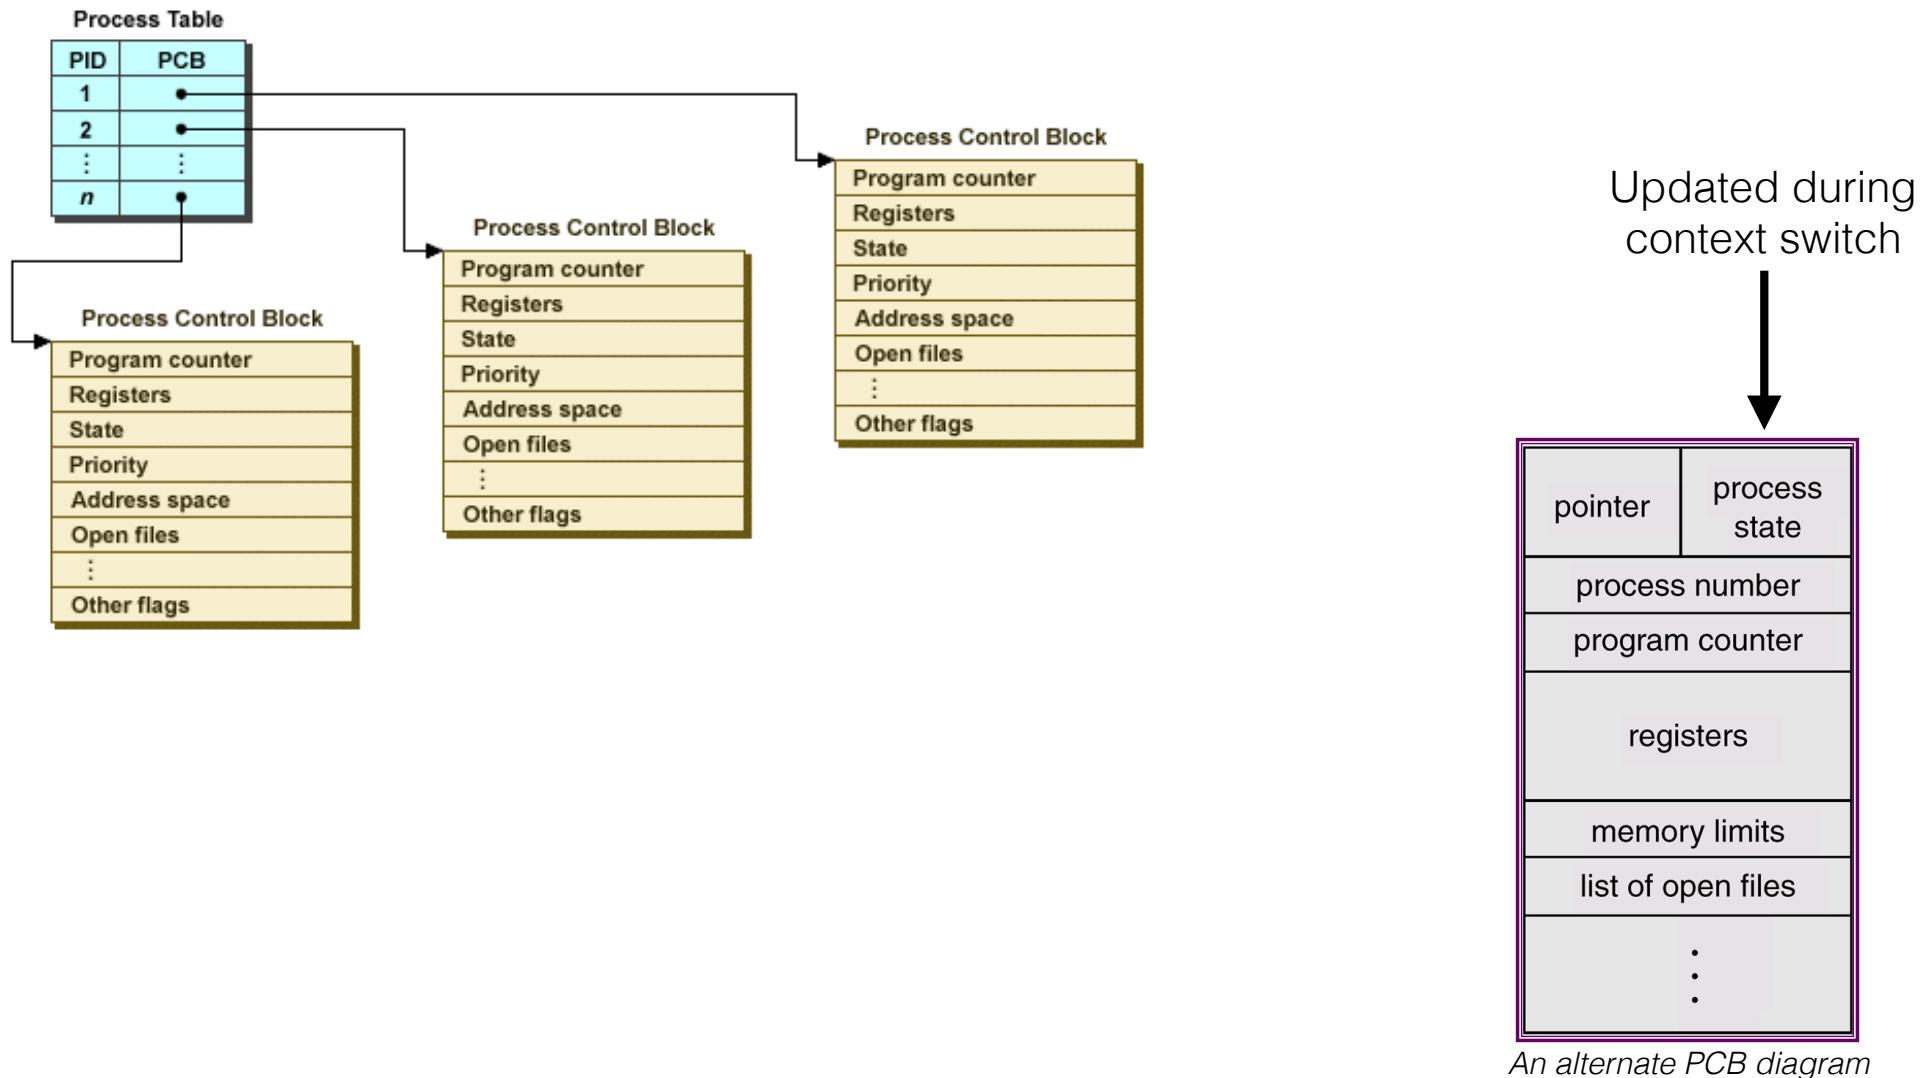


The Context Switch



Process Control Block

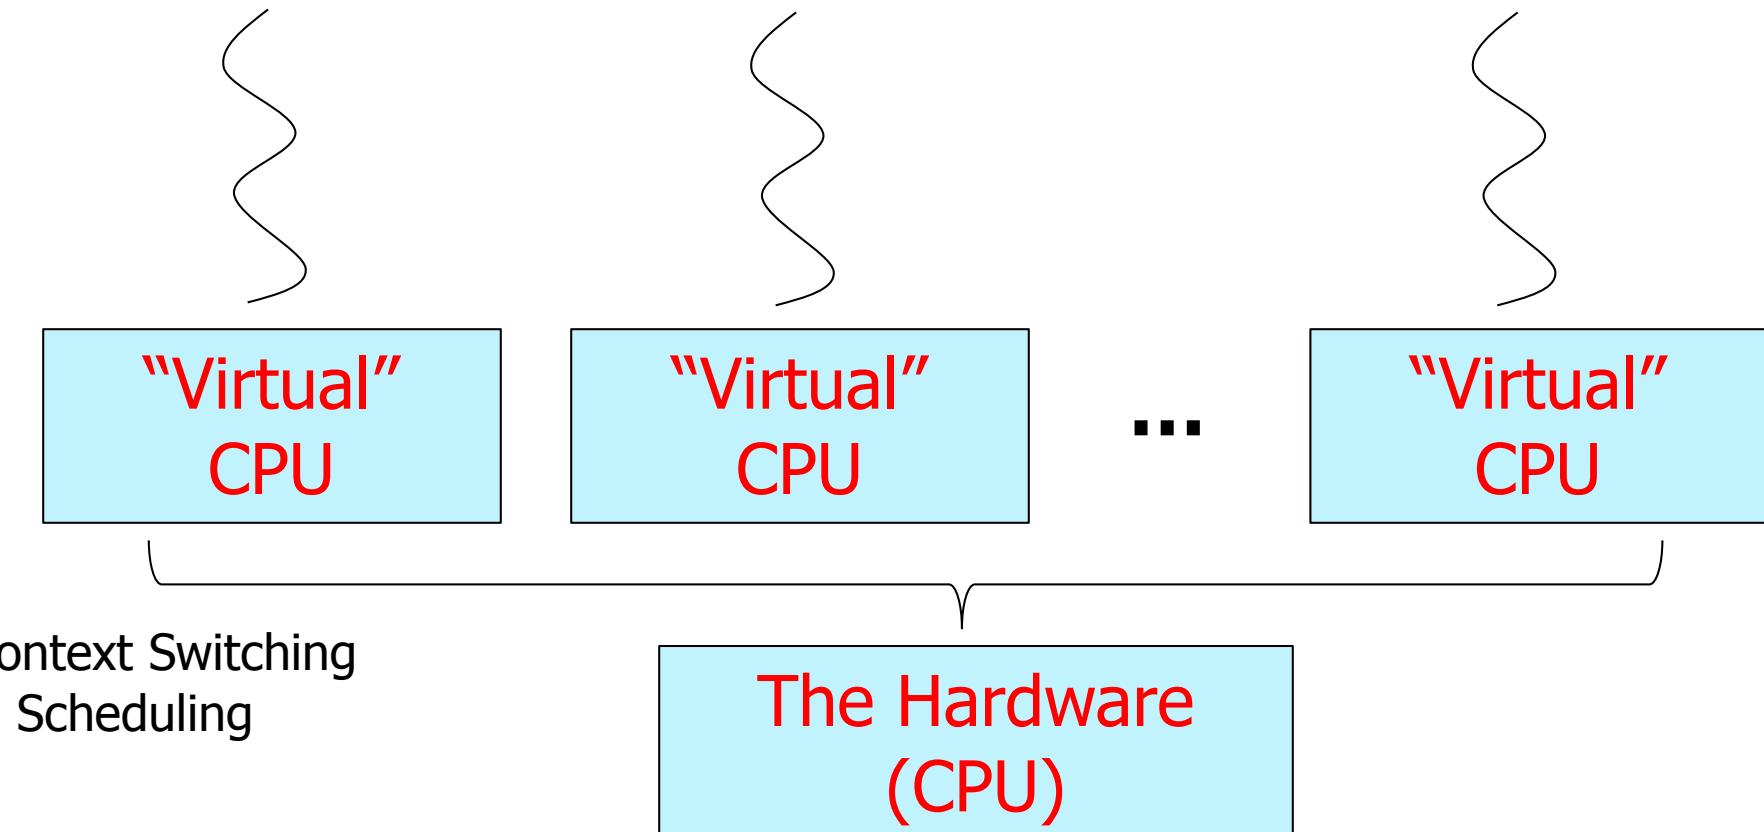
The state for processes that are not running on the CPU are maintained in the Process Control Block (PCB) data structure



Where We Are:

Last class, we discussed how context switches allow a single CPU to handle multiple tasks:

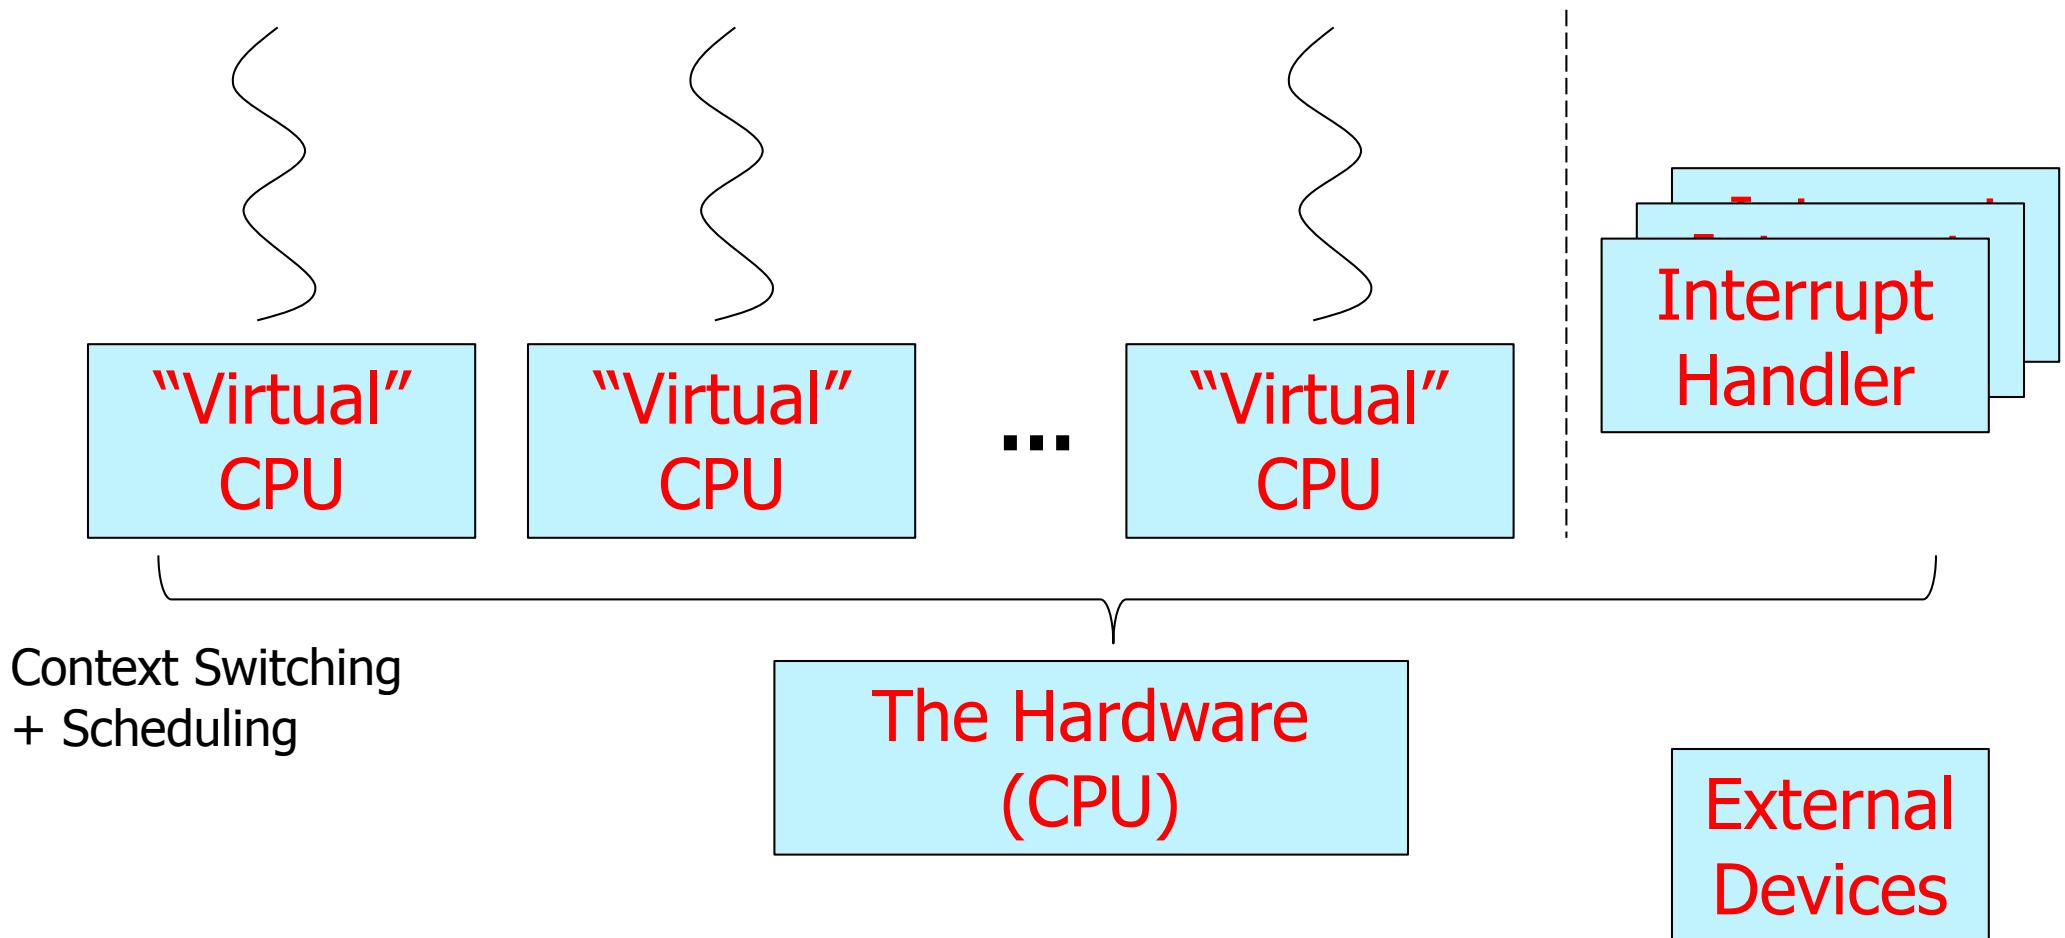
What's missing from this picture?



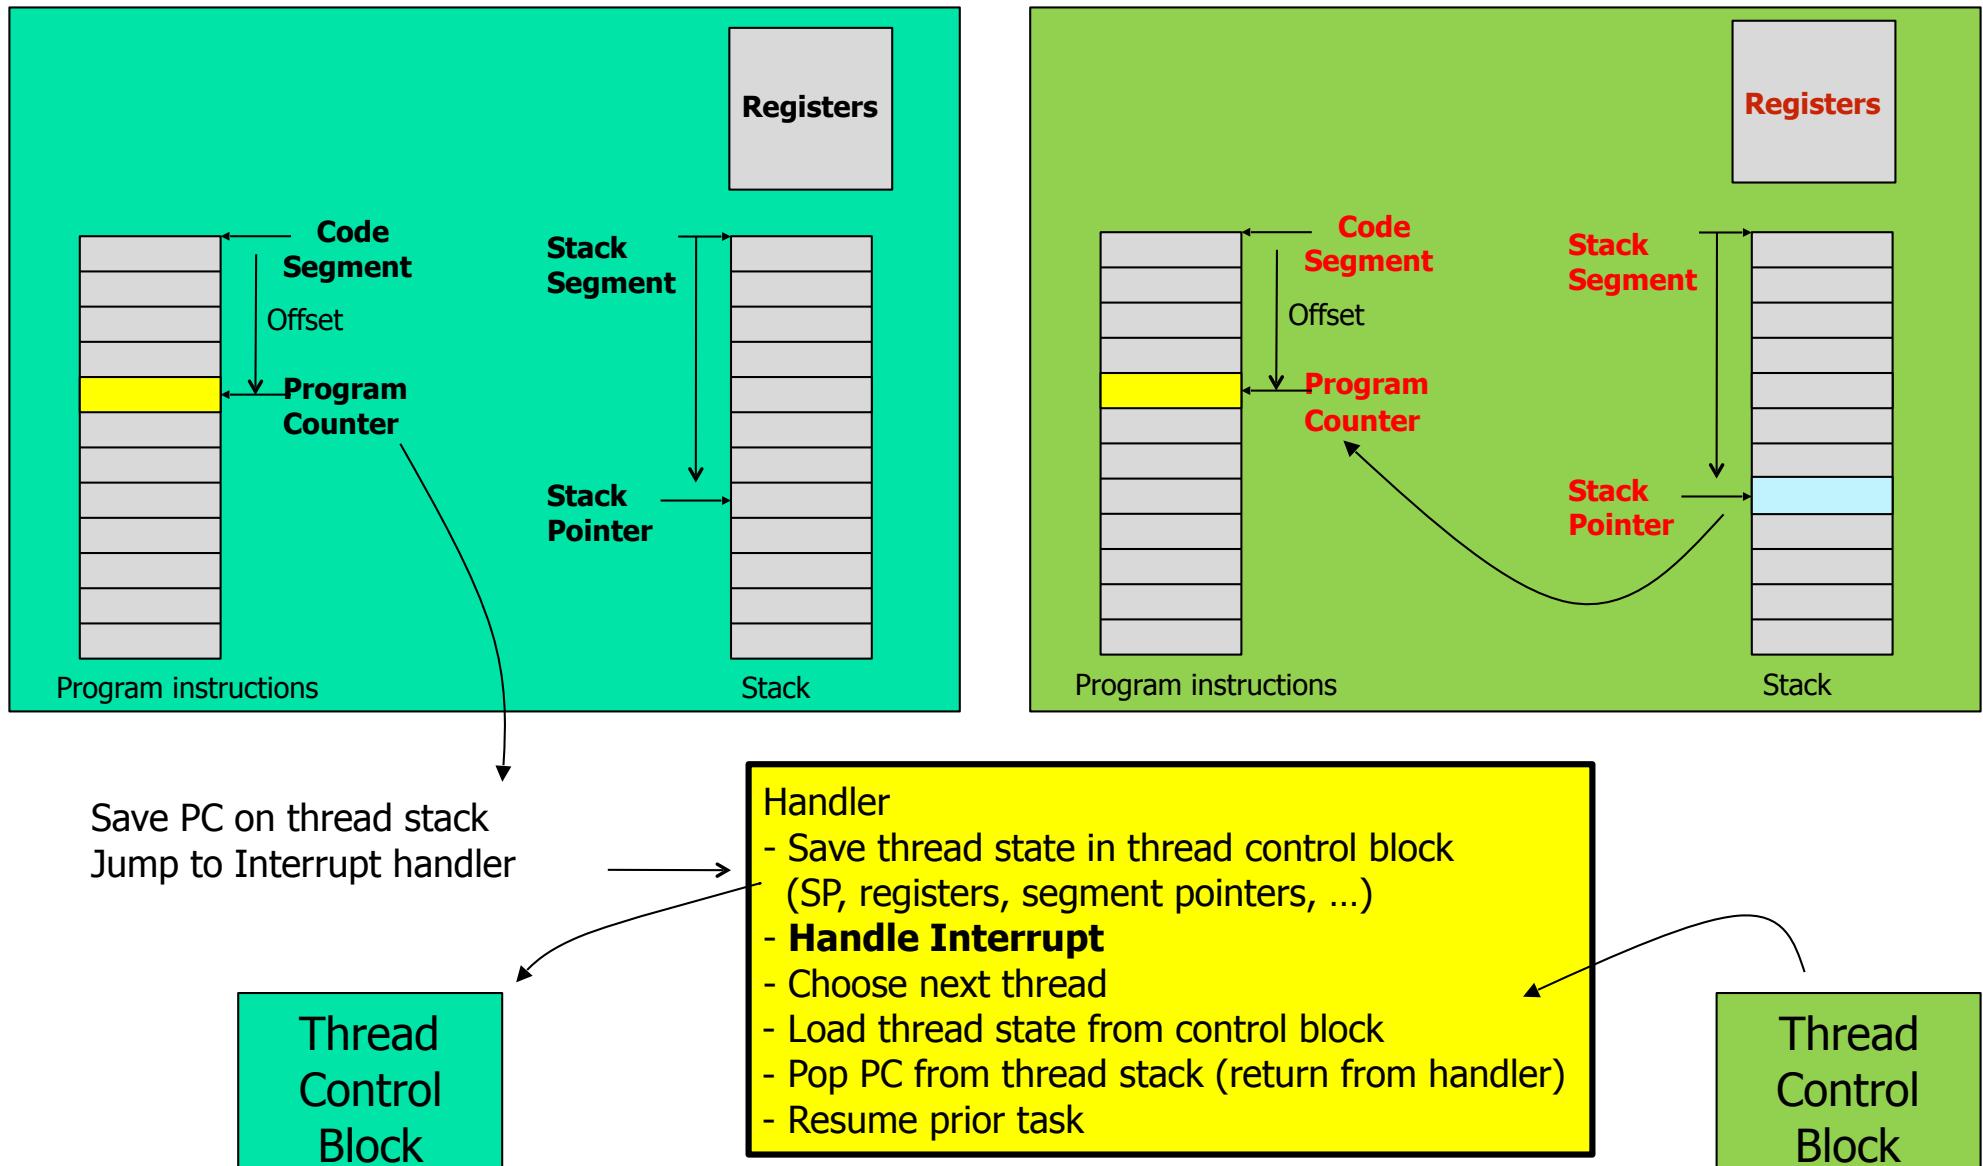
Where We Are:

Interrupts to drive scheduling decisions!

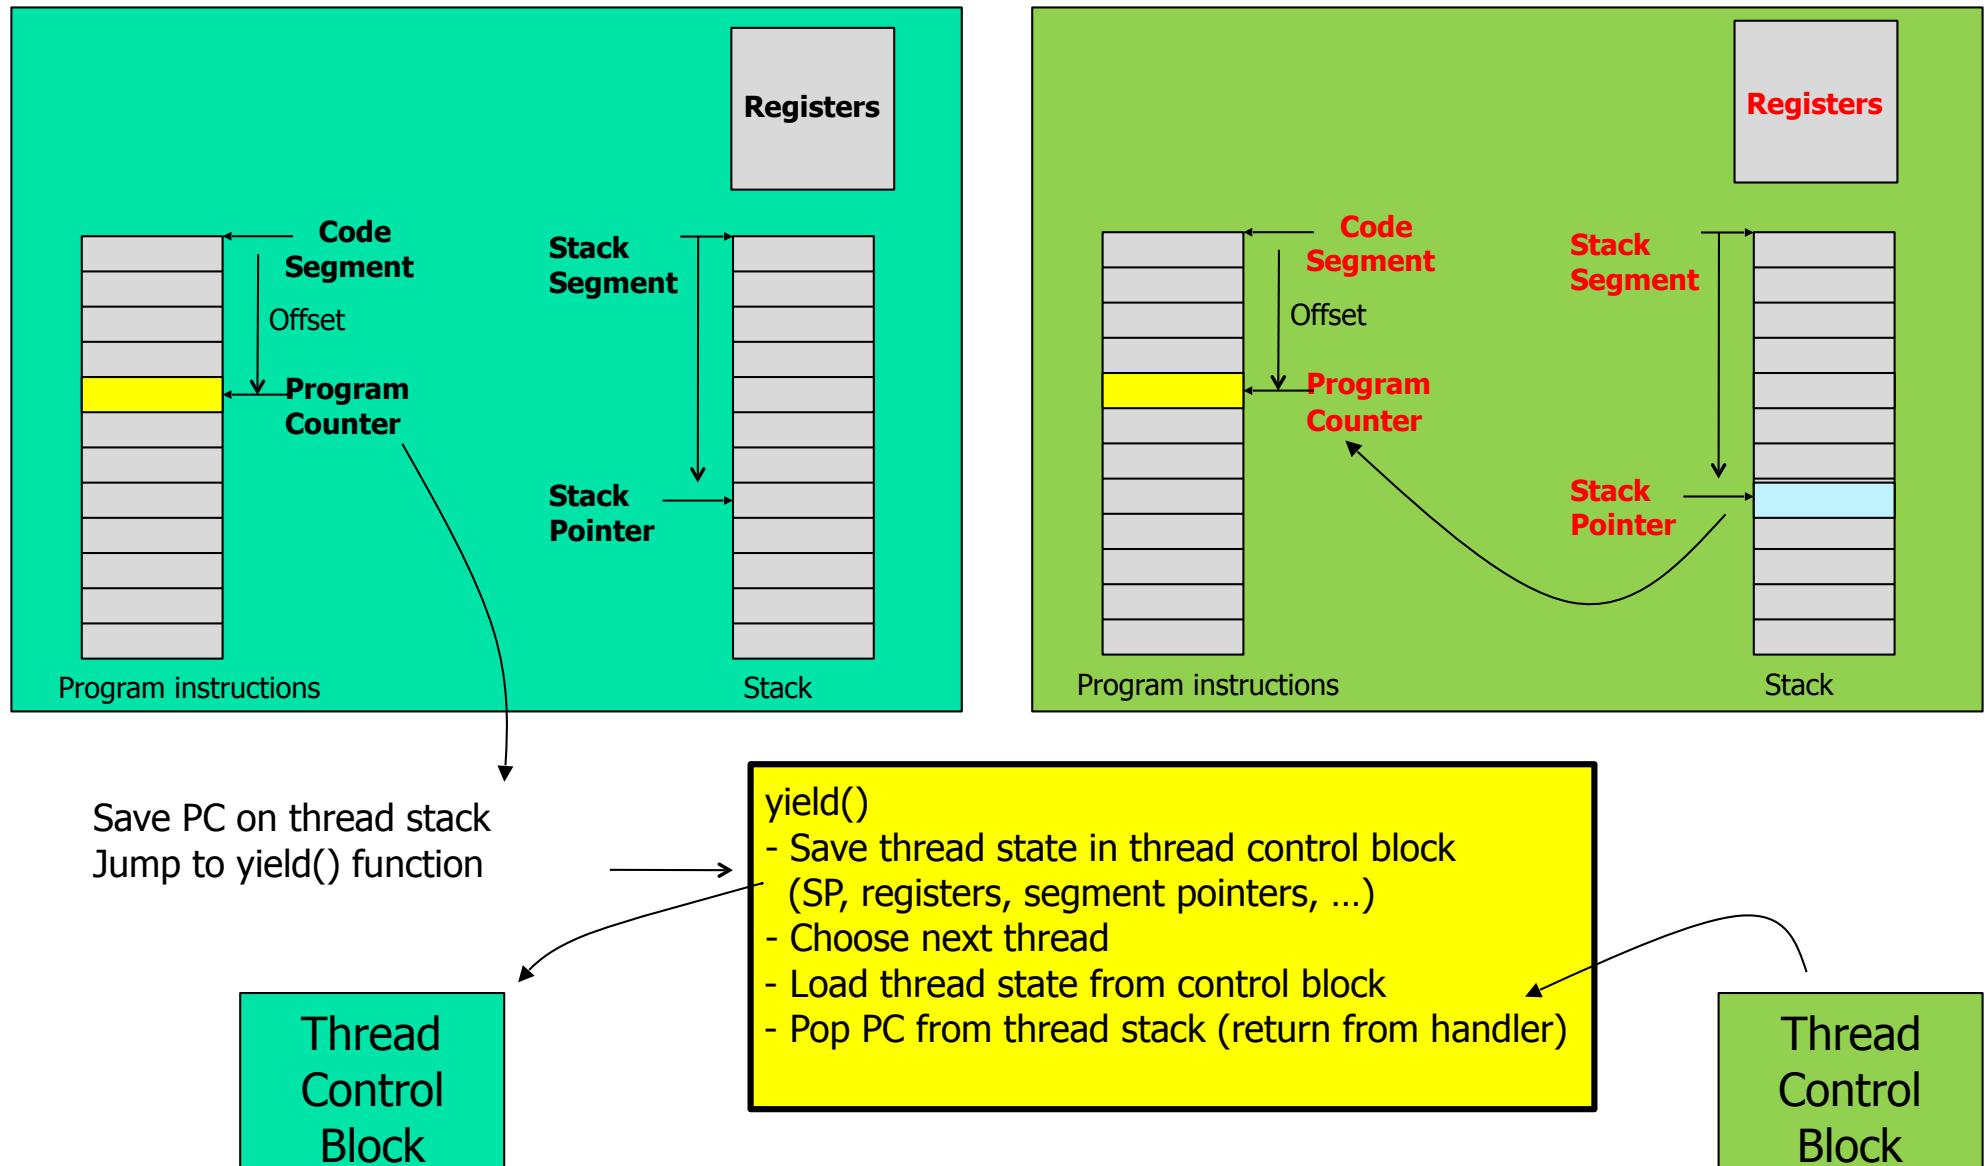
Interrupt handlers are also tasks that share the CPU.



CTX Switch: Interrupt



Can also CTX Switch from Yield



How do we take interrupts safely??

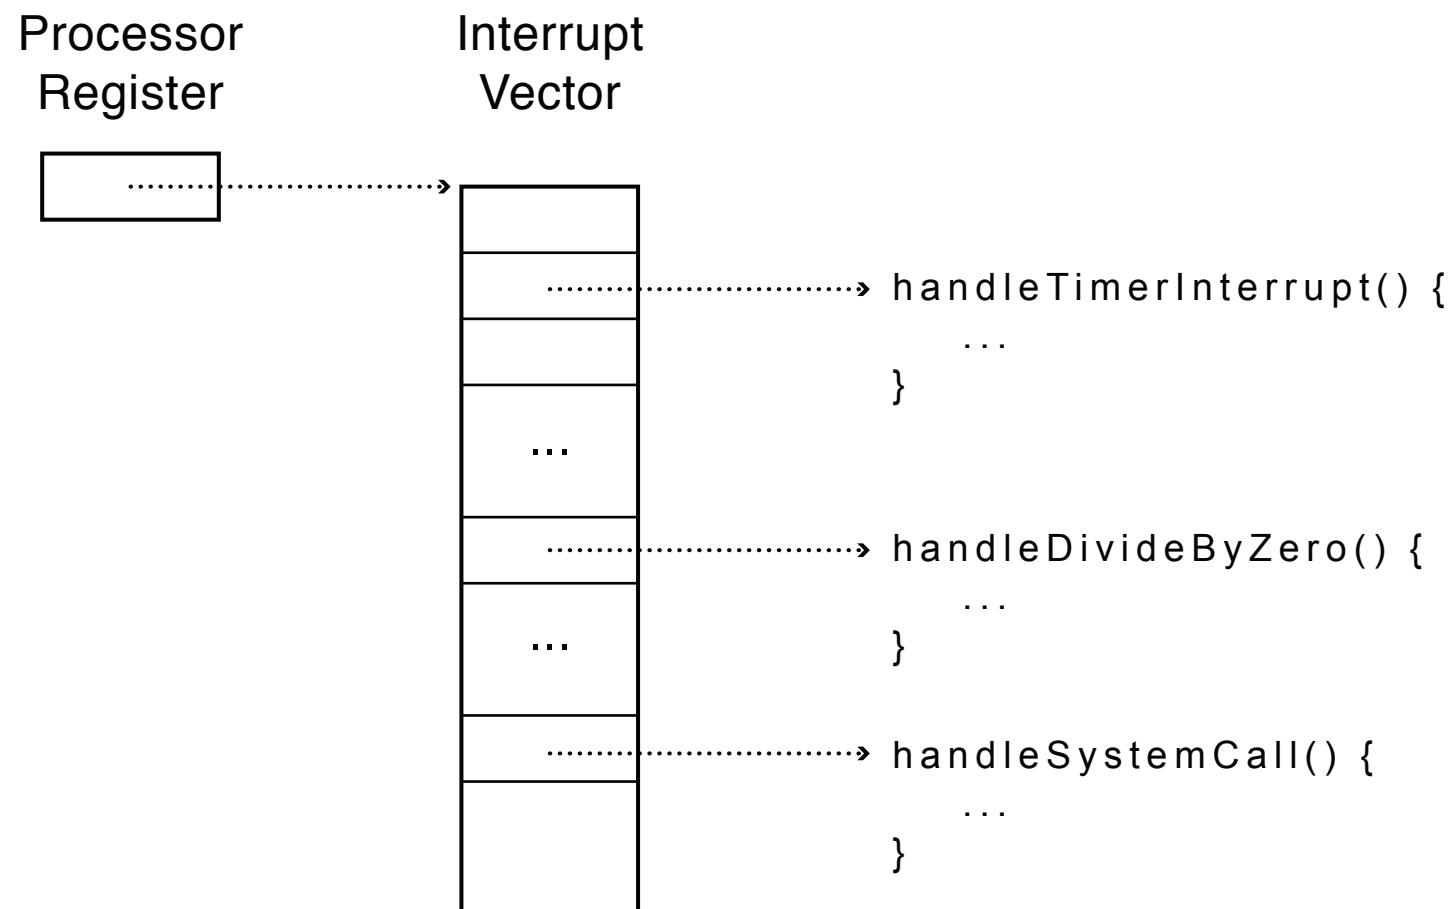


- **Interrupt vector**
 - Where the processor looks for a handler
 - Limited number of entry points into kernel
- **Atomic transfer of control**
 - Single instruction to change:
 - Program counter
 - Stack pointer
 - Memory protection
 - Kernel/user mode
- **Transparent restartable execution**
 - User program does not know interrupt occurred

Interrupt Vector Table



Table set up by OS kernel; pointers to code to run on different events

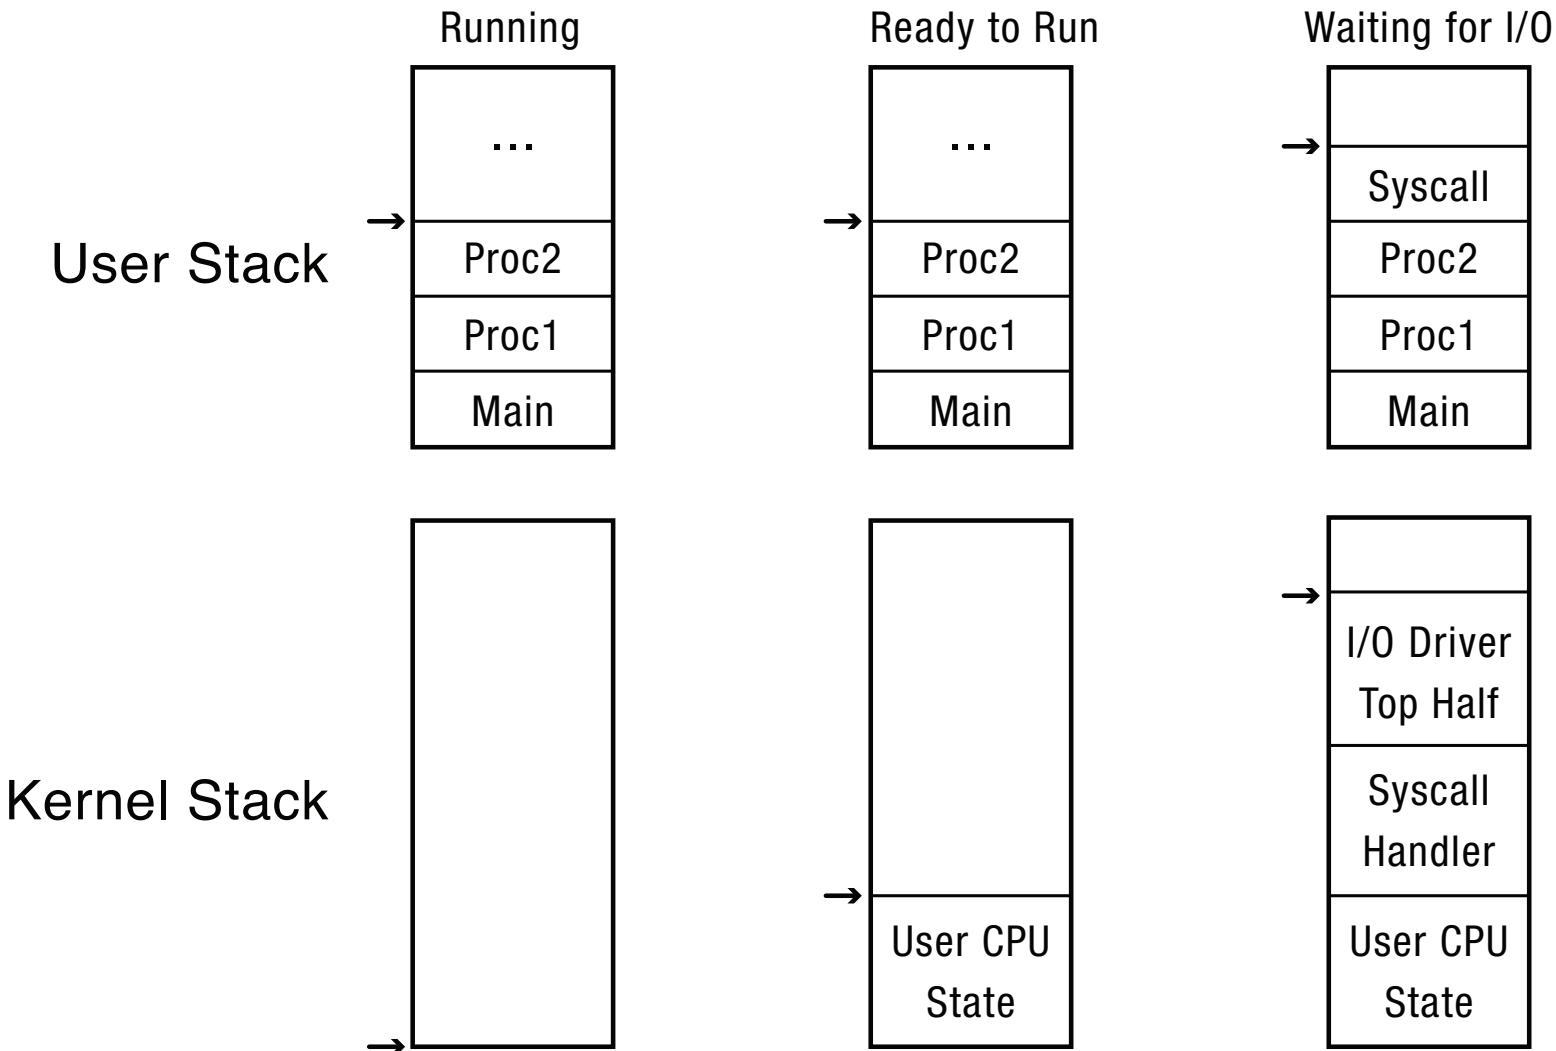


Interrupt Stack



- Per-processor, located in kernel (not user) memory
 - Fun fact! Usually a process/thread has both a kernel and user stack
- **Why can't the interrupt handler run on the stack of the interrupted user process?**

Interrupt Stack



Hardware Interrupts



- **Hardware generated:**
 - Different I/O devices are connected to different physical lines (pins) of an “Interrupt controller”
 - Device hardware signals the corresponding line
 - Interrupt controller signals the CPU (by signaling the Interrupt pin and passing an interrupt number)
 - CPU saves return address after next instruction and jumps to corresponding interrupt handler

Why Hardware INTs?



Why Hardware INTs?



- Hardware devices may need asynchronous and immediate service. For example:
 - Timer interrupt: Timers and time-dependent activities need to be updated with the passage of time at precise intervals

Why Hardware INTs?



- Hardware devices may need asynchronous and immediate service. For example:
 - Timer interrupt: Timers and time-dependent activities need to be updated with the passage of time at precise intervals
 - Network interrupt: The network card interrupts the CPU when data arrives from the network
 - I/O device interrupt: I/O devices (such as mouse and keyboard) issue hardware interrupts when they have input (e.g., a new character or mouse click)



Ex: Itanium 2 Pinout

	AH	AG	AF	AE	AD	AC	AB	AA	Y	W	V	U	T	R	P	N	M	L	K	J	H	G	F	E	D	C	B	A	
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	3	
4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	4	
5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	5	
6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	6	
7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	7	
8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	8	
9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	9	
10	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	10	
11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	11	
12	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	12	
13	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	13	
14	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	14	
15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15	
16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	16	
17	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	17	
18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	18	
19	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	19	
20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	20	
21	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	21	
22	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	22	
23	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	23	
24	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	24	
25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	25	
	AH	AG	AF	AE	AD	AC	AB	AA	Y	W	V	U	T	R	P	N	M	L	K	J	H	G	F	E	D	C	B	A	
	← Power Pad																												

UUU638b

Ex: Itanium 2 Pinout




	AH	AG	AF	AE	AD	AC	AB	AA	Y	W	V	U	R	P	N	M	L	K	J	H	G	F	E	D	C	B	A	
1	0	0	GND	0	GND	0	GND	0	GND	0	VC	0	TERM	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
2	0	0	0	0	0	0	0	0	0	0	VC	0	TERM	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2
3	0	0	TUNER(1)	TUNER(2)	TERM	GND	102#	GND	0	0	0	0	0	0	0	0	0	3										
4	0	GND	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	4
5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	5
6	0	0	GND	R2#	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	6
7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	7
8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	8
9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	9
10	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	10
11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	11
12	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	12
13	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	13
14	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	14
15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	15
16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	16
17	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	17
18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	18
19	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	19
20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	20
21	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	21
22	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	22
23	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	23
24	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	24
25	0	0	FERR	TH_TMR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	25
	AH	AG	AF	AE	AD	AC	AB	AA	Y	W	V	U	R	P	N	M	L	K	J	H	G	F	E	D	C	B	A	

Power Pad

UUU638b

Ex: Itanium 2 Pinout

12		GND	REQ#		REQ#		HTM#		A2#		A1#		D2#		VC
13	O	O	O	O	O	O	O	O	A2#	O	A1#	O	D2#		TERM
	BOLKN	BOLKP	GND	SSBY#	GND	RP#	GND	SSBY#	GND	A2#	GND	A1#	GND	D2#	
14		O	O		O	O		O	A3#		A3#	O	O		
		GND	TRDY#		O	REQ#		O	DEP#		A3#	VC	D2#		
15	O	O	O	O	O	O	O	O	A3#	O	A3#	O	D2#		TERM
	PWR	PROD	GND	LOCK#	GND	TND#	GND	BINT#	GND	A3#	GND	A2#	GND	D2#	
16		O	O		O	O		O	NC		A3#		A3#	O	VC
		GND	REQ#		O	REQ#		O			A3#		DEP#	1#	VC
17	O	O	O	O	O	O	O	O	O	O	O	O	O	O	TERM
	NC	NC	GND	NC	GND	NC	GND	A3#	GND	A3#	GND	BNR#	GND	D2#	
18		O	O		O	O		O	REQ#		A3#	O	O	O	
		GND	REQ#		O	NC		O	REQ#		A3#	A2#	VC	DEP#	1#
19	O	O	O	O	O	O	O	O	O	O	O	O	O	O	TERM
	NC	NC	GND	ERR#	GND	SSBY#	GND	SSBY#	GND	A3#	GND	A2#	GND	D2#	
20		O	O		O	O	O	O	AD#	GND	A3#	O	O	O	
		GND	PROD		O	RESET#		O	AD#	GND	A3#	A4#	GND	DEP#	1#
21	O	O	O	O	O	O	O	O	O	O	O	O	O	O	TERM
	NC	NC	GND	TRDY#	GND	NC	GND	DRDY#	GND	A4#	GND	A4#	GND	D12#	
22		O	O		O	O	O	O	A4#		A4#	O	O	O	
		GND	LINT0		O	BPV#		O	BPV#		A4#	A4#	VC	DEP#	1#
23	O	O	O	O	O	O	O	O	O	O	O	O	O	O	TERM
	A2#	IGN#		BPV#	GND	BPV#	GND	BPV#	GND	AP#	GND	A4#	GND	A4#	GND
24		O	O		O	O	O	O	O	O	O	O	O	O	
		GND	LINT1		O	BPV#	GND	BPV#	GND	A4#	O	A4#	GND	D12#	VC
25	O	O	O	O	O	O	O	O	O	O	O	O	O	O	TERM
	FERR#	THTRP#		PM#	GND	BPV#	GND	BPV#	GND	AP#	GND	A4#	GND	VC	GND
	AH	AG	AF	AE	AD	AC	AB	AA	Y	W	V	U	T	R	
															

LINTx — lines/pins for hardware interrupts.

In this case...

LINT0 — line for unmaskable interrupts

LINT1 — line for maskable interrupts

A Note on Multicore



A Note on Multicore



- How are interrupts handled on multicore machines?

A Note on Multicore

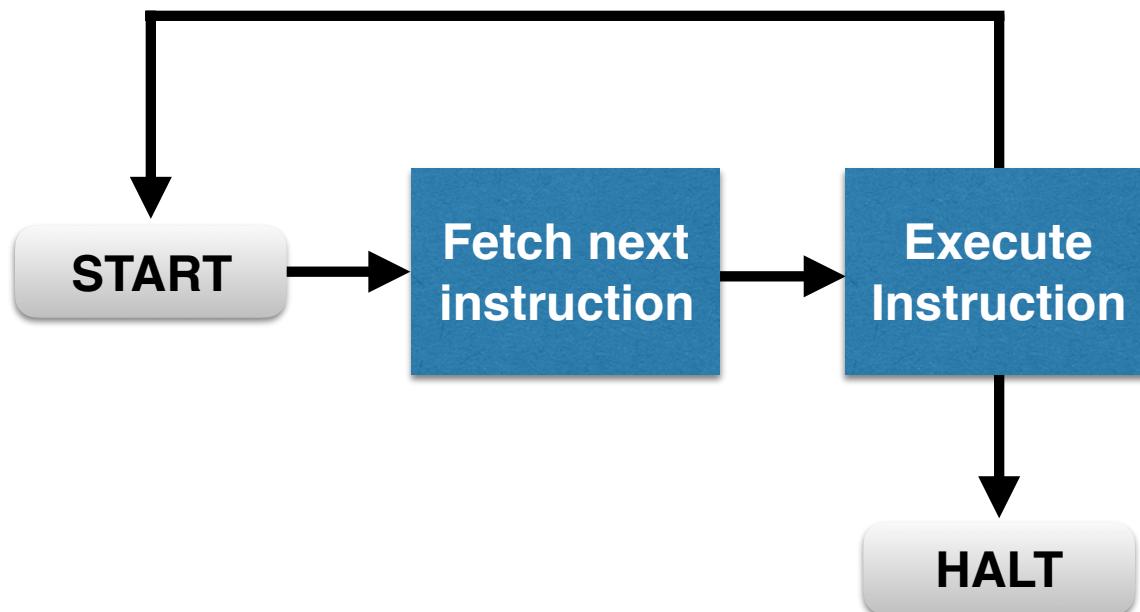


- How are interrupts handled on multicore machines?
 - On x86 systems each CPU gets its own local Advanced Programmable Interrupt Controller (APIC). They are wired in a way that allows routing device interrupts to any selected local APIC.
 - The OS can program the APICs to determine which interrupts get routed to which CPUs.
 - The default (unless OS states otherwise) is to route all interrupts to processor 0

Instruction Cycle



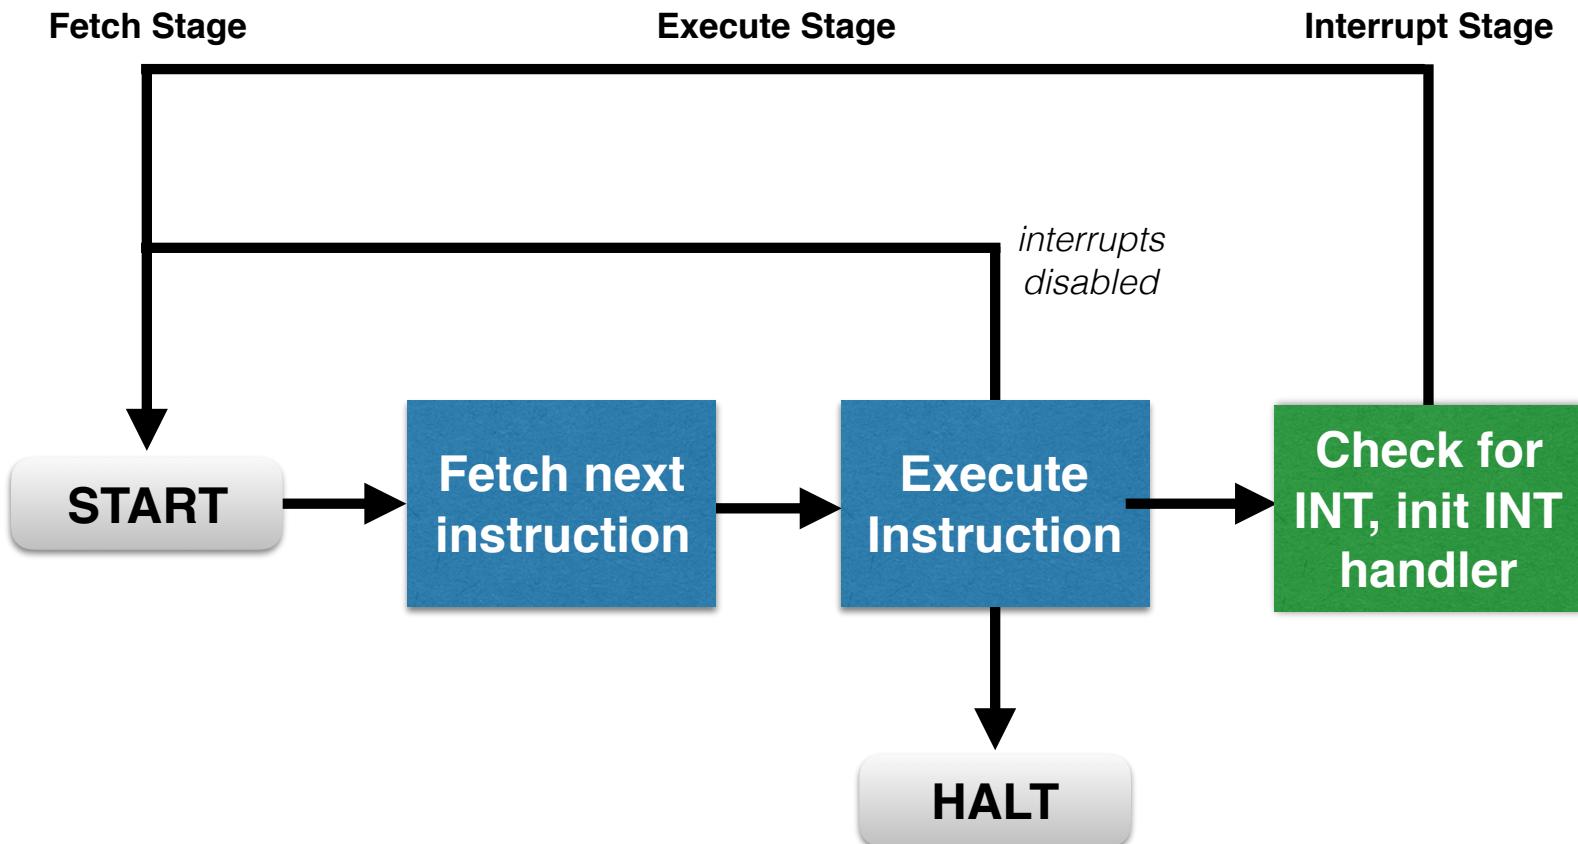
How does interrupt handling change the instruction cycle?



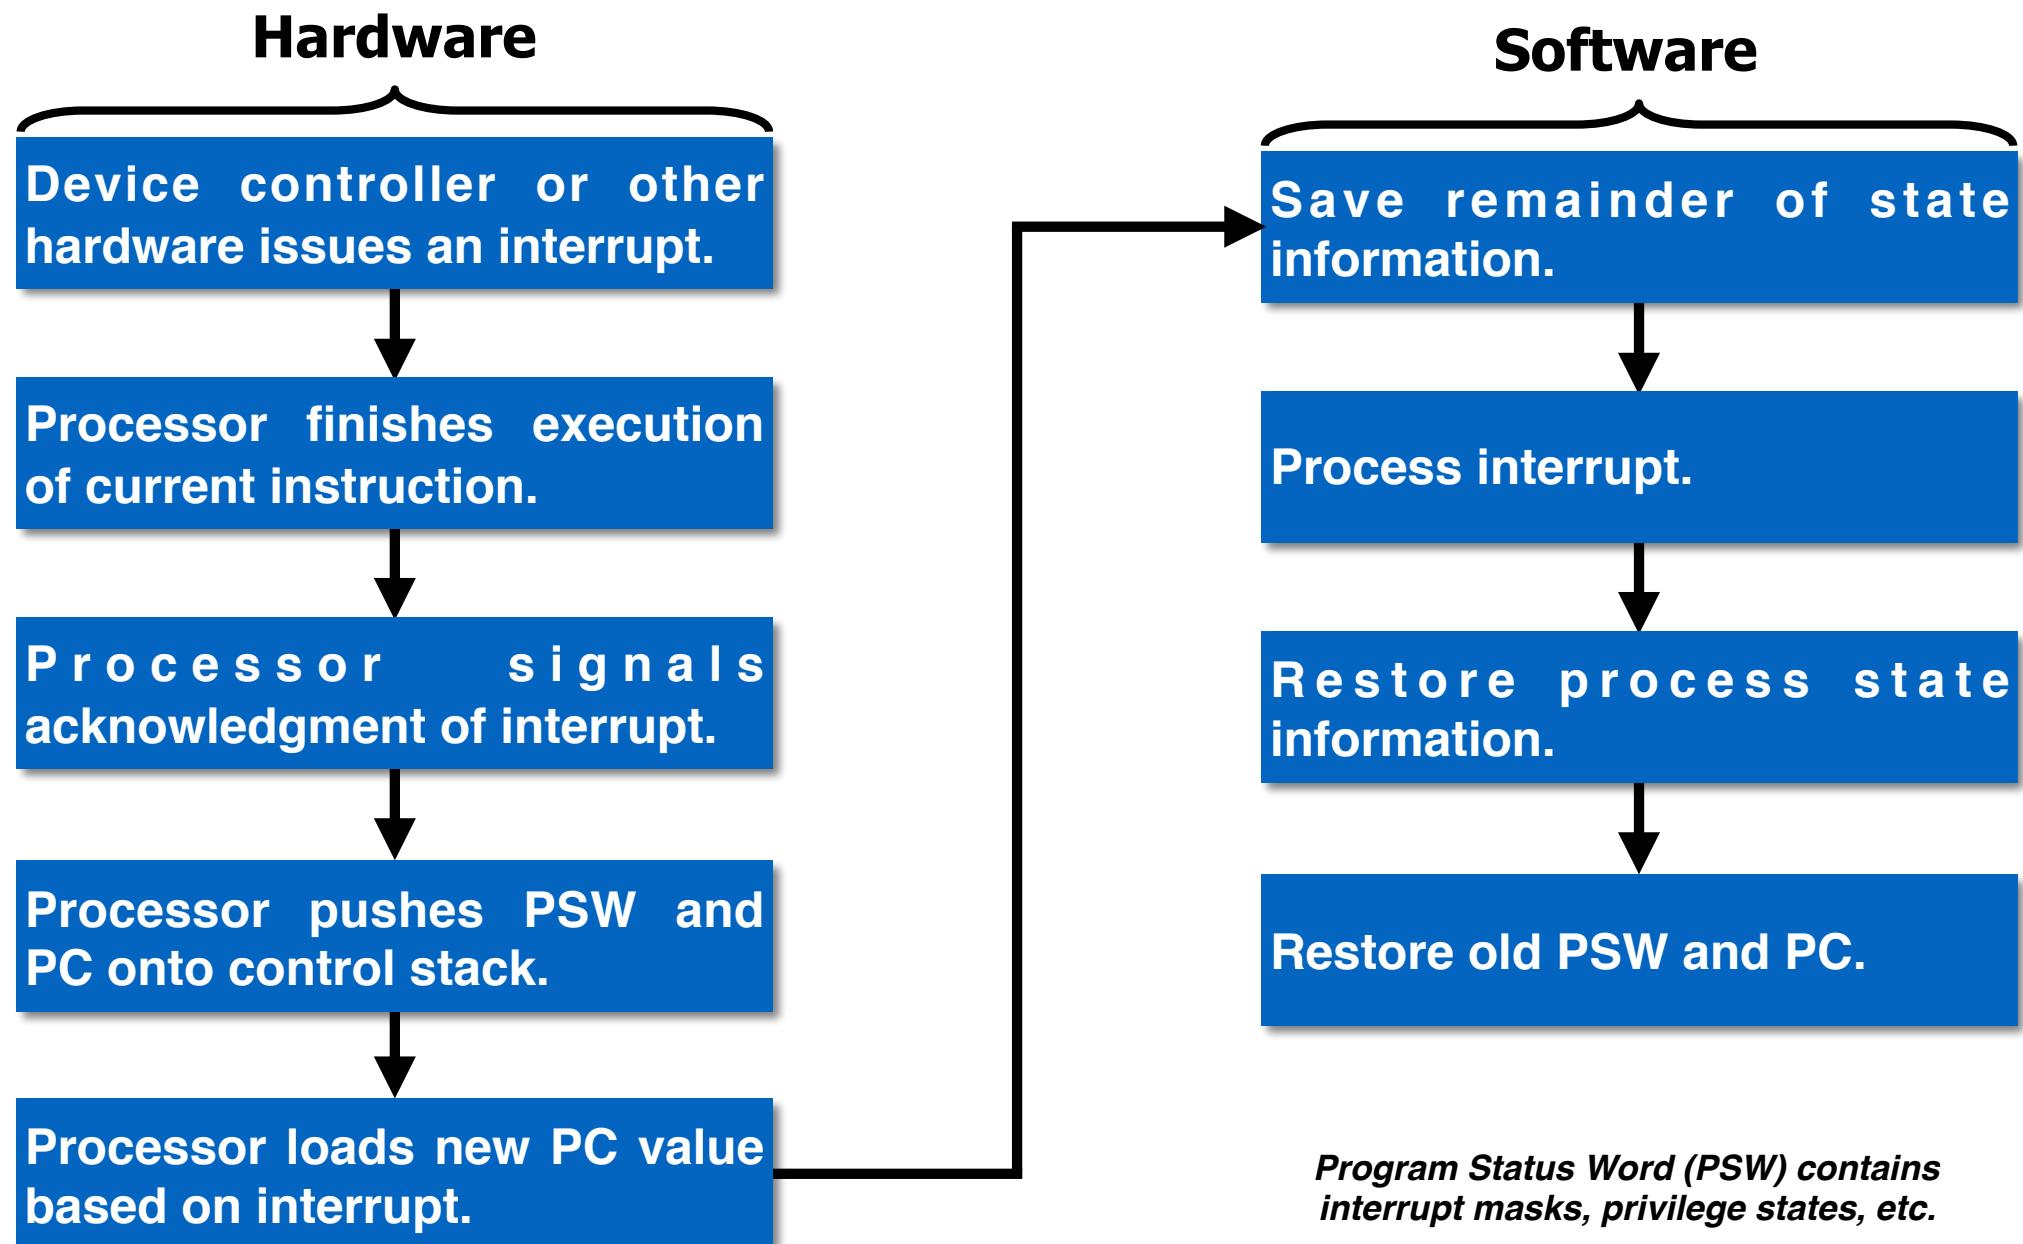
Instruction Cycle w/ INTs



How does interrupt handling change the instruction cycle?



Processing HW INT's



Other Interrupts



- Software Interrupts:
 - Interrupts caused by the execution of a software instruction:
 - `INT <interrupt_number>`
 - Used by the system call `interrupt()`
- Initiated by the running (user level) process
- Cause current processing to be interrupted and transfers control to the corresponding interrupt handler in the kernel

Other Interrupts



- **Exceptions**
 - Initiated by processor hardware itself
 - Example: divide by zero
- Like a software interrupt, they cause a transfer of control to the kernel to handle the exception

They're all interrupts



- HW -> CPU -> Kernel: Classic HW Interrupt
- User -> Kernel: SW Interrupt
- CPU -> Kernel: Exception
- Interrupt Handlers used in all 3 scenarios

INTs, Priorities, & Blocking



INTs, Priorities, & Blocking



- Interrupts (as the name suggests) have the highest priority (compared to user and kernel threads) and therefore run first
 - What are the implications on regular program execution?

INTs, Priorities, & Blocking



- Interrupts (as the name suggests) have the highest priority (compared to user and kernel threads) and therefore run first
 - What are the implications on regular program execution?
 - Must keep interrupt code short in order not to keep other processing stopped for a long time
 - Cannot block (regular processing does not resume until interrupt returns, so if the interrupt blocks in the middle the system “hangs”)

INTs, Priorities, & Blocking



INTs, Priorities, & Blocking



- Can an interrupt handler use malloc()?

INTs, Priorities, & Blocking



- Can an interrupt handler use malloc()?
- Can an interrupt handler write data to disk?

INTs, Priorities, & Blocking



- Can an interrupt handler use malloc()?
- Can an interrupt handler write data to disk?
- Can an interrupt handler use busy wait?
 - E.G. — `while (!event) loop;`

Interrupt Masking



- Interrupt handler runs with interrupts off
 - Re-enabled when interrupt completes
- OS kernel can also turn interrupts off
 - Eg., when determining the next process/thread to run

Interrupt Handlers



Designing an Interrupt Handler (Bottom Half):

- Since the interrupt handler must be minimal, all other processing related to the event that caused the interrupt must be deferred
 - Example:
 - Network interrupt causes packet to be copied from network card
 - Other processing on the packet should be deferred until its time comes
- The deferred portion of interrupt processing is called the “Bottom Half”



soft_irq's

- 32 handlers that must be statically defined in the Linux kernel.
- A hardware interrupt (before returning) uses `raise_softirq()` to mark that a given `soft_irq` must execute the bottom half
- At a later time, when scheduling permits, the marked `soft_irq` handler is executed
 - When a hardware interrupt is finished
 - When a process makes a system call
 - When a new process is scheduled



soft_irq types

- HI_SOFTIRQ
- TIMER_SOFTIRQ
- NET_TX_SOFTIRQ
- NET_RX_SOFTIRQ
- BLOCK_SOFTIRQ
- TASKLET_SOFTIRQ
- SCHED_SOFTIRQ
- ...



soft_irq types

- **HI_SOFTIRQ**
- **TIMER_SOFTIRQ**
- **NET_TX_SOFTIRQ**
- **NET_RX_SOFTIRQ**
- **BLOCK_SOFTIRQ**
- **TASKLET_SOFTIRQ**
- **SCHED_SOFTIRQ**
- ...



Tasklet

- Bottom halves multiplexed on top of soft_irq's
- Scheduled using
 - `tasklet_schedule()`
 - `tasklet_hi_schedule()`
- Same tasklet invocations are serialized
- Tasklets can be created or removed dynamically
- Cannot sleep (cannot save their context)

Work Queues



- Work deferred to its own thread
- Can be scheduled together with other threads according to priorities set by a scheduling policy
- Associated with its thread control block and hence can block (and save context)
 - `DECLARE_WORK(name, void (*func)(void *), void *data);`
 - `INIT_WORK(struct work_struct *work, void (*func)(void *), void *data);`
 - `schedule_work(&work);`