



System Calls, Signals, & Interrupts

CS 3410: Computer System Organization and Programming

Spring 2025



“Spring” Break!

- No lab this week!
- No work over break...except studying for Prelim 2
- Have fun, be safe!

Where are you going?



[PollEv.com/zjs](https://poll-ev.com/zjs)

Today's Goals

- Review: Processes
- System Calls: A C Perspective
- Exceptional Control Flow
 - Signals
 - Interrupts
- System Calls: How They Actually Work



Program




Storage/Disk



Memory



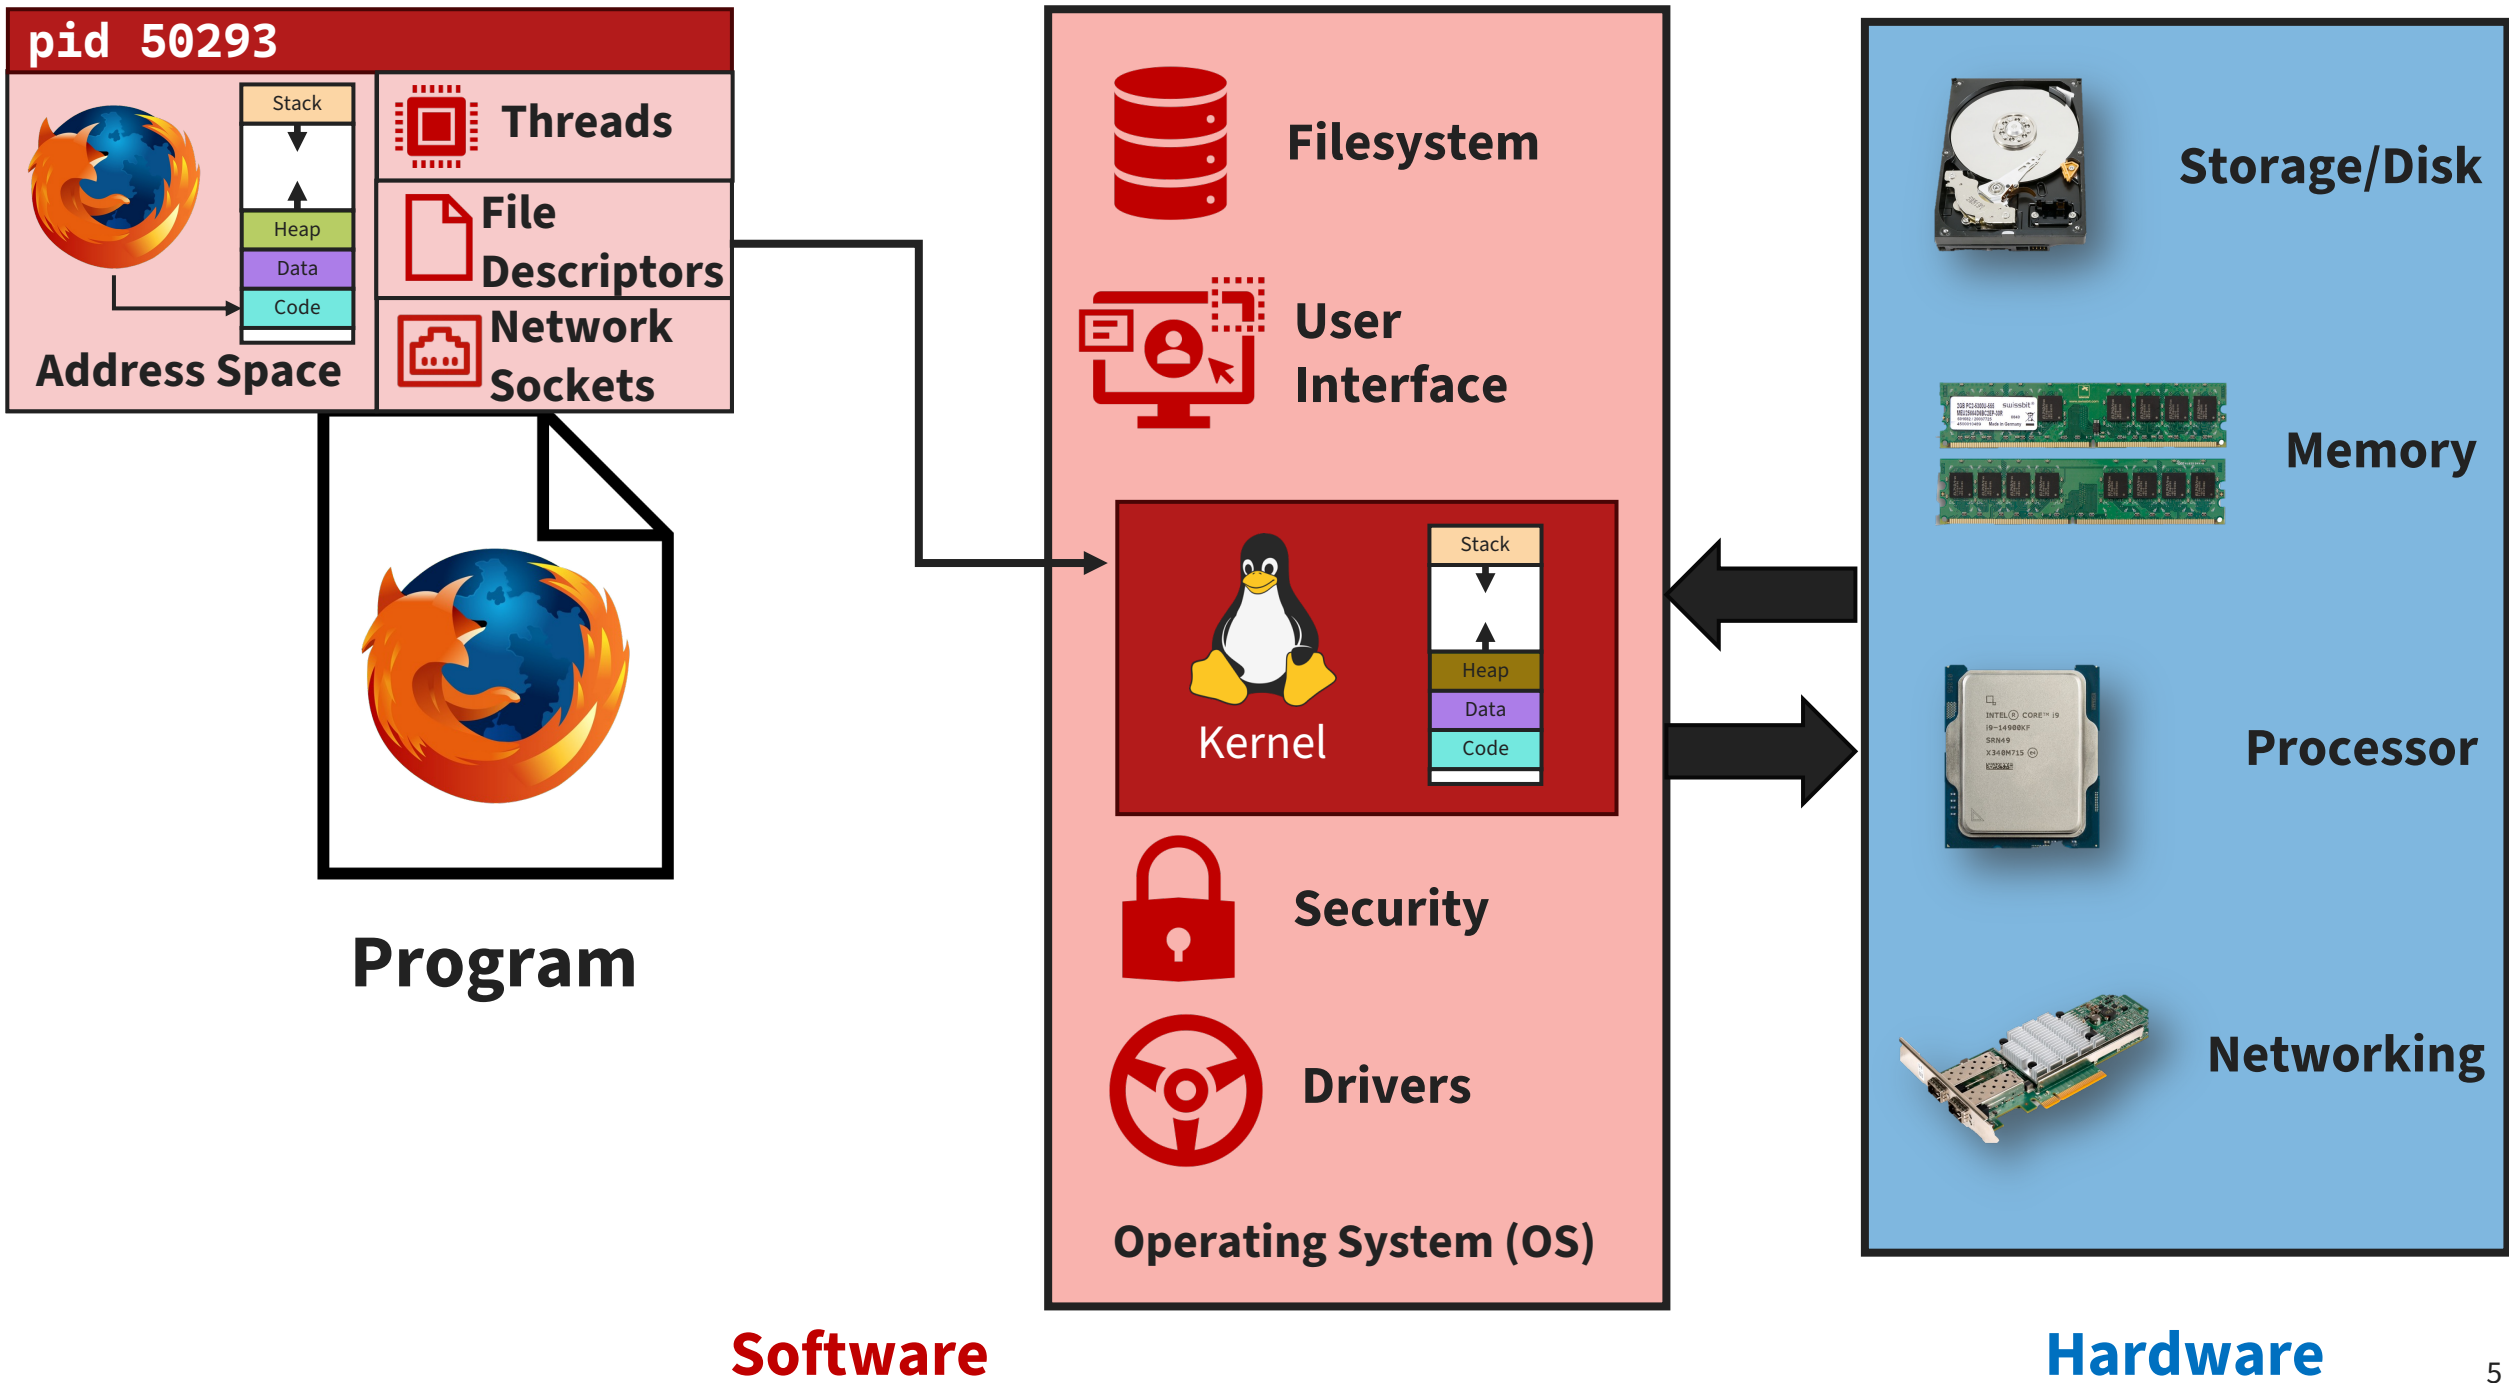
Processor



Networking






Software

Hardware



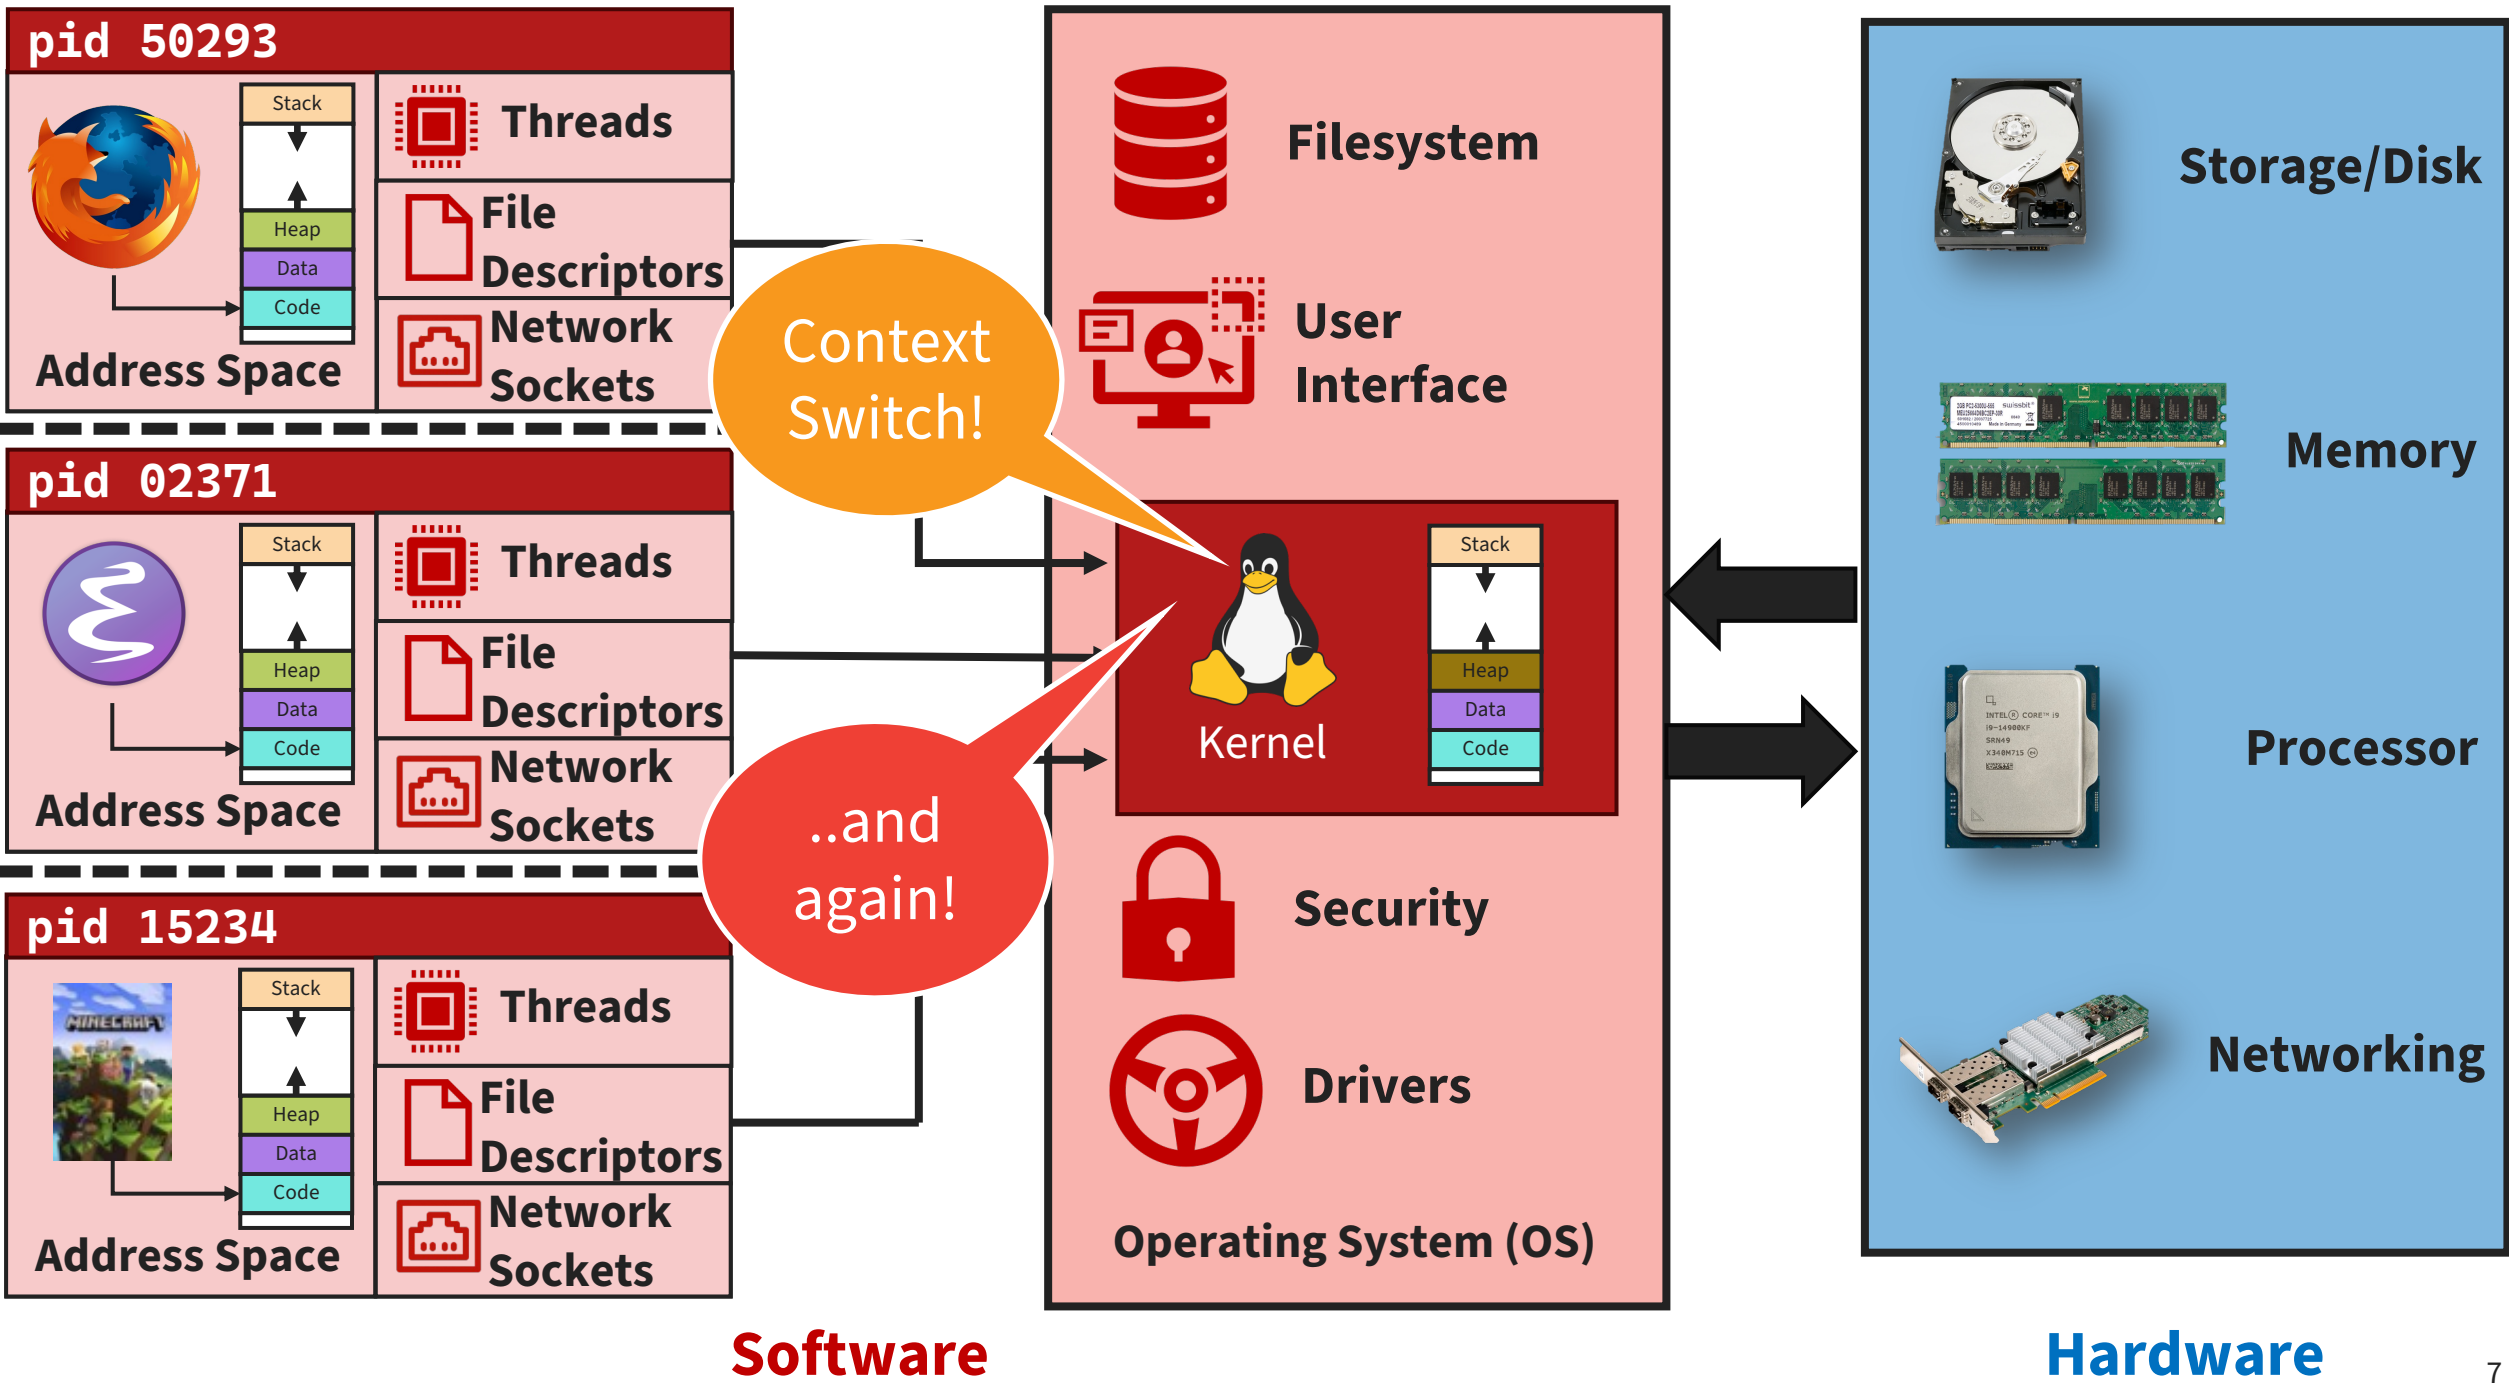
Software

Hardware



Software

Hardware



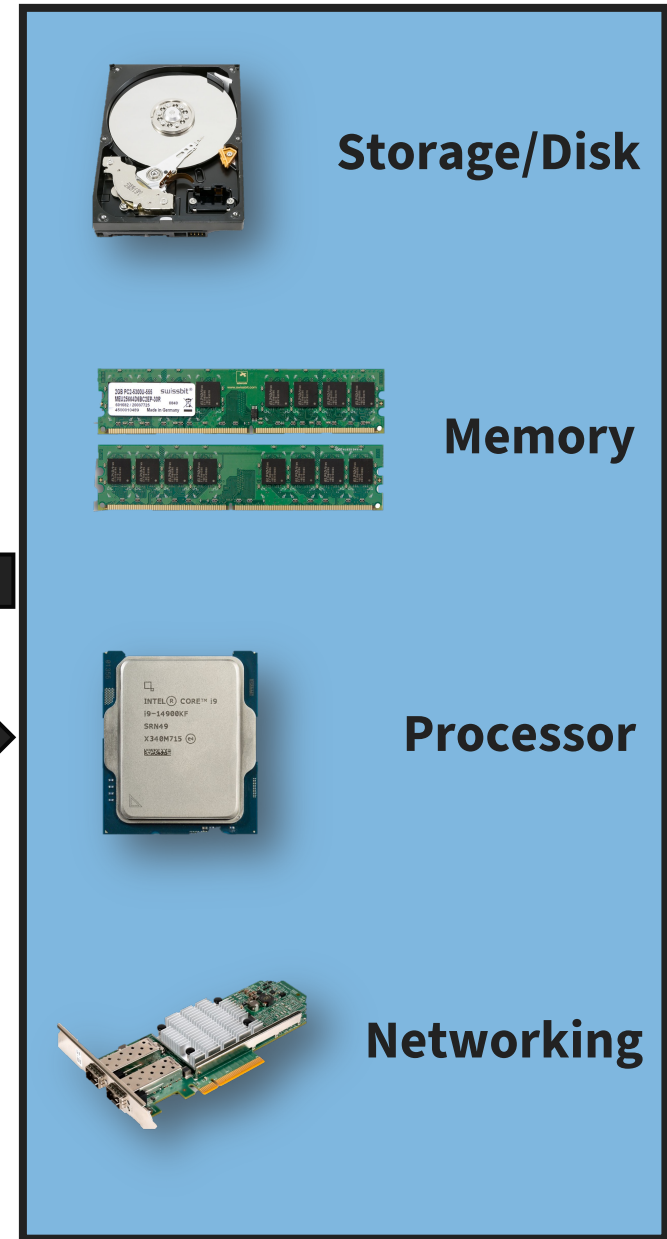

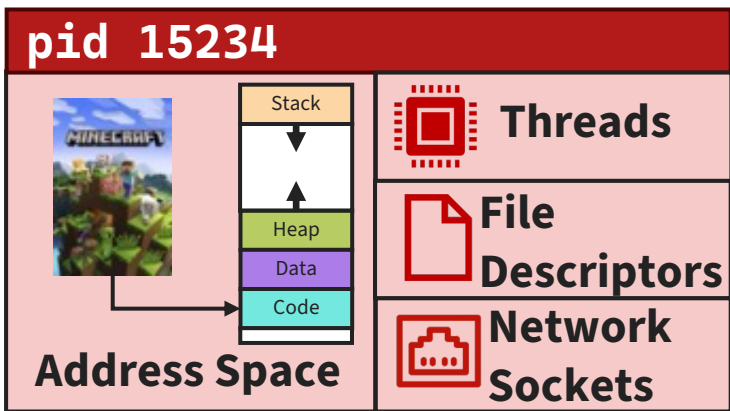


PollEverywhere

No matter what process is currently running on the CPU, the operating system is always actively running in the background.

- a) True
- b) False



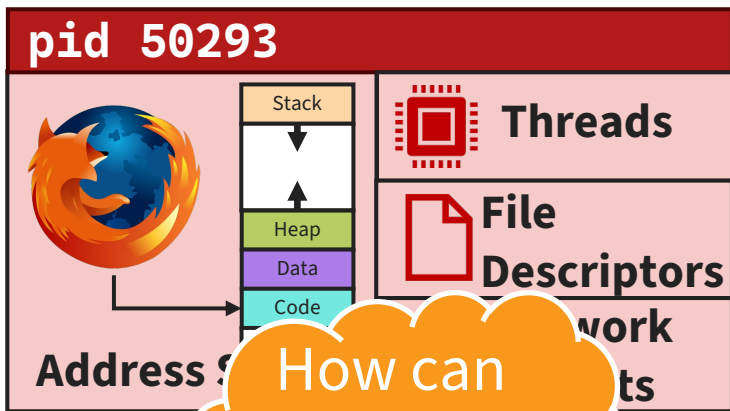


Pollev.com/zjs



Software


Hardware




Untrusted

Ask me to do it!


Make a system call!



Trusted

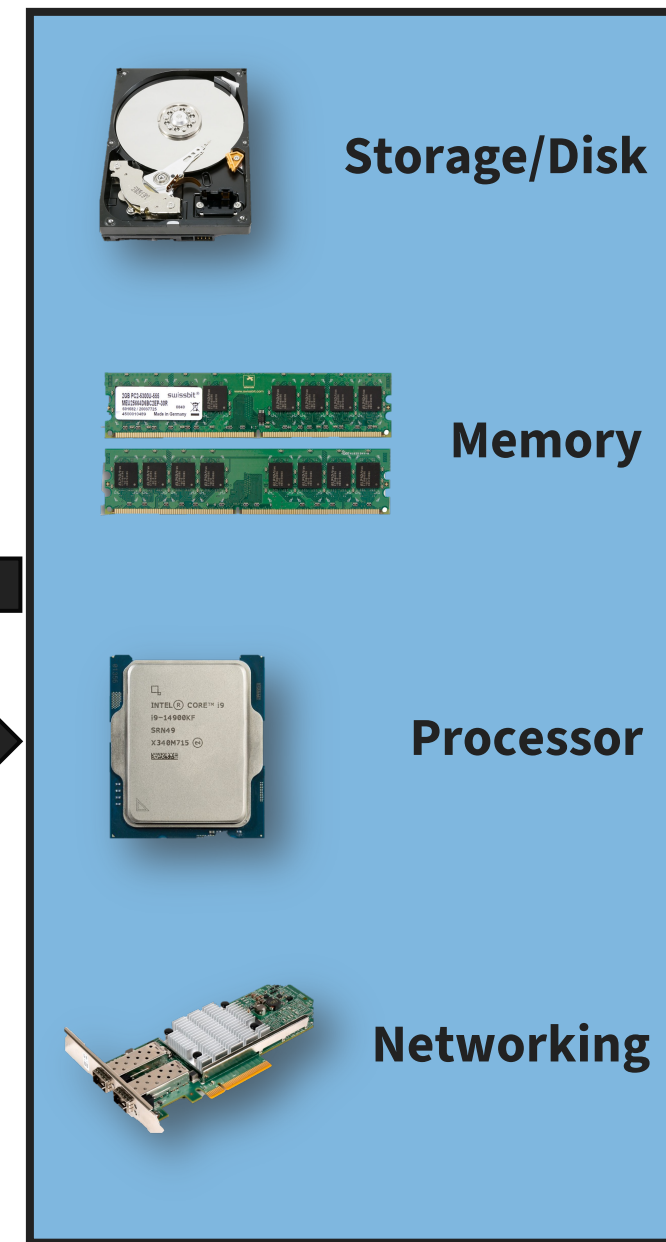



Hardware




`write()`

Asks the kernel to write an n -character long string to the file referenced by the provided file descriptor



Hardware

System Calls





System Calls

- A mechanism for processes to request the services from the kernel
 - File manipulation
 - Network communication
 - Process management
- Allows processes to perform **privileged operations** while running in user space
- OS defines **minimal** set of system calls
 - **Abstraction layer** between kernel and user code
- Why minimal?
 1. Portability: easier to implement and maintain
 2. Security: small “attack surface”; easier to protect against vulnerabilities



Portable Operating System Interface (POSIX)

- Originally published in 1988 under the name *IEEE Std 1003.1-1988*
 - Now 19 separate documents
- Colloquially referred to as Unix
- Implemented on multiple OSs
- `unistd.h`





Common System Calls

- **read()**: Reads data from a file descriptor
- **write()**: Writes data to a file descriptor
- **open()**: Opens a file and returns a file descriptor
- **close()**: Closes an open file descriptor
- **fork()**: Creates a new process
- **exec()**: Replaces the current process image with a new executable
- **waitpid()**: Waits for a specific child process to change state



Error Handling

- System calls often return -1 to indicate an error
- The global variable **errno** is set to indicate the specific error code
- The **perror()** function prints a human-readable error message based on the value of **errno**

```
1  #include <stdio.h>
2  #include <unistd.h>
3
4  const char msg[] = "CS 3410!";
5  int ret = write(STDOUT_FILENO, msg, sizeof(msg)-1);
6  if (ret == -1) {
7      perror("write");
8  }
```




Demo: A Tale of Three Syscalls

... named `fork()`, `exec()`, and `waitpid()`





How Are Processes Created?

- **fork()** :
 - Allocates process ID **pid**
 - Creates and initializes PCB
 - Creates and initializes new address space
 - Informs scheduler a new process is ready
- **exec(program, arguments)** :
 - Loads program into the address space
 - Copies arguments into memory address space
 - Initializes hardware context to start execution at “start”





How Do Processes Terminate?

- **exit()**: used by a process to terminate itself
- **abort()**: used by a parent process to terminate a child process
- **wait()** and **waitpid()**: used by a parent process to wait for a child process to terminate and retrieve its exit status



PollEverywhere

The operating system can interrupt user code using system calls.

- a) True
- b) False



Pollev.com/zjs



Signals



Signals

Signals are the mechanism for the kernel (or another process) to communicate with processes.

Common Signals:

- **SIGINT**: the **interrupt** signal
- **SIGTERM**: the (graceful) termination signal requests a process to terminate
- **SIGKILL**: the kill signal *forces* a process to terminate immediately
 - **Cannot be caught or ignored**
- **SIGSEGV**: the segmentation fault
- **SIGCHLD**: sent to parent process when a child process terminates or stops



Sending Signals

`kill(pid_t pid, int sig)`

`kill()` sends the signal **sig** to the process with the process ID **pid**

`raise(int sig)`

`raise()` sends the signal **sig** to the calling process

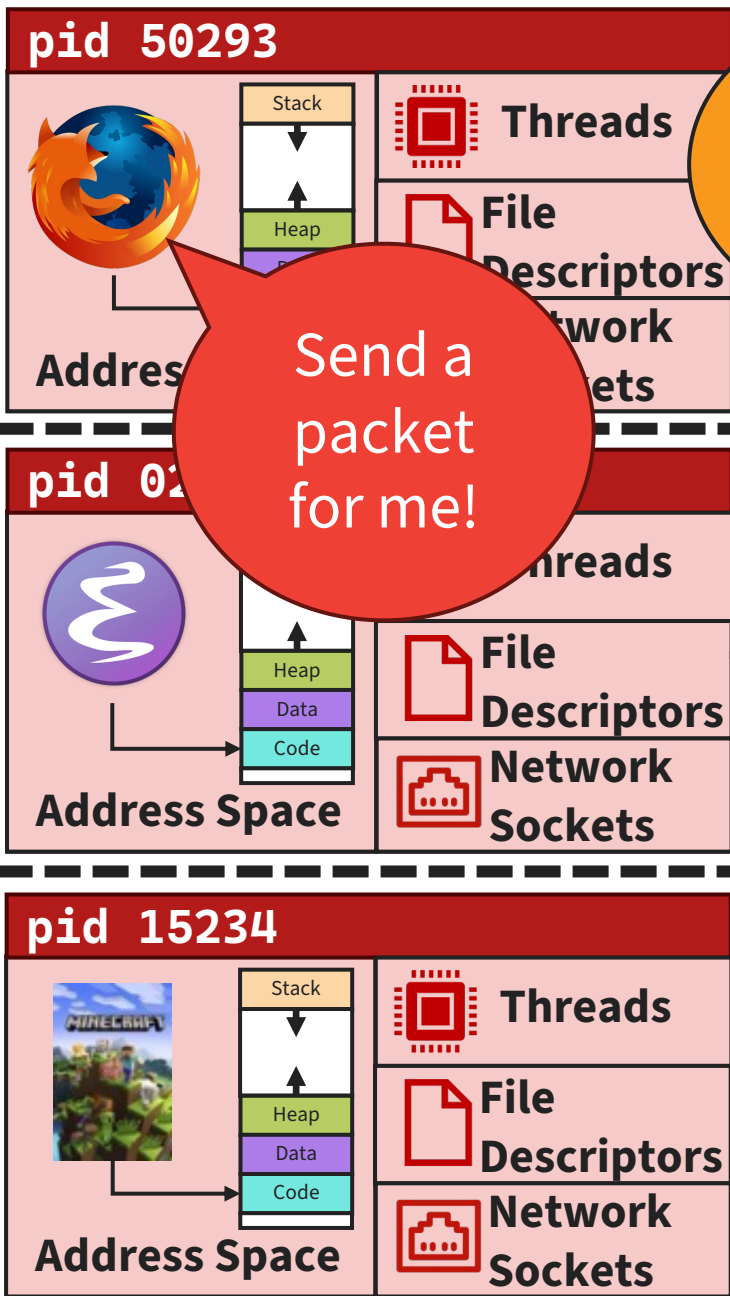
Handling Signals

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <signal.h>
4  #include <unistd.h>
5
6  void handle_signal(int sig) {
7      printf("Caught signal %d\n", sig);
8      exit(1);
9  }
10
11 int main() {
12     signal(SIGINT, handle_signal); // Set up the signal handler for SIGINT.
13     while (1) {
14         printf("Running. Press Ctrl+C to stop.\n");
15         sleep(1);
16     }
17     return 0;
18 }
```


Exceptional Control Flow

Traps, Faults, and Interrupts







Send a packet for me!

The user pressed CTRL-C!!

I'm out of free space!

I was asked to divide by zero... ☹️



Software

Hardware



Cornell Bowers C/S
Computer Science

Types of Interrupts

An **interrupt** is an unscheduled event that needs immediate attention which disrupts the normal execution of a program.

Hardware Interrupts:

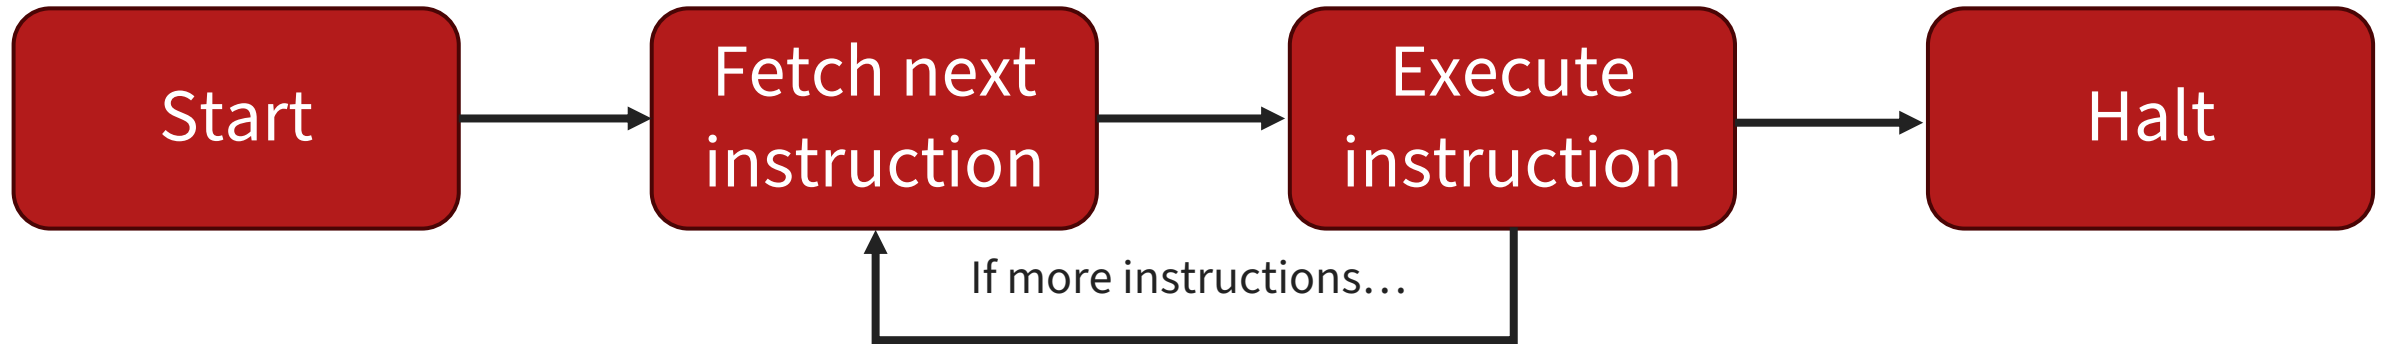
- Generated by hardware devices
 - Ex: keyboard input, hardware failure
- Some can be ignored (maskable); some can't (unmaskable)

Software Exceptions:

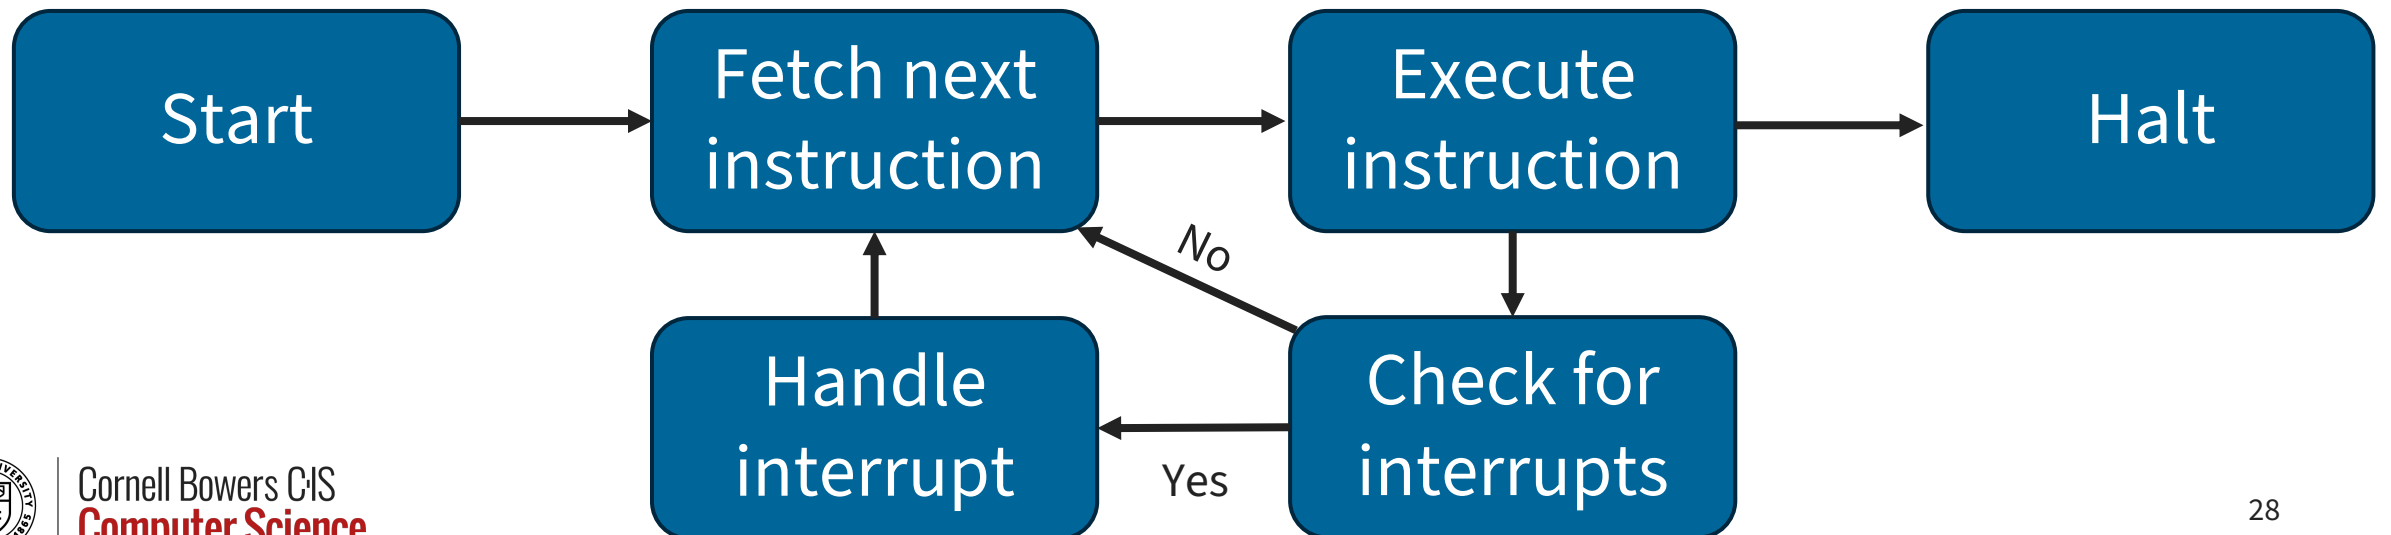
- Generated by programs (i.e., instructions!)
- Intentional interrupts are **traps** (e.g., system calls)
- Unintentional interrupts are **faults** (e.g., divide by zero, segmentation fault)

Instruction Cycle & Interrupts

Without Interrupts



With Interrupts



How Interrupts Work

- **Interrupt Signal**

- An interrupt signal is sent to the CPU by a hardware device or software

- **Saving State**

- The CPU saves the current state of the running process (e.g., program counter, registers) so it can resume execution later

- **Interrupt Handling**

- The CPU transfers control to the interrupt handler associated with the interrupt. The interrupt handler processes the event (e.g., reading data from a device)

- **Restoring State**

- The CPU restores the saved state and resumes execution of the interrupted process



What are Interrupts For?



EFFICIENCY



RESPONSIVENESS



MULTITASKING

How System Calls Actually Work

Let's write “Hello, World!” without the standard library!





Executing a System Call in RISC-V

Process (User Mode)

1. Call systems call function in standard library
2. Place function args. in **a0–a5** and/or **user stack**
3. Place syscall number in **a7**
4. Call **ecall**

...

Hardware

5. Save registers to **kernel stack** and current caller PC in special register **sepc**
6. Switch to **kernel mode**
7. Lookup syscall handler address in trap table
8. Jump to syscall handler
12. Restore registers from **kernel stack**
13. Switch to **user mode**
14. Jump to instr. after **ecall**

Kernel (Kernel Mode)

9. Do work of syscall
10. Place result in **a0**
11. Execute “supervisor exception return” **sret**



Unprivileged vs. Privileged Specifications

RISC-V Technical Specifications



Owned by [Jeff Scheel](#) ...
Last updated: Feb 21, 2025 • 5 min read

Below is a comprehensive list of all ratified technical publications.

[\[ISA Specifications \]](#) [\[Profiles \]](#) [\[Non-ISA Specifications \]](#) [\[Compatibility Test Framework \]](#)

ISA Specifications

These are the current, published versions of the ISA specifications. Prior published versions and the original ratification specifications for included extensions can be found on the [RISC-V Technical Specifications Archive](#) page.

Specification name (PDF link)	Version	Published	RISC-V Community	Source Repository
The RISC-V Instruction Set Manual Volume I: Unprivileged ISA	20240411	May 2024	Unprivileged Horizontal Committee	riscv/riscv-isa-manual
The RISC-V Instruction Set Manual Volume II: Privileged Architecture	20240411	May 2024	Privileged Horizontal Committee	riscv/riscv-isa-manual

Demo

```
.global printone
```

```
printone:
```

```
    mv t0, a0          # save the function  
argument: a character pointer
```

```
    # Make a system call: write(0, t0, 1)  
    addi a7, x0, 64    # syscall number: write  
    addi a0, x0, 0     # first argument: fd  
    mv    a1, t0       # second argument: buf  
    addi a2, x0, 1     # third argument: count  
    ecall
```

```
    ret
```

```
int printone(char* c);
```

```
int main() {  
    printone("h");  
    printone("i");  
    printone("\n");  
    return 0;  
}
```

Today's Goals

- Review: Processes
- System Calls: A C Perspective
- Exceptional Control Flow
 - Signals
 - Interrupts
- System Calls: How They Actually Work