


UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## System Control Flow and Processes

CSE 351 Winter 2018

**Instructor:**  
Mark Wyse

**Teaching Assistants:**  
Kevin Bi, Parker DeWilde, Emily Furst,  
Sarah House, Waylon Huang, Vinny Palaniappan



<http://xkcd.com/908/>

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Administrative

- ❖ Homework 4 due Friday (2/23)
- ❖ Lab 4 due next Wednesday (2/28)
  - Cache parameter puzzles and code optimizations

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Roadmap

**C:**

```
car *c = malloc(sizeof(car));
c->miles = 100;
c->gals = 17;
float mpg = get_mpg(c);
free(c);
```

**Java:**

```
Car c = new Car();
c.setMiles(100);
c.setGals(17);
float mpg = c.getMpg();
```

Memory & data  
Integers & floats  
x86 assembly  
Procedures & stacks  
Executables  
Arrays & structs  
Memory & caches  
**Processes**  
Virtual memory  
Memory allocation  
Java vs. C


**Assembly language:**

```
get_mpg:
    pushq %rbp
    movq %rsp, %rbp
    ...
    popq %rbp
    ret
```

**Machine code:**

```
0111010000011000
100011010000010000000010
1000100111000010
11000001111101000011111
```

**OS:** Windows 10, OS X, Linux

**Computer system:** 

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Leading Up to Processes

- ❖ System Control Flow
  - Control flow
  - Exceptional control flow
  - Asynchronous exceptions (interrupts)
  - Synchronous exceptions (traps & faults)

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Control Flow

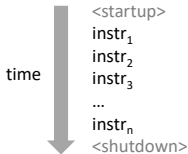
- ❖ **So far:** we've seen how the flow of control changes as a *single program* executes
- ❖ **Reality:** multiple programs running *concurrently*
  - How does control flow across the many components of the system?
  - In particular: More programs running than CPUs
- ❖ **Exceptional control flow** is basic mechanism used for:
  - Transferring control between *processes* and OS
  - Handling I/O and *virtual memory* within the OS
  - Implementing multi-process apps like shells and web servers
  - Implementing concurrency

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Control Flow

- ❖ Processors do only one thing:
  - From startup to shutdown, a CPU simply reads and executes (interprets) a sequence of instructions, one at a time
  - This sequence is the CPU's *control flow* (or *flow of control*)

*Physical control flow*



UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Altering the Control Flow

- Up to now, two ways to change control flow:
  - Jumps (conditional and unconditional)
  - Call and return
  - Both react to changes in *program state*
- Processor also needs to react to changes in *system state*
  - Unix/Linux user hits "Ctrl-C" at the keyboard
  - User clicks on a different application's window on the screen
  - Data arrives from a disk or a network adapter
  - Instruction divides by zero
  - System timer expires
- Can jumps and procedure calls achieve this?
  - No – the system needs mechanisms for "*exceptional*" control flow!

7

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Exceptional Control Flow

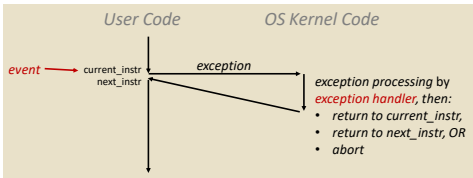
- Exists at all levels of a computer system
- Low level mechanisms
  - Exceptions**
    - Change in processor's control flow in response to a system event (i.e. change in system state, user-generated interrupt)
    - Implemented using a combination of hardware and OS software
- Higher level mechanisms
  - Process context switch**
    - Implemented by OS software and hardware timer
  - Signals**
    - Implemented by OS software
    - We won't cover these – see CSE451 and CSE/EE474

8

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Exceptions

- An *exception* is transfer of control to the operating system (OS) kernel in response to some *event* (i.e. change in processor state)
  - Kernel is the memory-resident part of the OS
  - Examples:** division by 0, page fault, I/O request completes, Ctrl-C



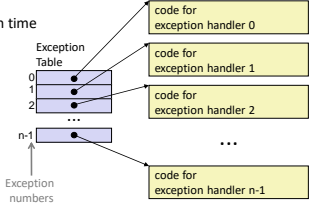
How does the system know where to jump to in the OS?

9

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Exception Table

- A jump table for exceptions (also called *Interrupt Vector Table*)
  - Each type of event has a unique exception number  $k$
  - $k$  = index into exception table (a.k.a interrupt vector)
  - Handler  $k$  is called each time exception  $k$  occurs



10

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Leading Up to Processes

- System Control Flow
  - Control flow
  - Exceptional control flow
  - Asynchronous exceptions (interrupts)**
  - Synchronous exceptions (traps & faults)**

11

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Asynchronous Exceptions (Interrupts)

- Caused by events external to the processor
  - Indicated by setting the processor's interrupt pin(s) (wire into CPU)
  - After interrupt handler runs, the handler returns to "next" instruction
- Examples:**
  - I/O interrupts
    - Hitting Ctrl-C on the keyboard
    - Clicking a mouse button or tapping a touchscreen
    - Arrival of a packet from a network
    - Arrival of data from a disk
  - Timer interrupt
    - Every few ms, an external timer chip triggers an interrupt
    - Used by the OS kernel to take back control from user programs

12

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Synchronous Exceptions

- Caused by events that occur as a result of executing an instruction:
  - Traps**
    - Intentional:** transfer control to OS to perform some function
    - Examples:** system calls, breakpoint traps, special instructions
    - Returns control to "next" instruction
  - Faults**
    - Unintentional** but possibly recoverable
    - Examples:** page faults, segment protection faults, integer divide-by-zero exceptions
    - Either re-executes faulting ("current") instruction or aborts
  - Aborts**
    - Unintentional** and unrecoverable
    - Examples:** parity error, machine check (hardware failure detected)
    - Aborts current program

13

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Traps Example: Opening File

- User calls `open(filename, options)`
- Calls `__open` function, which invokes system call instruction `syscall`

```

0000000000e5d70 <__open>:
...
e5d79: b8 02 00 00 00    mov  $0x2,%eax # open is syscall 2
e5d7e: 0f 05             syscall # return value in %rax
e5d80: 48 3d 01 f0 ff ff  cmp  $0xffffffffffff001,%rax
...
e5dfa: c3               retq
    
```

- `%rax` contains syscall number
- Other arguments in `%rdi`, `%rsi`, `%rdx`, `%r10`, `%r8`, `%r9`
- Return value in `%rax`
- Negative value is an error corresponding to negative `errno`

14

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Fault Example: Page Fault

- User writes to memory location
- That portion (page) of user's memory is currently on disk

```

int a[1000];
int main ()
{
    a[500] = 13;
}
    
```

```

80483b7: c7 05 10 9d 04 08 0d    movl  $0xd,0x8049d10
    
```

- Page fault handler must load page into physical memory
- Returns to faulting instruction: `mov` is executed again!
  - Successful on second try

15

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Fault Example: Invalid Memory Reference

```

int a[1000];
int main()
{
    a[5000] = 13;
}
    
```

```

80483b7: c7 05 60 e3 04 08 0d    movl  $0xd,0x804e360
    
```

- Page fault handler detects invalid address
- Sends `SIGSEGV` signal to user process
- User process exits with "segmentation fault"

16

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Summary (ECF)

- Exceptions
  - Events that require non-standard control flow
  - Generated externally (interrupts) or internally (traps and faults)
  - After an exception is handled, one of three things may happen:
    - Re-execute the current instruction
    - Resume execution with the next instruction
    - Abort the process that caused the exception

17

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Processes

- Processes and context switching**
- Creating new processes
  - `fork()`, `exec*()`, and `wait()`
- Zombies

18

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## What is a process?

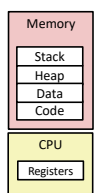
- ❖ Another *abstraction* in our computer system
  - Provided by the OS
  - OS uses a data structure to represent each process
  - Maintains the *interface* between the program and the underlying hardware (CPU + memory)
- ❖ What *do processes* have to do with *exceptional control flow*?
  - Exceptional control flow is the *mechanism* the OS uses to enable **multiple processes** to run on the same system

20

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Processes

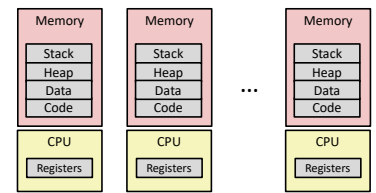
- ❖ A **process** is an instance of a running program
  - One of the most profound ideas in computer science
  - Not the same as “program” or “processor”
- ❖ Process provides each program with **two key abstractions**:
  - **Logical control flow**
    - Each process seems to have exclusive use of the CPU
    - Provided by kernel mechanism called **context switching**
  - **Private address space**
    - Each process seems to have exclusive use of main memory
    - Provided by kernel mechanism called **virtual memory**



21

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Multiprocessing: The Illusion

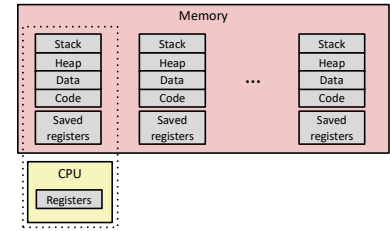


- ❖ Computer runs many processes simultaneously
  - Applications for one or more users
    - Web browsers, email clients, editors, ...
  - Background tasks
    - Monitoring network & I/O devices

22

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Multiprocessing: The Reality

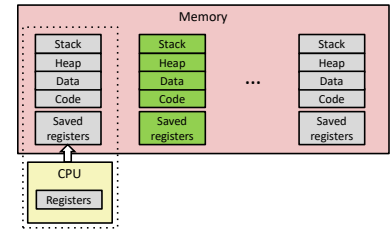


- ❖ Single processor executes multiple processes *concurrently*
  - Process executions interleaved, CPU runs *one at a time*
  - Address spaces managed by virtual memory system (later in course)
  - *Execution context* (register values, stack, ...) for other processes saved in memory

23

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Multiprocessing

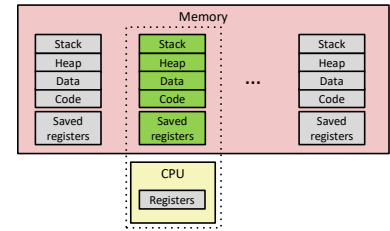


- ❖ **Context switch**
  - 1) Save current registers in memory

24

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Multiprocessing



- ❖ **Context switch**
  - 1) Save current registers in memory
  - 2) Schedule next process for execution (**OS decides**)

25

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Multiprocessing

- Context switch
  - Save current registers in memory
  - Schedule next process for execution (OS decides)
  - Load saved registers and switch address space

26

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

Assume only one CPU

## Concurrent Processes

- Each process is a logical control flow
- Two processes *run concurrently* (are concurrent) if their instruction executions (flows) overlap in time
  - Otherwise, they are *sequential*
- Example: (running on single core)
  - Concurrent: A & B, A & C
  - Sequential: B & C

27

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

Assume only one CPU

## User's View of Concurrency

- Control flows for concurrent processes are physically disjoint in time
  - CPU only executes instructions for one process at a time
- However, the user can *think of* concurrent processes as executing at the same time, in *parallel*

28

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

Assume only one CPU

## Context Switching

- Processes are managed by a *shared* chunk of OS code called the **kernel**
  - The kernel is not a separate process, but rather runs as part of a user process
- In x86-64 Linux:
  - Same address in each process refers to same shared memory location

29

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

Assume only one CPU

## Context Switching

- Processes are managed by a *shared* chunk of OS code called the **kernel**
  - The kernel is not a separate process, but rather runs as part of a user process
- Context switch passes control flow from one process to another and is performed using kernel code

30

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Creating Processes & Reaping Zombies

- Processes and context switching
- Creating new processes
  - `fork()`, `exec*()`, and `wait()`
- Zombies

31

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Creating New Processes & Programs

32

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Creating New Processes & Programs

- ❖ fork-exec model (Linux):
  - `fork()` creates a copy of the current process
  - `exec*()` replaces the current process' code and address space with the code for a different program
    - Family: `execv`, `execl`, `execve`, `execle`, `execvp`, `execlp`
  - `fork()` and `execve()` are *system calls*
- ❖ Other system calls for process management:
  - `getpid()`
  - `exit()`
  - `wait()`, `waitpid()`

33

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## fork: Creating New Processes

- ❖ `pid_t fork(void)`
  - Creates a new "child" process that is *identical* to the calling "parent" process, including all state (memory, registers, etc.)
  - Returns 0 to the **child** process
  - Returns child's **process ID (PID)** to the **parent** process
- ❖ Child is *almost* identical to parent:
  - Child gets an identical (but separate) copy of the parent's virtual address space
  - Child has a different PID than the parent

```
pid_t pid = fork();
if (pid == 0) {
    printf("hello from child\n");
} else {
    printf("hello from parent\n");
}
```
- ❖ `fork` is unique (and often confusing) because it is called **once** but returns **"twice"**

34

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Fork Example

```
void fork1() {
    int x = 1;
    pid_t pid = fork();
    if (pid == 0)
        printf("Child has x = %d\n", ++x);
    else
        printf("Parent has x = %d\n", --x);
    printf("Bye from process %d with x = %d\n", getpid(), x);
}
```

- ❖ Both processes continue/start execution after `fork`
  - Child starts at instruction after the call to `fork` (storing into `pid`)
- ❖ Can't predict execution order of parent and child
- ❖ Both processes start with `x=1`
  - Subsequent changes to `x` are independent
- ❖ Shared open files: `stdout` is the same in both parent and child

35

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Fork-Exec

Note: the return values of `fork` and `exec*` should be checked for errors

- ❖ fork-exec model:
  - `fork()` creates a copy of the current process
  - `exec*()` replaces the current process' code and address space with the code for a different program
    - Whole family of `exec` calls – see **exec(3)** and **execve(2)**

```
// Example arguments: path="/usr/bin/ls",
// argv[0]="/usr/bin/ls", argv[1]="-ahl", argv[2]=NULL
void fork_exec(char *path, char *argv[]) {
    pid_t pid = fork();
    if (pid != 0) {
        printf("Parent: created a child %d\n", pid);
    } else {
        printf("Child: about to exec a new program\n");
        execv(path, argv);
    }
    printf("This line printed by parent only!\n");
}
```

36

UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Exec-ing a new program

Very high-level diagram of what happens when you run the command "ls" in a Linux shell:

- ❖ This is the loading part of CALL!

37

W UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## exit: Ending a process

- ❖ `void exit(int status)`
  - Exits a process
    - Status code: 0 is used for a normal exit, nonzero for abnormal exit

40

W UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Zombies

- ❖ When a process terminates, it still consumes system resources
  - Various tables maintained by OS
  - Called a “zombie” (a living corpse, half alive and half dead)
- ❖ *Reaping* is performed by parent on terminated child
  - Parent is given exit status information and kernel then deletes zombie child process
- ❖ What if parent doesn't reap?
  - If any parent terminates without reaping a child, then the orphaned child will be reaped by `init` process (`pid == 1`)
    - **Note:** on more recent Linux systems, `init` has been renamed `systemd`
  - In long-running processes (e.g. shells, servers) we need *explicit* reaping

41

W UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## wait: Synchronizing with Children

- ❖ `int wait(int *child_status)`
  - Suspends current process (*i.e.* the parent) until one of its children terminates
  - Return value is the PID of the child process that terminated
    - *On successful return, the child process is reaped*
  - If `child_status != NULL`, then the `*child_status` value indicates why the child process terminated
    - Special macros for interpreting this status – see `man wait (2)`
- ❖ **Note:** If parent process has multiple children, `wait` will return when *any* of the children terminates
  - `waitpid` can be used to wait on a specific child process

42

W UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Process Management Summary

- ❖ `fork` makes two copies of the same process (parent & child)
  - Returns different values to the two processes
- ❖ `exec*` replaces current process from file (new program)
  - Two-process program:
    - First `fork()`
    - `if (pid == 0) { /* child code */ } else { /* parent code */ }`
  - Two different programs:
    - First `fork()`
    - `if (pid == 0) { execv(...) } else { /* parent code */ }`
- ❖ `wait` or `waitpid` used to synchronize parent/child execution and to reap child process

43

W UNIVERSITY of WASHINGTON L19: System Control Flow and Processes CSE351, Winter 2018

## Summary

- ❖ Processes
  - At any given time, system has multiple active processes
  - On a one-CPU system, only one can execute at a time, but each process appears to have total control of the processor
  - OS periodically “context switches” between active processes
    - Implemented using *exceptional control flow*
- ❖ Process management
  - `fork`: one call, two returns
  - `execve`: one call, usually no return
  - `wait` or `waitpid`: synchronization
  - `exit`: one call, no return

44