

Exceptional Control Flow: Exceptions and Processes

CSci 2021: Machine Architecture and Organization
December 3rd, 2018

Your instructor: Stephen McCamant


Based on slides originally by:
Randy Bryant, Dave O'Hallaron

Today

- Exceptional Control Flow
- Exceptions
- Processes
- Process Control

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*)



Altering the Control Flow

- Up to now: two mechanisms for changing control flow:
 - Jumps and branches
 - Call and returnReact to changes in *program state*
- Insufficient for a useful system:
Difficult to react to changes in *system state*
 - Data arrives from a disk or a network adapter
 - Instruction divides by zero
 - User hits Ctrl-C at the keyboard
 - System timer expires
- System needs mechanisms for "exceptional control flow"

Exceptional Control Flow


- Exists at all levels of a computer system
- Low level mechanisms
 - 1. **Exceptions**
 - Change in control flow in response to a system event (i.e., change in system state)
 - Implemented using combination of hardware and OS software
- Higher level mechanisms
 - 2. **Process context switch**
 - Implemented by OS software and hardware timer
 - 3. **Signals**
 - Implemented by OS software
 - 4. **Nonlocal jumps**: `setjmp()` and `longjmp()`
 - Implemented by C runtime library

Today

- Exceptional Control Flow
- Exceptions
- Processes
- Process Control


Exceptions

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



Exception Tables

Exception numbers



- 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

Asynchronous Exceptions (Interrupts)

- **Caused by events external to the processor**
 - Indicated by setting the processor's *interrupt pin*
 - Handler returns to "next" instruction
- **Examples:**
 - Timer interrupt
 - Every few ms, an external timer chip triggers an interrupt
 - Used by the kernel to take back control from user programs
 - I/O interrupt from external device
 - Hitting Ctrl-C at the keyboard
 - Arrival of a packet from a network
 - Arrival of data from a disk

Synchronous Exceptions

- **Caused by events that occur as a result of executing an instruction:**
 - **Traps**
 - Intentional
 - Examples: **system calls**, breakpoint traps, special instructions
 - Returns control to "next" instruction
 - **Faults**
 - Unintentional but possibly recoverable
 - Examples: page faults (recoverable), protection faults (unrecoverable), floating point exceptions
 - Either re-executes faulting ("current") instruction or aborts
 - **Aborts**
 - Unintentional and unrecoverable
 - Examples: illegal instruction, parity error, machine check
 - Aborts current program

System Calls

- Each x86-64 system call has a **unique ID number**
- **Examples:**


Number	Name	Description
0	read	Read file
1	write	Write file
2	open	Open file
3	close	Close file
4	stat	Get info about file
57	fork	Create process
59	execve	Execute a program
60	_exit	Terminate process
62	kill	Send signal to process

System Call 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 $0xfffffffffffff001,%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`

Fault Example: Page Fault

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

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


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



Fault Example: Invalid Memory Reference

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

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




- Sends SIGSEGV signal to user process
- User process exits with "segmentation fault"

Today


- Exceptional Control Flow
- Exceptions
- Processes
- Process Control

Processes

- Definition: 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 program seems to have exclusive use of the CPU
 - Provided by kernel mechanism called *context switching*
 - Private address space**
 - Each program seems to have exclusive use of main memory.
 - Provided by kernel mechanism called *virtual memory*



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


Multiprocessing Example

```
Process: 123 total, 5 running, 9 stuck, 109 sleeping, 611 threads
Load Avg: 1.03, 1.13, 1.14 CPU usage: 3.27% user, 5.35% sys, 91.56% idle
SharedLibs: 578K resident, 0B data, 0B linkedit.
MemRegions: 27369 total, 11271 resident, 35H private, 434H shared,
PhysMem: 1629H wired, 1374H active, 1062H inactive, 4078H used, 18H free.
VM: 280G vmize, 1091H framework vmize, 23075213(1) pageins, 584337(0) pageouts.
Networks: packets: 41046228/115 in, 66483096/776 out.
Disks: 1787433/2495 read, 1294727/2546 write.
```

PID	COMMAND	%CPU	TIME	MM	MPORT	#INEG	PPRVT	PSHED	RSIZE	VRPNT	VSIZE	
98217	Microsoft Of	0.0	02:28.34	4	1	202	418	21M	24M	21M	65H	763H
99161	usbmuxd	0.0	00:04.10	3	1	47	66	430K	218K	480K	60H	2422H
99406	iTunesHelper	0.0	00:01.23	2	1	55	78	728K	3124K	1124K	43H	2423H
94298	bash	0.0	00:00.11	1	0	20	24	224K	720K	484K	17H	2378H
84285	xterm	0.0	00:00.83	1	0	32	72	656K	872K	630K	3728K	2382H
95333	Microsoft Ex	0.3	21:58.37	10	3	360	954	16H	60H	40H	114H	1057H
54751	sleep	0.0	00:00.00	1	0	17	20	80K	212K	390K	8620K	2370H
54733	launchdadd	0.0	00:00.00	2	1	33	50	488K	220K	1736K	48H	2409H
54737	top	6.5	00:02.53	1/1	0	30	29	1416K	216K	2124K	17H	2379H
54719	automountd	0.0	00:00.02	7	1	53	64	880K	218K	2184K	53H	2413H
54701	ocspd	0.0	00:00.05	4	1	61	54	1268K	2644K	3132K	50H	2426H
54651	Grab	0.6	00:02.75	6	3	222	389	150H	261H	400H	75H	2526H
54659	cooliepd	0.0	00:00.15	2	1	40	61	3316K	224K	4088K	62H	2411H
82818	automountd	0.0	00:01.57	4	1	82	91	7628K	7419H	15M	48H	2430H


- Running program "top" on Mac
 - System has 123 processes, 5 of which are active
 - Identified by Process ID (PID)

Multiprocessing: The (Traditional) Reality



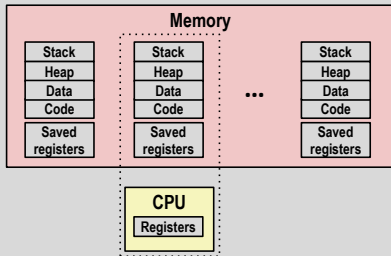
- Single processor executes multiple processes concurrently
 - Process executions interleaved (multitasking)
 - Address spaces managed by virtual memory system (later in course)
 - Register values for nonexecuting processes saved in memory

Multiprocessing: The (Traditional) Reality



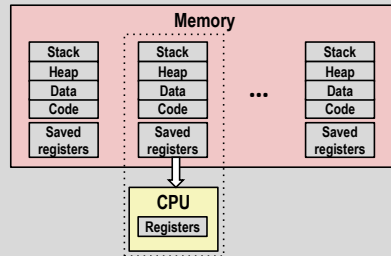
- Save current registers in memory

Multiprocessing: The (Traditional) Reality




- Schedule next process for execution

Multiprocessing: The (Traditional) Reality



- Load saved registers and switch address space (context switch)


Multiprocessing: The (Modern) Reality



- Multicore processors
 - Multiple CPUs on single chip
 - Share main memory (and some of the caches)
 - Each can execute a separate process
 - Scheduling of processors onto cores done by kernel


Concurrent Processes

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




User View of Concurrent Processes

- Control flows for concurrent processes are physically disjoint in time
- However, we can think of concurrent processes as running in parallel with each other



Context Switching

- Processes are managed by a shared chunk of memory-resident OS code called the *kernel*
 - Important: the kernel is not a separate process, but rather runs as part of some existing process.
- Control flow passes from one process to another via a *context switch*



Today

- Exceptional Control Flow
- Exceptions
- Processes
- System calls and process startup

System Call Error Handling

- On error, Linux system-level functions typically return `-1` and set global variable `errno` to indicate cause.
- Hard and fast rule:
 - You must check the return status of every system-level function
 - Only exception is the handful of functions that return `void`
- Example:

```
if ((pid = fork()) < 0) {
    fprintf(stderr, "fork error: %s\n", strerror(errno));
    exit(0);
}
```

Error-reporting functions

- Can simplify somewhat using an *error-reporting function*:

```
void unix_error(char *msg) /* Unix-style error */
{
    fprintf(stderr, "%s: %s\n", msg, strerror(errno));
    exit(0);
}
```

```
if ((pid = fork()) < 0)
    unix_error("fork error");
```

Error-handling Wrappers

- We simplify the code we present to you even further by using Stevens-style error-handling wrappers:

```
pid_t Fork(void)
{
    pid_t pid;

    if ((pid = fork()) < 0)
        unix_error("Fork error");
    return pid;
}
```

```
pid = Fork();
```

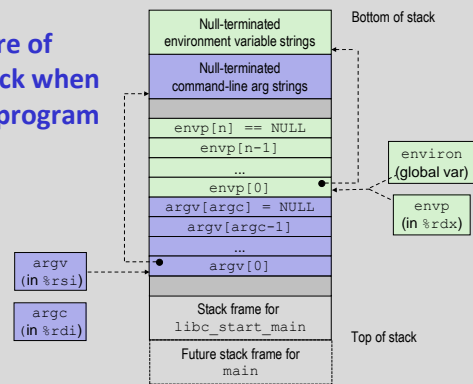
execve : Loading and Running Programs

- `int execve(char *filename, char *argv[], char *envp[])`
- **Loads and runs in the current process:**
 - Executable file `filename`
 - Can be object file or script file beginning with `#!` interpreter (e.g., `#!/bin/bash`)
 - ...with argument list `argv`
 - By convention `argv[0]==filename`
 - ...and environment variable list `envp`
 - "name=value" strings (e.g., `USER=droh`)
 - `getenv`, `putenv`, `printenv`
- **Overwrites code, data, and stack**
 - Retains PID, open files and signal context
- **Called once and never returns**
 - ...except if there is an error

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

49

Structure of the stack when a new program starts




Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

50

execve Example

- Executes `"/bin/ls -lt /usr/include"` in child process using current environment:



```

if ((pid = Fork()) == 0) { /* Child runs program */
    if (execve(myargv[0], myargv, environ) < 0) {
        printf("%s: Command not found.\n", myargv[0]);
        exit(1);
    }
}

```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

51

Summary

- **Exceptions**
 - Events that require nonstandard control flow
 - Generated externally (interrupts) or internally (traps and faults)
- **Processes**
 - At any given time, system has multiple active processes
 - Only one can execute at a time on a single core, though
 - Each process appears to have total control of processor + private memory space

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

52