


# Exceptions and Processes

Samira Khan  
April 18, 2017

## 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*)



2

## Altering the Control Flow

- Up to now: two mechanisms for changing control flow:
  - Jumps and branches
  - Call and return
 React 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”

3

## 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
    - Overlaps execution with useful work from other process
  - 3. **Signals**
    - Implemented by OS software

4


## Today

- Exceptional Control Flow
- **Exceptions**
- Processes
- Process Control

5


## 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



6

## Exception Tables




7

## Running the Exception Handler

- Hardware saves the old program counter
- Identifies location of exception handler via table
- Then jumps to that location
- OS code can save registers, etc.

8

## (partial) Taxonomy



9

## 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

10

## 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

11

## Note on Terminology

- Real world does not use consistent terms for exceptions
- We will follow textbook's terms in this course
- However, in real world:
  - 'interrupt' meaning what we call 'exception' (x86)
  - 'exception' meaning what we call 'fault'
  - 'fault' meaning what we call 'fault' or 'abort' (ARM)
  - ... and more

12

## 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

13


## 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    mov $0x2,%eax # open is syscall #2
e5d7e: 0f 05          syscall # Return value in %rax
e5d80: 48 3d 01 ff ff cmp $0xfffffffffff001,%rax
...
e5dfa: c3            retq

```



14

## System Call


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

```

0000000000e5d70: <__open>:
...
e5d79: b8 02 00 00    mov $0x2,%eax # open is syscall #2
e5d7e: 0f 05          syscall # Return value in %rax
e5d80: 48 3d 01 ff ff cmp $0xfffffffffff001,%rax
...
e5dfa: c3            retq

```

- Almost like a function call
  - Transfer of control
  - On return, executes next instruction
  - Passes arguments using calling convention
  - Gets result in `%rax`
- One Important exception!
  - Executed by Kernel
  - Different set of privileges
- And other differences:
  - E.g., "address" of "function" is in `%rax`
  - Uses `errno`
  - Etc.



15

## Fault Example: Page Fault

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

```

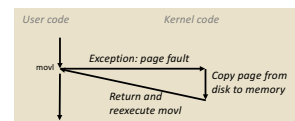
int a[1000];
main()
{
    a[500] = 13;
}

```

```

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

```




16

### Fault Example: Invalid Memory Reference

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

```
80483b7: c7 05 60 e3 04 08 0d movl 0x04e360,0x04e360
```



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

17


### Today

- Exceptional Control Flow
- Exceptions
- Processes
- Process Control

18


### 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*



19

### 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

20

## Multiprocessing Example

```

xterm
Processes: 123 total, 5 running, 9 stuck, 109 sleeping, 611 threads
Load Avg: 1.05, 1.15, 1.14 (CPU usage: 3.07% user, 5.15% sys, 91.78% idle)
SharedLib: 570K resident, 0B data, 0B inedit.
MemRegions: 22969 total, 1270 resident, 20R private, 494K shared,
PhysMem: 1039M wired, 1974M active, 1062M inactive, 4070M used, 19M free,
Virt: 20G vmem, 105M framework vmem, 230702513 pageins, 59433070 pageouts,
Network: packets: 4104628/116 in, 6603936/776 out,
Tasks: 176/4261/946: read, 1204/375/946: written.


PID  COMMAND   CPU TIME  WTH  WIO  WPOUT  WFREE  WFWIT  RWNO  RSIZE  WFWIT  WSIZE
99517  Microsoft  0.0 00:10.34 4 1 202 418 218 241 231 60R 762M
99518  vim        0.0 00:00.10 3 1 47 86 430K 210K 40R 0M 242M
99006  iTerm2     0.0 00:01.22 2 1 95 70 720K 3124K 1124K 43M 242M
94306  bash      0.0 00:00.11 1 0 20 24 224K 72K 40K 17M 227M
94305  xterm     0.0 00:00.05 1 0 32 73 656K 872K 65K 972M 230M
99539  Microsoft  0.3 21:58.37 10 3 360 954 10R 891 40R 114M 1057M
54781  sleep     0.0 00:00.00 1 0 17 20 92K 212K 30K 862K 227M
54739  launchd   0.0 00:00.00 2 1 33 50 488K 220K 173K 48M 240M
54727  top       6.5 00:00.33 1/1 0 30 29 1410K 210K 2234K 17M 227M
54719  autowatch 0.0 00:00.02 7 1 53 64 860K 210K 2184K 63M 2413M
54710  occid     0.0 00:00.06 4 1 61 64 1308K 2644K 3132K 50M 2452M
54681  Erab      0.6 00:02.75 6 3 222 389 159K 39M 40M 75M 2559M
54659  combined 0.0 00:00.15 2 1 40 62 2310K 224K 4000K 42M 2411M

```

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

21


## 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

22


## Multiprocessing: The (Traditional) Reality



- Save current registers in memory

23


## Multiprocessing: The (Traditional) Reality



- Schedule next process for execution

24


## Multiprocessing: The (Traditional) Reality



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

25

## Multiprocessing: The (Modern) Reality




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

26

## 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



27

## 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



28

## 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*



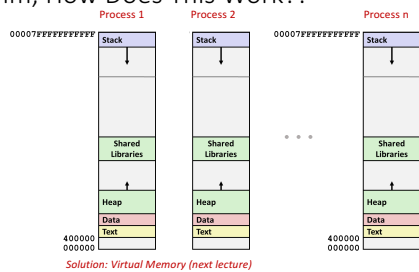
29

## Context

- all registers values
- %rax %rbx, ..., %rsp, ...
- condition codes
- program counter
- i.e. all visible state in your CPU except memory

30

## Hmmm, How Does This Work?!



31

## Context

- all registers values
- %rax %rbx, ..., %rsp, ...
- condition codes
- program counter
- address space: map from program to real addresses

32



## Today

- Exceptional Control Flow
- Exceptions
- Processes
- **Process Control**

33

## 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(-1);
}
```

34

## 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(-1);
}
```

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

35

## Error-handling Wrappers

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

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

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

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

- NOT what you generally want to do in a real application

36

## Creating and Terminating Processes

From a programmer's perspective, we can think of a process as being in one of three states

- **Running**
  - Process is either executing, or waiting to be executed and will eventually be *scheduled* (i.e., chosen to execute) by the kernel
- **Stopped**
  - Process execution is *suspended* and will not be scheduled until further notice (next lecture when we study signals)
- **Terminated**
  - Process is stopped permanently

37

## Terminating Processes

- Process becomes terminated for one of three reasons:
  - Receiving a signal whose default action is to terminate (next lecture)
  - Returning from the **main** routine
  - Calling the **exit** function
- `void exit(int status)`
  - Terminates with an *exit status* of **status**
  - Convention: normal return status is 0, nonzero on error
  - Another way to explicitly set the exit status is to return an integer value from the main routine
- **exit** is called **once** but **never** returns.

38

## Creating Processes

- *Parent process* creates a new running *child process* by calling **fork**
- `int fork(void)`
  - Returns 0 to the child process, child's PID to parent process
  - Child is *almost* identical to parent:
    - Child get an identical (but separate) copy of the parent's virtual address space.
    - Child has a different PID than the parent
- **fork** is interesting (and often confusing) because it is called **once** but returns **twice**

39

## fork Example

```
int main(int argc, char** argv)
{
    pid_t pid;
    int x = 1;

    pid = Fork();
    if (pid == 0) { /* Child */
        printf("child : x=%d\n", ++x);
        return 0;
    }

    /* Parent */
    printf("parent: x=%d\n", --x);
    return 0;
}
```

fork.c

- Call once, return twice
- Concurrent execution
  - Can't predict execution order of parent and child

linux> ./fork parent: x=0 child : x=2	linux> ./fork child : x=2 parent: x=0	linux> ./fork parent: x=0 child : x=2	linux> ./fork parent: x=0 child : x=2
---	---	---	---

40

## fork Example

```
int main(int argc, char** argv)
{
    pid_t pid;
    int x = 1;

    pid = Fork();
    if (pid == 0) { /* Child */
        printf("child : x=%d\n", ++x);
        printf("child : x=%d\n", ++x);
        return 0;
    }
    /* Parent */
    printf("parent: x=%d\n", --x);
    printf("parent: x=%d\n", --x);
    return 0;
}
```

- Call once, return twice
- Concurrent execution
  - Can't predict execution order of parent and child
- Duplicate but separate address space
  - x has a value of 1 when fork returns in parent and child
  - Subsequent changes to x are independent

```
linux> ./fork
parent: x=0
child : x=0
parent: x=-1
child : x=3
```

43

## fork Example

```
int main(int argc, char** argv)
{
    pid_t pid;
    int x = 1;

    pid = Fork();
    if (pid == 0) { /* Child */
        printf("child : x=%d\n", ++x);
        return 0;
    }
    /* Parent */
    printf("parent: x=%d\n", --x);
    return 0;
}
```

- Call once, return twice
- Concurrent execution
  - Can't predict execution order of parent and child
- Duplicate but separate address space
  - x has a value of 1 when fork returns in parent and child
  - Subsequent changes to x are independent
  - stdout is the same in both parent and child

```
linux> ./fork
parent: x=0
child : x=2
```

42

## Modeling fork with Process Graphs


- A *process graph* is a useful tool for capturing the partial ordering of statements in a concurrent program:
  - Each vertex is the execution of a statement
  - a → b means a happens before b
  - Edges can be labeled with current value of variables

43

## Process Graph Example

```
int main(int argc, char** argv)
{
    pid_t pid;
    int x = 1;

    pid = Fork();
    if (pid == 0) { /* Child */
        printf("child : x=%d\n", ++x);
        return 0;
    }
    /* Parent */
    printf("parent: x=%d\n", --x);
    return 0;
}
```




44

## fork Example: Two consecutive forks

```
void fork2()
{
    printf("L0\n");
    fork();
    printf("L1\n");
    fork();
    printf("Bye\n");
}
```

*forks.c*



Feasible output:

```
L0
L1
Bye
Bye
L1
Bye
Bye
```

Infeasible output:


```
L0
L0
L1
Bye
Bye
L1
Bye
Bye
```

45

## fork Example: Nested forks in parent

```
void fork4()
{
    printf("L0\n");
    if (fork() != 0) {
        printf("L1\n");
        if (fork() != 0) {
            printf("L2\n");
        }
    }
    printf("Bye\n");
}
```

*forks.c*



Feasible output:

```
L0
L1
Bye
Bye
L2
Bye
```

Infeasible output:


```
L0
L0
Bye
Bye
L2
Bye
L2
```

46

## fork Example: Nested forks in children

```
void fork5()
{
    printf("L0\n");
    if (fork() == 0) {
        printf("L1\n");
        if (fork() == 0) {
            printf("L2\n");
        }
    }
    printf("Bye\n");
}
```

*forks.c*



Feasible output:

```
L0
L0
Bye
L1
Bye
L2
Bye
```

Infeasible output:

```
L0
L0
Bye
L1
Bye
L2
L2
```

47

## 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 (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`)

## • Overwrites code, data, and stack

- Retains PID, open files and signal context


• Called **once** and **never** returns

- ...except if there is an error

48

## execve Example

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



```

1 if ((pid = Fork()) == 0) { /* Child runs program */
2   if (execve(myargv[0], myargv, environ) < 0) {
3     printf("%s: Command not found.\n", myargv[0]);
4     exit(1);
5   }
6 }

```

40

## 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 any single core
  - Each process appears to have total control of processor + private memory space

50

## Summary (cont.)

- Spawning processes
  - Call `fork`
  - One call, two returns
- Process completion
  - Call `exit`
  - One call, no return
- Loading and running programs
  - Call `execve` (or variant)
  - One call, (normally) no return

51

## Exceptions and Processes

Samira Khan  
April 18, 2017