

15-213

Exceptional Control Flow I

March 12, 2002

Topics

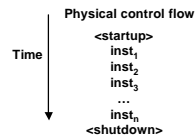
- Exceptions
- Process context switches
- Creating and destroying processes
- Reading: 8.1-8.5
- Problems: 8.10-8.11

class16.ppt

Control flow

From startup to shutdown, a CPU simply reads and executes (interprets) a sequence of instructions, one at a time.

This sequence is the system's physical *control flow* (or *flow of control*).



class16.ppt

- 2 -

CS 213 F'01

Altering the Control Flow

We've discussed two mechanisms for changing the control flow:

- Jumps and branches
- Call and return using the stack discipline.
- Both react to changes in program state.

Insufficient for a useful system

- difficult for the CPU 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”

class16.ppt

- 3 -

CS 213 F'01

Exceptional control flow

Mechanisms for exceptional control flow exists at all levels of a computer system.

Low level mechanism:

- exceptions
 - change in control flow in response to a system event (i.e., change in system state)
- Implemented as a combination of both hardware and OS software

Higher level mechanisms:

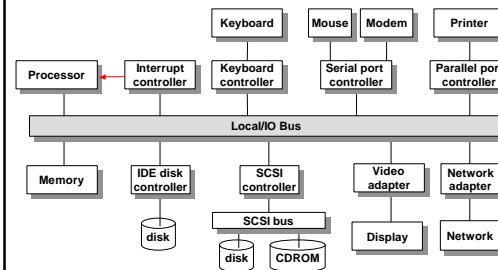
- process context switch
- signals
- nonlocal jumps (setjmp/longjmp)
 - Implemented by either:
 - OS software (context switch and signals).
 - C language runtime library: nonlocal jumps.

class16.ppt

- 4 -

CS 213 F'01

System context for exceptions



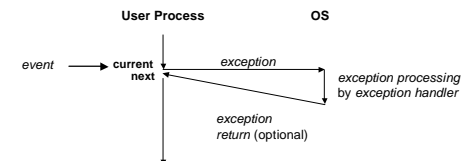
class16.ppt

- 5 -

CS 213 F'01

Exceptions

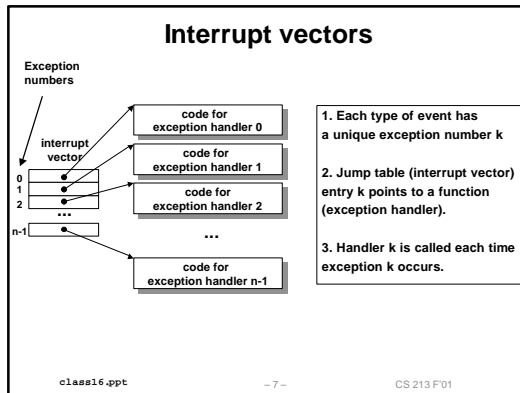
An *exception* is a transfer of control to the OS in response to some *event* (i.e., change in processor state)



class16.ppt

- 6 -

CS 213 F'01



Asynchronous exceptions (interrupts)

Caused by events (changes in state) external to the processor

- Indicated by setting the processor's interrupt pin
- handler returns to "next" instruction.

Examples:

- I/O interrupts**
 - hitting ctrl-c at the keyboard
 - arrival of a packet from a network
 - arrival of a data sector from a disk
- Hard reset interrupt**
 - hitting the reset button
- Soft reset interrupt**
 - hitting ctrl-alt-delete on a PC

class16.ppt - 8 - CS 213 F'01

Synchronous exceptions

Caused by events (changes in state) that occur as a result of executing an instruction:

- Traps**
 - intentional
 - returns control to "next" instruction
 - Examples: system calls, breakpoint traps
- Faults**
 - unintentional but possibly recoverable
 - either re-executes faulting ("current") instruction or aborts.
 - Examples: page faults (recoverable), protection faults (unrecoverable).
- Aborts**
 - unintentional and unrecoverable
 - aborts current program
 - Examples: parity error, machine check.

class16.ppt - 9 - CS 213 F'01

Trap Example

Opening a File

- User calls `open(filename, options)`

```

0804d070 <_libc_open>:
. . .
804d082: cd 80          int    $0x80
804d084: 5b            pop     %ebx
. . .
  
```

— Function `open` executes system call instruction `int`

- OS must find or create file, get it ready for reading or writing
- Returns integer file descriptor

class16.ppt - 10 - CS 213 F'01

Fault Example #1

Memory Reference

- 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
  
```

- Page handler must load page into physical memory
- Returns to faulting instruction
- Successful on second try

class16.ppt - 11 - CS 213 F'01

Fault Example #2

Memory Reference

- User writes to memory location
- Address is not valid

```

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

```

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

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

class16.ppt - 12 - CS 213 F'01

Processes

Def: A process is an instance of a running program.

- One of the most profound ideas in computer science.

Process provides each program with two key abstractions:

- Logical control flow**
 - gives each program the illusion that it has exclusive use of the CPU.
- Private address space**
 - gives each program the illusion that it has exclusive use of main memory.

How is this illusion maintained?

- Process executions interleaved (multitasking)
- Address spaces managed by virtual memory system

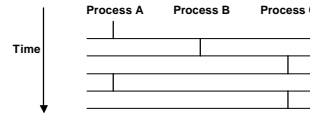
class16.ppt

- 13 -

CS 213 F'01

Logical control flows

Each process has its own logical control flow



class16.ppt

- 14 -

CS 213 F'01

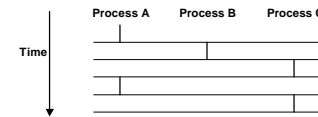
Concurrent processes

Two processes *run concurrently (are concurrent)* if their flows overlap in time.

Otherwise, they are *sequential*.

Examples:

- Concurrent: A & B, A&C
- Sequential: B & C



class16.ppt

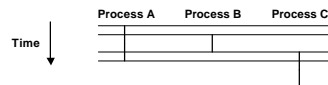
- 15 -

CS 213 F'01

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.



class16.ppt

- 16 -

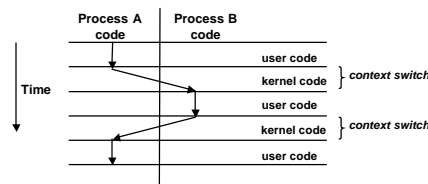
CS 213 F'01

Context switching

Processes are managed by a shared chunk of OS code called the *kernel*

- Important: the kernel is not a separate process, but rather runs as part of some user process

Control flow passes from one process to another via a *context switch*.



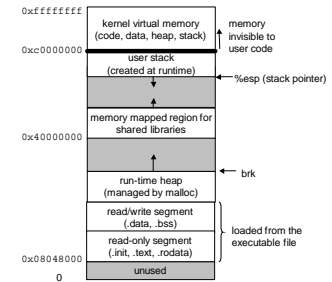
class16.ppt

- 17 -

CS 213 F'01

Private address spaces

Each process has its own private address space.



class16.ppt

- 18 -

CS 213 F'01

fork: Creating new processes

`int fork(void)`

- creates a new process (child process) that is identical to the calling process (parent process)
- returns 0 to the child process
- returns child's pid to the parent process

```
if (fork() == 0) {
    printf("hello from child\n");
} else {
    printf("hello from parent\n");
}
```

Fork is interesting
(and often confusing)
because it is called
once but returns *twice*

class16.ppt

- 19 -

CS 213 F'01

Fork Example #1

Key Points

- Parent and child both run same code
 - Distinguish parent from child by return value from `fork`
- Start with same state, but each has private copy
 - Including shared output file descriptor
 - Relative ordering of their print statements undefined

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

class16.ppt

- 20 -


CS 213 F'01

Fork Example #2

Key Points

- Both parent and child can continue forking

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



class16.ppt

- 21 -


CS 213 F'01

Fork Example #3

Key Points

- Both parent and child can continue forking

```
void fork3()
{
    printf("L0\n");
    fork();
    printf("L1\n");
    fork();
    printf("L2\n");
    fork();
    printf("Bye\n");
}
```



class16.ppt

- 22 -


CS 213 F'01

Fork Example #4

Key Points

- Both parent and child can continue forking

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



class16.ppt

- 23 -


CS 213 F'01

Fork Example #5

Key Points

- Both parent and child can continue forking

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



class16.ppt

- 24 -

CS 213 F'01

exit: Destroying process

void exit(int status)

- **exits a process**
 - Normally return with status 0
- **atexit()** registers functions to be executed upon exit

```
void cleanup(void) {
    printf("cleaning up\n");
}

void fork6() {
    atexit(cleanup);
    fork();
    exit(0);
}
```

class16.ppt

– 25 –

CS 213 F'01

Zombies

Idea

- What process terminates, still consumes system resources
 - Various tables maintained by OS
- Called a "zombie"
 - Living corpse, half alive and half dead

Reaping

- Performed by parent on terminated child
- Parent is given exit status information
- Kernel discards process

What if Parent Doesn't Reap?

- If any parent terminates without reaping a child, then child will be reaped by init process
- Only need explicit reaping for long-running processes
 - E.g., shells and servers

class16.ppt

– 26 –

CS 213 F'01

Zombie Example

```
void fork7()
{
    if (fork() == 0) {
        /* Child */
        printf("Terminating Child, PID = %d\n",
            getpid());
        exit(0);
    } else {
        printf("Running Parent, PID = %d\n",
            getpid());
        while (1)
            ; /* Infinite loop */
    }
}
```

```
linux> ./forks 7 &
[1] 6639
Running Parent, PID = 6639
Terminating Child, PID = 6640
linux> ps
  PID TTY          TIME CMD
 6585 ttys9    00:00:00 tcsh
 6639 ttys9    00:00:03 forks
 6640 ttys9    00:00:00 forks <defunct>
 6641 ttys9    00:00:00 ps
linux> kill 6639
[1] Terminated
linux> ps
  PID TTY          TIME CMD
 6585 ttys9    00:00:00 tcsh
 6642 ttys9    00:00:00 ps
```

- ps shows child process as "defunct"
- Killing parent allows child to be reaped

class16.ppt

– 27 –

CS 213 F'01

Nonterminating Child Example

```
void fork8()
{
    if (fork() == 0) {
        /* Child */
        printf("Running Child, PID = %d\n",
            getpid());
        while (1)
            ; /* Infinite loop */
    } else {
        /* Parent */
        printf("Terminating Parent, PID = %d\n",
            getpid());
        exit(0);
    }
}
```

```
linux> ./forks 8
Terminating Parent, PID = 667
Running Child, PID = 6676
linux> ps
  PID TTY          TIME CMD
 6585 ttys9    00:00:00 tcsh
 6676 ttys9    00:00:06 forks
 6677 ttys9    00:00:00 ps
linux> kill 6676
linux> ps
  PID TTY          TIME CMD
 6585 ttys9    00:00:00 tcsh
 6678 ttys9    00:00:00 ps
```

- ps shows child process as "defunct"
- Killing parent allows child to be reaped

class16.ppt

– 28 –

CS 213 F'01

wait: Synchronizing with children

int wait(int *child_status)

- suspends current process until one of its children terminates
- return value = the pid of the child process that terminated
- if child_status != NULL, then the object it points to will be set to a status indicating why the child process terminated

```
void fork8() {
    int child_status;

    if (fork() == 0) {
        printf("HC: hello from child\n");
    }
    else {
        printf("HP: hello from parent\n");
        wait(&child_status);
        printf("CT: child has terminate\n");
    }
    printf("Bye\n");
    exit();
}
```

class16.ppt

– 29 –

CS 213 F'01

Wait Example

- If multiple child completed, will take in arbitrary order
- Can use macros WIFEXITED and WEXITSTATUS to get information about exit status

```
void fork9()
{
    pid_t pid[N];
    int i;
    int child_status;
    for (i = 0; i < N; i++)
        if ((pid[i] = fork()) == 0)
            exit(100+i); /* Child */
    for (i = 0; i < N; i++) {
        pid_t wpid = wait(&child_status);
        if (WIFEXITED(child_status))
            printf("Child %d terminated with exit status %d",
                wpid, WEXITSTATUS(child_status));
        else
            printf("Child %d terminate abnormally\n", wpid);
    }
}
```

class16.ppt

– 30 –

CS 213 F'01

Waitpid

```
• waitpid(pid, &status, options)
  -Can wait for specific process
  -Various options

void fork10()
{
    pid_t pid[N];
    int i;
    int child_status;
    for (i = 0; i < N; i++)
        if ((pid[i] = fork()) == 0)
            exit(100+i); /* Child */
    for (i = 0; i < N; i++) {
        pid_t wpid = waitpid(pid[i], &child_status, 0);
        if (WIFEXITED(child_status))
            printf("Child %d terminated with exit status %d",
                wpid, WEXITSTATUS(child_status));
        else
            printf("Child %d terminated abnormally", wpid);
    }
}
```

class16.ppt - 31 - CS 213 F'01

Wait/Waitpid Example Outputs

Using wait (fork9)

```
Child 3565 terminated with exit status 103
Child 3564 terminated with exit status 102
Child 3563 terminated with exit status 101
Child 3562 terminated with exit status 100
Child 3566 terminated with exit status 104
```

Using waitpid (fork10)

```
Child 3568 terminated with exit status 100
Child 3569 terminated with exit status 101
Child 3570 terminated with exit status 102
Child 3571 terminated with exit status 103
Child 3572 terminated with exit status 104
```

class16.ppt - 32 - CS 213 F'01

exec: Running new programs

```
int execl(char *path, char *arg0, char *arg1, ..., 0)

• loads and runs executable at path with args arg0, arg1, ...
  -path is the complete path of an executable
  -arg0 becomes the name of the process
    » typically arg0 is either identical to path, or else it contains only the
      executable filename from path
  -"real" arguments to the executable start with arg1, etc.
  -list of args is terminated by a (char *)0 argument

• returns -1 if error, otherwise doesn't return!
```

```
main() {
    if (fork() == 0) {
        execl("/usr/bin/cp", "cp", "foo", "bar", 0);
    }
    wait(NULL);
    printf("copy completed\n");
    exit();
}
```

class16.ppt - 33 - CS 213 F'01

Summarizing

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, though
- Each process appears to have total control of processor + private memory space

Spawning Processes

- Call to fork
- One call, two returns

Terminating Processes

- Call exit

Reaping Processes

- Call wait or waitpid
- class16.ppt - 34 - CS 213 F'01