

15-213
"The course that gives CMU its Zip!"
Exceptional Control Flow
Part II
October 12, 2007

Topics

- Process Hierarchy
- Shells
- Signals
- Nonlocal jumps

lecture-14.ppt

ECF Exists at All Levels of a System

<p>Exceptions</p> <ul style="list-style-type: none"> ■ Hardware and operating system kernel software <p>Concurrent processes</p> <ul style="list-style-type: none"> ■ Hardware timer and kernel software <p>Signals</p> <ul style="list-style-type: none"> ■ Kernel software <p>Non-local jumps</p> <ul style="list-style-type: none"> ■ Application code 	<p>} Previous Lecture</p> <p>} This Lecture</p>
---	---

- 2 - 15-213, F07

The World of Multitasking

System Runs Many Processes Concurrently

- Process: executing program
 - State consists of memory image + register values + program counter
- Continually switches from one process to another
 - Suspend process when it needs I/O resource or timer event occurs
 - Resume process when I/O available or given scheduling priority
- Appears to user(s) as if all processes executing simultaneously
 - Even though most systems can only execute one process at a time
 - Except possibly with lower performance than if running alone

- 3 - 15-213, F07

Programmer's Model of Multitasking

Basic Functions

- `fork()` spawns new process
 - Called once, returns twice
- `exit()` terminates own process
 - Called once, never returns
 - Puts it into "zombie" status
- `wait()` and `waitpid()` wait for and reap terminated children
- `exec1()` and `execve()` run a new program in an existing process
 - Called once, (normally) never returns

Programming Challenge

- Understanding the nonstandard semantics of the functions
- Avoiding improper use of system resources
 - E.g. "Fork bombs" can disable a system

- 4 - 15-213, F07

wait: Synchronizing with Children

```
int wait(int *child_status)
```

- suspends current process until one of its children terminates
- return value is 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

- 5 - 15-213, F07

wait: Synchronizing with Children

```
void fork9() {
    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 terminated\n");
    }
    printf("Bye\n");
    exit();
}
```

```

sequenceDiagram
    participant HP as HP
    participant HC as HC
    Note over HP: HP: hello from parent
    HP-->>HC: fork()
    Note over HC: HC: hello from child
    Note over HC: HC terminates
    Note over HP: HP: wait()
    Note over HP: HP: CT: child has terminated
    Note over HP: HP: Bye
  
```

- 6 - 15-213, F07

wait() Example

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

```
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 = wait(&child_status);
        if (WIFEXITED(child_status))
            printf("Child %d terminated with exit status %d\n",
                wpid, WEXITSTATUS(child_status));
        else
            printf("Child %d terminate abnormally\n", wpid);
    }
}
```

waitpid(): Waiting for a Specific Process

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

```
void fork11()
{
    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\n",
                wpid, WEXITSTATUS(child_status));
        else
            printf("Child %d terminated abnormally\n", wpid);
    }
}
```

-8-

15-213, F07

exec: Loading and Running 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 object file
 - By convention, arg0 is the name of the executable object file
 - "Real" arguments to the program start with arg1, etc.
 - List of args is terminated by a (char *)0 argument
 - Environment taken from char **environ, which points to an array of "name=value" strings:
 - » USER=droh
 - » LOGNAME=droh
 - » HOME=/afs/cs.cmu.edu/user/droh
- Returns -1 if error, otherwise doesn't return!
- Family of functions includes `execv`, `execve` (base function), `execvp`, `execl`, `execle`, and `execlp`

-9-

15-213, F07

exec: Loading and Running Programs

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

-10-

15-213, F07

Shell Programs

A **shell** is an application program that runs programs on behalf of the user.

- sh - Original Unix Bourne Shell
- csh - BSD Unix C Shell, tcsh - Enhanced C Shell
- bash - Bourne-Again Shell

```
int main()
{
    char cmdline[MAXLINE];
    while (1) {
        /* read */
        printf("> ");
        fgets(cmdline, MAXLINE, stdin);
        if (feof(stdin))
            exit(0);
        /* evaluate */
        eval(cmdline);
    }
}
```

Execution is a sequence of read/evaluate steps

-11-

15-213, F07

Simple Shell eval Function

```
void eval(char *cmdline)
{
    char *argv[MAXARGS]; /* argv for execve() */
    int bg; /* should the job run in bg or fg? */
    pid_t pid; /* process id */

    bg = parseline(cmdline, argv);
    if (!builtin_command(argv)) {
        if ((pid = Fork()) == 0) { /* child runs user job */
            if (execve(argv[0], argv, environ) < 0) {
                printf("%s: Command not found.\n", argv[0]);
                exit(0);
            }
        }
        if (!bg) { /* parent waits for fg job to terminate */
            int status;
            if (waitpid(pid, &status, 0) < 0)
                unix_error("waitpid: waitpid error");
        }
        else /* otherwise, don't wait for bg job */
            printf("%d %s", pid, cmdline);
    }
}
```

-12-

15-213, F07

Problem with Simple Shell Example

Shell correctly waits for and reaps foreground jobs.

But, what about background jobs?

- Will become zombies when they terminate
- Will never be reaped because shell (typically) will not terminate
- Creates a memory leak that will eventually crash the kernel when it runs out of memory

Solution: Reaping background jobs requires a mechanism called a **signal**

- 13 -

15-213, F07

Signals

A **signal** is a small message that notifies a process that an event of some type has occurred in the system.

- akin to exceptions and interrupts
- sent from the kernel (sometimes at the request of another process) to a process
- signal type is identified by small integer ID's (1-30)
- the only information in a signal is its ID and the fact that it arrived

ID	Name	Default Action	Corresponding Event
2	SIGINT	Terminate	Interrupt (e.g., <code>ctrl-c</code> from keyboard)
9	SIGKILL	Terminate	Kill program (cannot override or ignore)
11	SIGSEGV	Terminate & Dump	Segmentation violation
14	SIGALRM	Terminate	Timer signal
17	SIGCHLD	Ignore	Child stopped or terminated

- 14 -

15-213, F07

Signal Concepts

Sending a signal

- Kernel **sends** (delivers) a signal to a **destination process** by updating some state in the context of the destination process.
- Kernel sends a signal for one of the following reasons:
 - Kernel has detected a system event such as divide-by-zero (SIGFPE) or the termination of a child process (SIGCHLD)
 - Another process has invoked the `kill` system call to explicitly request the kernel to send a signal to the destination process.

- 15 -

15-213, F07

Signal Concepts (continued)

Receiving a signal

- A destination process **receives** a signal when it is forced by the kernel to react in some way to the delivery of the signal.
- Three possible ways to react:
 - Ignore the signal (do nothing)
 - Terminate the process (with optional core dump).
 - **Catch** the signal by executing a user-level function called a **signal handler**.
 - » Akin to a hardware exception handler being called in response to an asynchronous interrupt.

- 16 -

15-213, F07

Signal Concepts (continued)

A signal is **pending** if it has been sent but not yet received.

- There can be at most one pending signal of any particular type.
- Important: Signals are not queued
 - If a process has a pending signal of type `k`, then subsequent signals of type `k` that are sent to that process are discarded.

A process can **block** the receipt of certain signals.

- Blocked signals can be delivered, but will not be received until the signal is unblocked.

A pending signal is received at most once.

- 17 -

15-213, F07

Signal Concepts

Kernel maintains **pending** and **blocked** bit vectors in the context of each process.

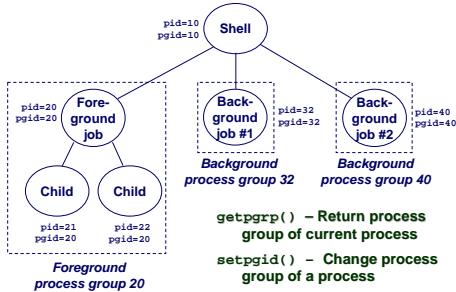
- **pending** – represents the set of pending signals
 - Kernel sets bit `k` in **pending** whenever a signal of type `k` is delivered.
 - Kernel clears bit `k` in **pending** whenever a signal of type `k` is received
- **blocked** – represents the set of blocked signals
 - Can be set and cleared by the application using the `sigprocmask` function.

- 18 -

15-213, F07

Process Groups

Every process belongs to exactly one process group



- 19 -

15-213, F07

Sending Signals with kill Program

kill program sends arbitrary signal to a process or process group

Examples

- kill -9 24818
 - Send SIGKILL to process 24818
- kill -9 -24817
 - Send SIGKILL to every process in process group 24817.

```
linux> ./forks 16
linux> Child1: pid=24818 pgrp=24817
Child2: pid=24819 pgrp=24817

linux> ps
  PID TTY          TIME CMD
 24788 pts/2    00:00:00 tcsh
 24818 pts/2    00:00:02 forks
 24819 pts/2    00:00:02 forks
 24820 pts/2    00:00:00 ps
linux> kill -9 -24817
linux> ps
  PID TTY          TIME CMD
 24788 pts/2    00:00:00 tcsh
 24823 pts/2    00:00:00 ps
linux>
```


- 20 -

15-213, F07

Sending Signals from the Keyboard

Typing ctrl-c (ctrl-z) sends a SIGINT (SIGTSTP) to every job in the foreground process group.

- SIGINT - default action is to terminate each process
- SIGTSTP - default action is to stop (suspend) each process



- 21 -

15-213, F07

Example of ctrl-c and ctrl-z

```
bluefish> ./forks 17
Child: pid=28108 pgrp=28107
Parent: pid=28107 pgrp=28107
<types ctrl-z>
Suspended
bluefish> ps w
  PID TTY          STAT      TIME COMMAND
 27699 pts/8      Ss         0:00 -tcsh
 28107 pts/8      T          0:01 ./Forks 17
 28108 pts/8      T          0:01 ./Forks 17
 28109 pts/8      R+         0:00 ps w
bluefish> fg
./forks 17
<types ctrl-c>
bluefish> ps w
  PID TTY          STAT      TIME COMMAND
 27699 pts/8      Ss         0:00 -tcsh
 28110 pts/8      R+         0:00 ps w
```

STAT (process state)
Legend:

First letter:
S: sleeping
T: stopped
R: running

Second letter:
s: session leader
+: foreground proc group

See "man ps" for more details

- 22 -

15-213, F07

Sending Signals with kill Function

```
void fork12()
{
    pid_t pid[N];
    int i, child_status;
    for (i = 0; i < N; i++)
        if ((pid[i] = fork()) == 0)
            while(1); /* Child infinite loop */

    /* Parent terminates the child processes */
    for (i = 0; i < N; i++) {
        printf("Killing process %d\n", pid[i]);
        kill(pid[i], SIGINT);
    }

    /* Parent reaps terminated children */
    for (i = 0; i < N; i++) {
        pid_t wpid = wait(&child_status);
        if (WIFEXITED(child_status))
            printf("Child %d terminated with exit status %d\n",
                wpid, WEXITSTATUS(child_status));
        else
            printf("Child %d terminated abnormally\n", wpid);
    }
}
```

- 23 -

15-213, F07

Receiving Signals

Suppose kernel is returning from an exception handler and is ready to pass control to process *p*.

Kernel computes *pnb* = pending & ~blocked

- The set of pending nonblocked signals for process *p*

If (*pnb* == 0)

- Pass control to next instruction in the logical flow for *p*.

Else

- Choose least nonzero bit *k* in *pnb* and force process *p* to receive signal *k*.
- The receipt of the signal triggers some *action* by *p*
- Repeat for all nonzero *k* in *pnb*.
- Pass control to next instruction in logical flow for *p*.

- 24 -

15-213, F07

Default Actions

Each signal type has a predefined **default action**, which is one of:

- The process terminates
- The process terminates and dumps core.
- The process stops until restarted by a SIGCONT signal.
- The process ignores the signal.

- 25 -

15-213, F07

Installing Signal Handlers

The `signal` function modifies the default action associated with the receipt of signal `signum`:

- `handler_t *signal(int signum, handler_t *handler)`

Different values for `handler`:

- `SIG_IGN`: ignore signals of type `signum`
- `SIG_DFL`: revert to the default action on receipt of signals of type `signum`.
- Otherwise, `handler` is the address of a **signal handler**
 - Called when process receives signal of type `signum`
 - Referred to as "**installing**" the handler.
 - Executing handler is called "**catching**" or "**handling**" the signal.
 - When the handler executes its return statement, control passes back to instruction in the control flow of the process that was interrupted by receipt of the signal.

- 26 -

15-213, F07

Signal Handling Example

```
void int_handler(int sig)
{
    printf("Process %d received signal %d\n",
        getpid(), sig);
    exit(0);
}

void fork13()
{
    pid_t pid[N];
    int i, child_status;
    signal(SIGINT, int_handler);
    ...
}
```


```
linux> ./forks 13
Killing process 24973
Killing process 24974
Killing process 24975
Killing process 24976
Killing process 24977
Process 24977 received signal 2
Child 24977 terminated with exit status 0
Process 24976 received signal 2
Child 24976 terminated with exit status 0
Process 24975 received signal 2
Child 24975 terminated with exit status 0
Process 24974 received signal 2
Child 24974 terminated with exit status 0
Process 24973 received signal 2
Child 24973 terminated with exit status 0
linux>
```

- 27 -

15-213, F07

Signals Handlers as Concurrent Flows


A signal handler is a separate logical flow (thread) that runs concurrently with the main program.



- 28 -

15-213, F07

Another View of Signal Handlers as Concurrent Flows



- 29 -

15-213, F07

Nonlocal Jumps: `setjmp`/`longjmp`

Powerful (but dangerous) user-level mechanism for transferring control to an arbitrary location.

- Controlled to way to break the procedure call / return discipline
- Useful for error recovery and signal handling

```
int setjmp(jmp_buf j)
```

- Must be called before `longjmp`
- Identifies a return site for a subsequent `longjmp`.
- Called once, returns one or more times

Implementation:

- Remember where you are by storing the current register context, stack pointer, and PC value in `jmp_buf`.
- Return 0

- 30 -

15-213, F07

setjmp/longjmp (cont)

```
void longjmp(jmp_buf j, int i)
```

- Meaning:
 - return from the setjmp remembered by jump buffer j again...
 - ...this time returning i instead of 0
- Called after setjmp
- Called once, but never returns

longjmp Implementation:

- Restore register context from jump buffer j
- Set %eax (the return value) to i
- Jump to the location indicated by the PC stored in jump buf j.

- 31 -

15-213, F07

setjmp/longjmp Example

```
#include <setjmp.h>
jmp_buf buf;

main() {
    if (setjmp(buf) != 0) {
        printf("back in main due to error\n");
    } else {
        printf("first time through\n");
        p1(); /* p1 calls p2, which calls p3 */
    }
    ...
    p3() {
        <error checking code>
        if (error)
            longjmp(buf, 1)
    }
}
```

- 32 -

15-213, F07

Limitations of Nonlocal Jumps

Works within stack discipline

- Can only long jump to environment of function that has been called but not yet completed

```
jmp_buf env;

P1()
{
    if (setjmp(env)) {
        /* Long Jump to here */
    } else {
        P2();
    }
}

P2()
{
    . . . P2(); . . . P3();
}

P3()
{
    longjmp(env, 1);
}
```

- 33 -

15-213, F07

Limitations of Long Jumps (cont.)

Works within stack discipline

- Can only long jump to environment of function that has been called but not yet completed

```
jmp_buf env;

P1()
{
    P2(); P3();
}

P2()
{
    if (setjmp(env)) {
        /* Long Jump to here */
    }
}

P3()
{
    longjmp(env, 1);
}
```

- 34 -

15-213, F07

Putting It All Together: A Program That Restarts Itself When ctrl-c'd

```
#include <stdio.h>
#include <signal.h>
#include <setjmp.h>

sigjmp_buf buf;

void handler(int sig) {
    siglongjmp(buf, 1);
}

main() {
    signal(SIGINT, handler);
    if (!sigsetjmp(buf, 1))
        printf("starting\n");
    else
        printf("restarting\n");
}
```

```
while(1) {
    sleep(1);
    printf("processing...\n");
}
```

```
base> a.out
starting
processing...
processing...
restarting
processing...
processing...
restarting
processing...
```

- 35 -

15-213, F07

Summary

Signals provide process-level exception handling

- Can generate from user programs
- Can define effect by declaring signal handler

Some caveats

- Very high overhead
 - >10,000 clock cycles
 - Only use for exceptional conditions
- Don't have queues
 - Just one bit for each pending signal type

Nonlocal jumps provide exceptional control flow within process

- Within constraints of stack discipline

- 36 -

15-213, F07