

Figure 1. Fork system call. With this original illustration, Melvin Conway envisioned "fork and join" system calls in 1965.

1. from handout week2.a

1.a C code

```

...
31 uint64_t f(uint64_t *ptr)
32 {
33     uint64_t x = 0;
34     x = g(*ptr);
35     return x + 1;
36 }
37
38 uint64_t g(uint64_t a)
39 {
40     uint64_t x = 2*a;
41     q = &x; // <--- THIS IS AN ERROR (AKA BUG)
42     return x;
43 }
...
    
```


① 8B } 16B why 32B for f?
 ② 2(*ptr)+1

1.b assembly code (by "gcc -O0")

```

...
28 f:
29     pushq %rbp          # prologue: store caller's frame pointer
30     movq %rsp, %rbp    # prologue: set frame pointer for new frame
31
32     subq $32, %rsp     # make stack space
33     movq %rdi, -24(%rbp) # Move ptr to the stack
34     movq %0, -8(%rbp)  # ptr now lives at rbp - 24
35     movq %0, -8(%rbp)  # x = 0 (x's address is rbp - 8)
36
37     movq -24(%rbp), %r8 # move 'ptr' to %r8
38     movq (%r8), %r9    # dereference 'ptr' and save value to %r9
39     movq %r9, %rdi     # Move the value of *ptr to rdi,
40                       # so we can call g
41
42     call g             # invoke g
43
44     movq %rax, -8(%rbp) # x = (return value of g)
45     movq -8(%rbp), %r10 # compute x + 1, part I
46     addq $1, %r10      # compute x + 1, part II
47     movq %r10, %rax    # Get ready to return x + 1
48
49     movq %rbp, %rsp    # epilogue: undo stack frame
50     popq %rbp         # epilogue: restore frame pointer from caller
51     ret               # return
...
    
```

"no opt"




2. "gcc -O3 -S example.c"

```

...
subq $24, %rsp          // push frame: allocate stack frame 24B
                        // [Security]
                        // canary value (%fs:40) for detecting stack
                        // smashing attacks
                        // put canary at 8(%rsp)
movq %fs:40, %rax
movq %rax, 8(%rsp)
xorl %eax, %eax        // %eax=0 [%eax is the low 32bit of %rax]
movq %rsp, %rdx        // %rdx = %rsp
movq (%rdi), %rax      // %rax = *ptr (%rdi contains the first arg to
                        // function f, which is "ptr")
movq %rdx, q(%rip)    // "q" is the global variable; set "q" to %rdx:

                        // [Security]
                        // copy the canary at stack
                        // check if the value has been changed
                        // if so, alert!!!
movq 8(%rsp), %rcx
xorq %fs:40, %rcx
jne .L9
leaq 1(%rax,%rax), %rax // %rax = %rax + %rax + 1
                        // (!!! this is function g!)
addq $24, %rsp         // pop frame
ret                   // return to main
...
    
```

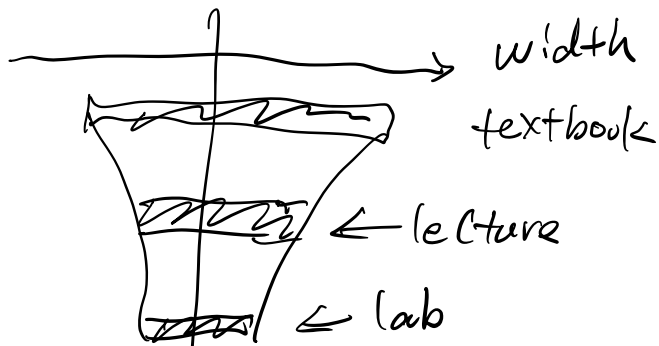


② no "call g"


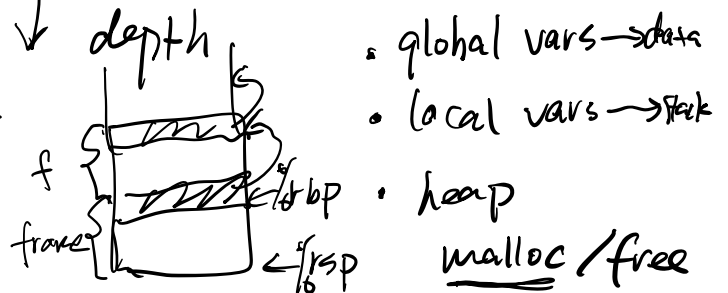
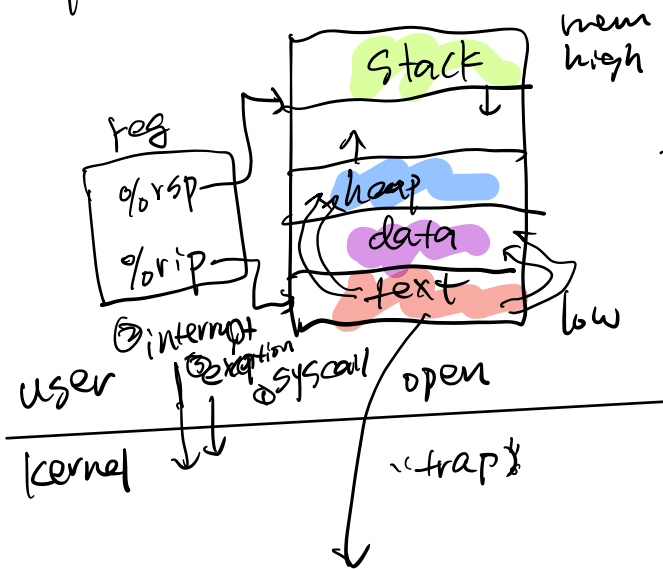
ret = 2 * %rax + 1

0. Admin
1. Last time
2. Process birth
3. Shell, crash course
4. Shell internals I
5. File descriptors
6. Shell internals II

textbook



process



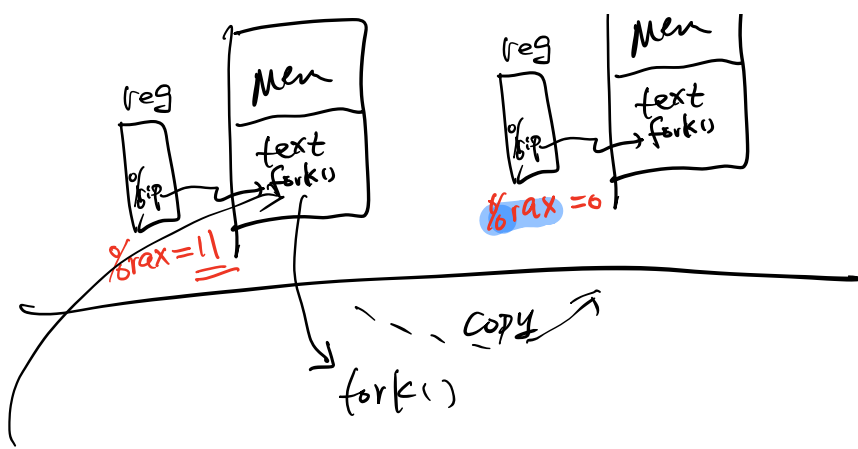
Calling Conventions

process birth

syscall: fork()

process 10 (Parent) process 11 (child)

only one diff



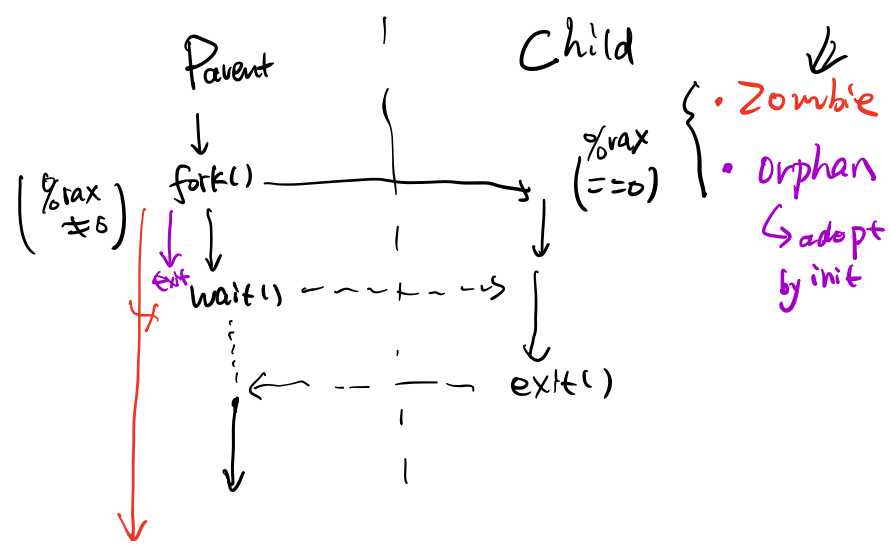
Create Process(...)

↳ fork/exec separation

```

if (fork()
    == 0) {
    // do A ← process 1
} else {
    // do B ← process 0
}

```




```

$ ./hellworld &
└─┬─ shell


```

File descriptor

int
0 : Std in ←
1 : Std out ← printf
2 : Std err



```
main() {  
    get_input() ←  
    read(0, ...) ←  
    printf(...) }  
    write(1, ...)
```



Crash course, Shell

1. run cmd
"\$ ls" and "\$ ls -a"

shell
fork
"ls"

2. output redirection
"\$ ls" prints to screen
"\$ ls > files.txt"

3. backgrounding
"\$ web-server &
\$ "

bash
zshell

4. pipe
-- "\$ cat students.txt | shuf -n 1"
-- equivalent to
"\$ cat students.txt > /tmp/tmpfile
\$ shuf -n 1 /tmp/tmpfile
\$ rm /tmp/tmpfile"

5. Shell builtin cmds vs. program
-- "echo/pwd/which" vs. "ls"
-- use "which" to tell
program: "\$ which ls" => "/bin/ls"
built-in: "\$ which which" => "which: shell
built-in command"

"cd"

```

1 CS5600
2 Handout week03a
3
4 The handout is meant to:
5
6 --illustrate how the shell itself uses syscalls
7
8 --communicate the power of the fork()/exec() separation
9
10 --give an example of how small, modular pieces (file descriptors,
11 pipes, fork(), exec()) can be combined to achieve complex behavior
12 far beyond what any single application designer could or would have
13 specified at design time.
14

```

1. Pseudocode for a very simple shell

```

15 while (1) {
16     write(1, "$ ", 2); // $
17     readcommand(command, args); // parse input
18     if ((pid = fork()) == 0) { // child?
19         execve(command, args, 0); // write (arg1, ...)
20     } else if (pid > 0) { // parent?
21         wait(0); // wait for child
22     } else {
23         perror("failed to fork");
24     }
25 }

```

\$./ls -ld
write (arg1, ...)
file descript

2. Now add two features to this simple shell: output redirection and backgrounding

By output redirection, we mean, for example:
\$ ls > list.txt
 By backgrounding, we mean, for example:
\$ myprog &

\$./ls > list.txt &
fd table

0	/dev/tty
1	/dev/tty → list.txt
2	/dev/tty

```

37 while (1) {
38     write(1, "$ ", 2);
39     readcommand(command, args); // parse input
40     if ((pid = fork()) == 0) { // child?
41         if (output_redirected) {
42             close(1);
43             open(redirect_file, O_CREAT | O_TRUNC | O_WRONLY, 0666);
44         }
45         // when command runs, fd 1 will refer to the redirected file
46         execve(command, args, 0);
47     } else if (pid > 0) { // parent?
48         if (foreground_process) {
49             wait(0); //wait for child
50         }
51     } else {
52         perror("failed to fork");
53     }
54 }

```

create process
fork()
 ...
open
exec()

3. Another syscall example: pipe()

The pipe() syscall is used by the shell to implement pipelines, such as
\$ ls | sort | head -4
 We will see this in a moment; for now, here is an example use of pipes.

```

61 // C fragment with simple use of pipes
62
63 int fdarray[2];
64 char buf[512];
65 int n;
66
67 pipe(fdarray);
68 write(fdarray[1], "hello", 5);
69 n = read(fdarray[0], buf, sizeof(buf));
70 // buf[] now contains 'h', 'e', 'l', 'l', 'o'
71

```

4. File descriptors are inherited across fork

```

72 // C fragment showing how two processes can communicate over a pipe
73
74 int fdarray[2];
75 char buf[512];
76 int n, pid;
77
78 pipe(fdarray);
79 pid = fork();
80 if (pid > 0) {
81     write(fdarray[1], "hello", 5);
82 } else {
83     n = read(fdarray[0], buf, sizeof(buf));
84 }
85

```

fork + exec
posix_spawn()
 ✓ fork() / ✓ vfork()

```

90 5. Putting it all together: implementing shell pipelines using
91 fork(), exec(), and pipe().
92
93
94 // Pseudocode for a Unix shell that can run processes in the
95 // background, redirect the output of commands, and implement
96 // two element pipelines, such as "ls | sort"
97
98 void main_loop() {
99
100     while (1) {
101         write(1, "$ ", 2);
102         readcommand(command, args); // parse input
103         if ((pid = fork()) == 0) { // child?
104             if (pipeline_requested) {
105                 handle_pipeline(left_command, right_command)
106             } else {
107                 if (output_redirected) {
108                     close(1);
109                     open(redirect_file, O_CREAT | O_TRUNC | O_WRONLY, 0666);
110                 }
111                 exec(command, args, 0);
112             }
113         } else if (pid > 0) { // parent?
114             if (foreground_process) {
115                 wait(0); // wait for child
116             }
117         } else {
118             perror("failed to fork");
119         }
120     }
121 }
122
123 void handle_pipeline(left_command, right_command) {
124
125     int fdarray[2];
126
127     if (pipe(fdarray) < 0) panic("error");
128     if ((pid = fork ()) == 0) { // child (left end of pipe)
129
130         dup2 (fdarray[1], 1); // make fd 1 the same as fdarray[1],
131                             // which is the write end of the
132                             // pipe. implies close (1).
133
134         close (fdarray[0]);
135         close (fdarray[1]);
136         parse(command1, args1, left_command);
137         exec (command1, args1, 0);
138     } else if (pid > 0) { // parent (right end of pipe)
139
140         dup2 (fdarray[0], 0); // make fd 0 the same as fdarray[0],
141                             // which is the read end of the pipe.
142                             // implies close (0).
143
144         close (fdarray[0]);
145         close (fdarray[1]);
146         parse(command2, args2, right_command);
147         exec (command2, args2, 0);
148     } else {
149         printf ("Unable to fork\n");
150     }
151 }

```

```

151
152
153 6. Commentary
154
155 Why is this interesting? Because pipelines and output redirection
156 are accomplished by manipulating the child's environment, not by
157 asking a program author to implement a complex set of behaviors.
158 That is, the identical code* for "ls" can result in printing to the
159 screen ("ls -l"), writing to a file ("ls -l > output.txt"), or
160 getting ls's output formatted by a sorting program ("ls -l | sort").
161
162 This concept is powerful indeed. Consider what would be needed if it
163 weren't for redirection: the author of ls would have had to
164 anticipate every possible output mode and would have had to build in
165 an interface by which the user could specify exactly how the output
166 is treated.
167
168 What makes it work is that the author of ls expressed their
169 code in terms of a file descriptor:
170 write(1, "some output", byte_count);
171 This author does not, and cannot, know what the file descriptor will
172 represent at runtime. Meanwhile, the shell has the opportunity, *in
173 between fork() and exec(*), to arrange to have that file descriptor
174 represent a pipe, a file to write to, the console, etc.

```