

1. pipelines
  2. what makes a good abstraction?
  3. process memory layout
  4. Crash course in x86-64 assembly
  5. Stack frames
- 

fork/exec  
 wait/exit

Q: fork() ret val?

Q: fd 0/1/2?

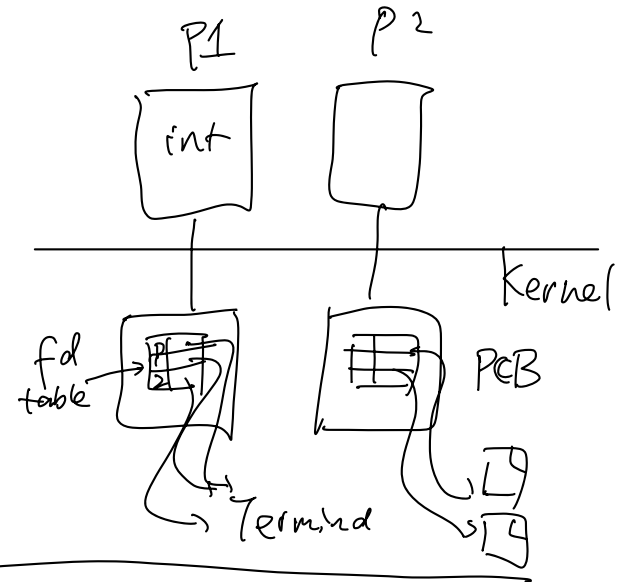
Q: "ls > log.txt"?

0: std in  
 1: stdout  
 2: std err

"cat"

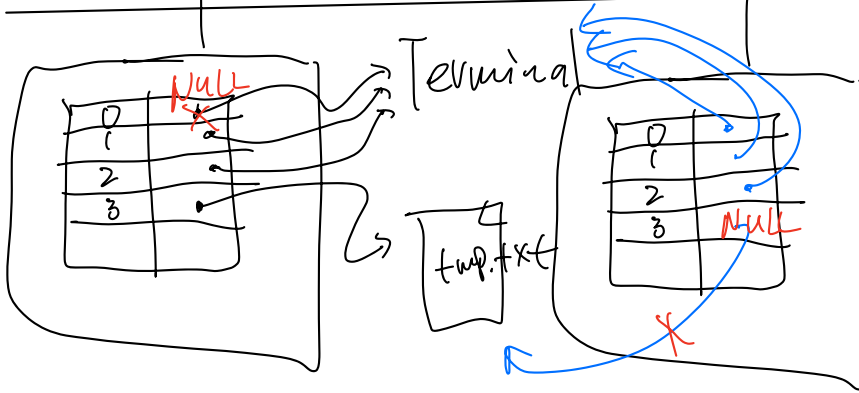


~~Terminal~~



```
main {
  int fd = 3;
}
```

```
main {
  int fd = 3;
}
```



display:

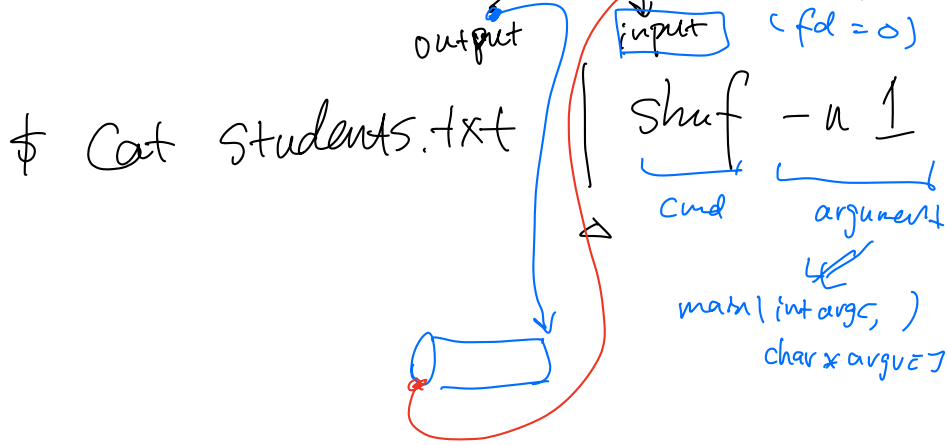
[child] fd = 3

[parent] fd = 3

file:



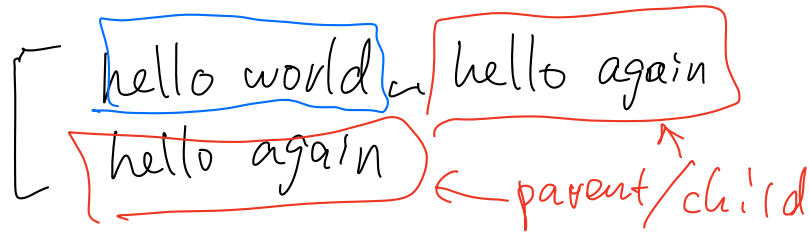
• pipe




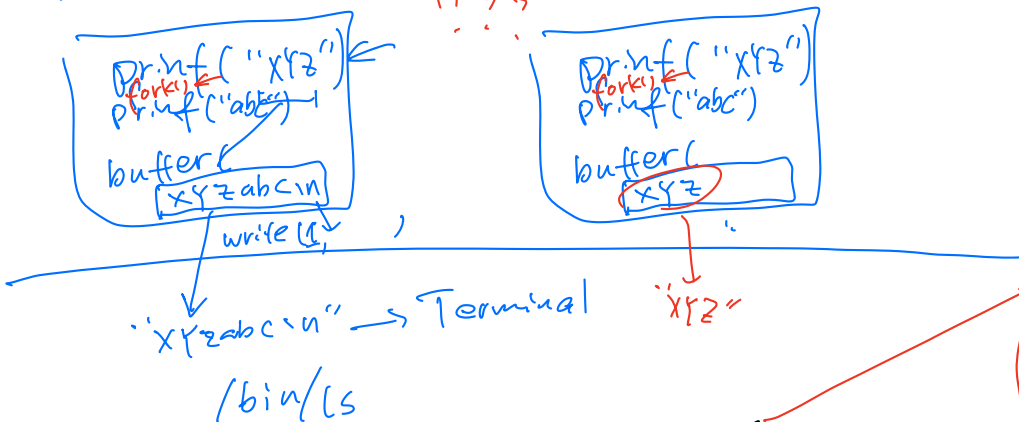

"A fork() in the road" (HotOS'19)

"Fork today is a convenient API for a single-threaded process with a small memory footprint and simple memory layout that requires fine-grained control over the execution environment of its children but does not need to be strongly isolated from them. In other words, a shell."

parent only




Truth ⇒



Compile



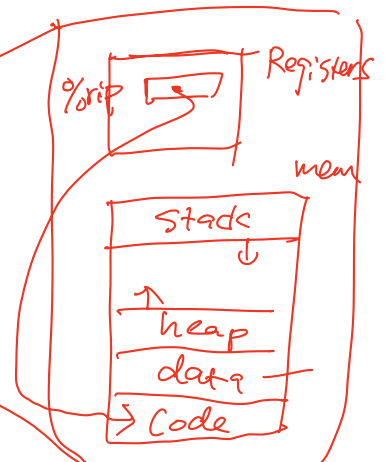
exec



← shell  
fork/exec/...

← process  
mem layout

← NEXT →



```

1 CS3650 24spring
2 Handout week.04a
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 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.

```

#### 1. Pseudocode for a very simple shell

```

16 while (1) {
17     write(1, "$ ", 2);
18     readcommand(command, args); // parse input
19     if ((pid = fork()) == 0) { // child?
20         execve(command, args, 0);
21     } else if (pid > 0) { // parent?
22         wait(0); //wait for child
23     } else {
24         perror("failed to fork");
25     }
26 }

```

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

```

29 By output redirection, we mean, for example:
30 $ ls > list.txt
31
32 while (1) {
33     write(1, "$ ", 2);
34     readcommand(command, args); // parse input
35     if ((pid = fork()) == 0) { // child?
36         if (output_redirected) {
37             close(1);
38             open(redirect_file, O_CREAT | O_TRUNC | O_WRONLY, 0666);
39         }
40         // when command runs, fd 1 will refer to the redirected file
41         execve(command, args, 0);
42     } else if (pid > 0) { // parent?
43         wait(0); //wait for child
44     } else {
45         perror("failed to fork");
46     }
47 }
48 }
49

```

x86-32

#### 3. Another syscall example: pipe()

```

52 The pipe() syscall is used by the shell to implement pipelines, such as
53 $ ls | sort | head -4
54 We will see this in a moment; for now, here is an example use of
55 pipes.

```

```

56 // C fragment with simple use of pipes
57
58
59 int fdarray[2];
60 char buf[512];
61 int n;

```

```

62 pipe(fdarray);
63 write(fdarray[1], "hello", 5);
64 n = read(fdarray[0], buf, sizeof(buf));
65 // buf[] now contains 'h', 'e', 'l', 'l', 'o'
66
67
68
69

```

#### 4. File descriptors are inherited across fork

```

70 // C fragment showing how two processes can communicate over a pipe
71
72
73 int fdarray[2];
74 char buf[512];
75 int n, pid;

```

```

76 pipe(fdarray);
77 pid = fork();
78 if (pid > 0) {
79     write(fdarray[1], "hello", 5);
80 } else {
81     n = read(fdarray[0], buf, sizeof(buf));
82 }
83
84

```

#### 5. Commentary

86 Why is this interesting? Because pipelines and output redirection  
87 are accomplished by manipulating the child's environment, not by  
88 asking a program author to implement a complex set of behaviors.  
89 That is, the \*identical code\* for "ls" can result in printing to the  
90 screen ("ls -l"), writing to a file ("ls -l > output.txt"), or  
91 getting ls's output formatted by a sorting program ("ls -l | sort").  
92

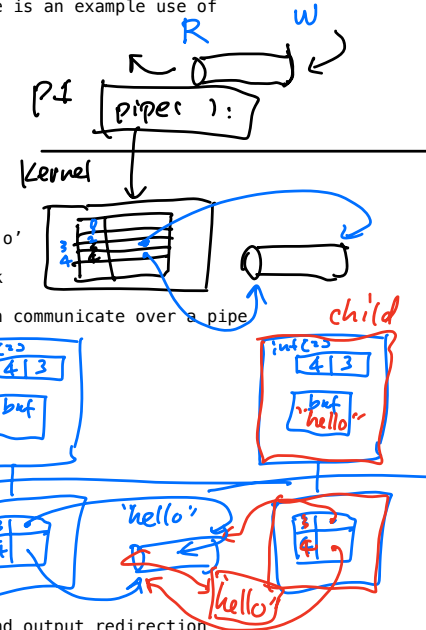
93 This concept is powerful indeed. Consider what would be needed if it  
94 weren't for redirection: the author of ls would have had to  
95 anticipate every possible output mode and would have had to build in  
96 an interface by which the user could specify exactly how the output  
97 is treated.  
98

99 What makes it work is that the author of ls expressed their  
100 code in terms of a file descriptor:

```

101 write(1, "some output", byte_count);
102 This author does not, and cannot, know what the file descriptor will
103 represent at runtime. Meanwhile, the shell has the opportunity, *in
104 between fork() and exec()* , to arrange to have that file descriptor
105 represent a pipe, a file to write to, the console, etc.

```



bytes

# Crash course of x86-64 assembly

ISA CPU

- x86: Intel, AMD
- ARM: Apple
- RISC-V:

opCode      operands


\* movq PLACE1, PLACE2

← register  
← memory  
← immediate

\* pushq %rax [ subq \$8, %rsp  
                  "            "  
                  123            movq %rax, (%rsp) ]

← 8 bytes

\* popq %rax [ movq (%rsp), %rax  
                  addq \$8, %rsp ]



\* call 0x12345 [ pushq %rip  
                  movq \$0x12345, %rip ]

\* ret [ popq %rip ]

