

160 / 5

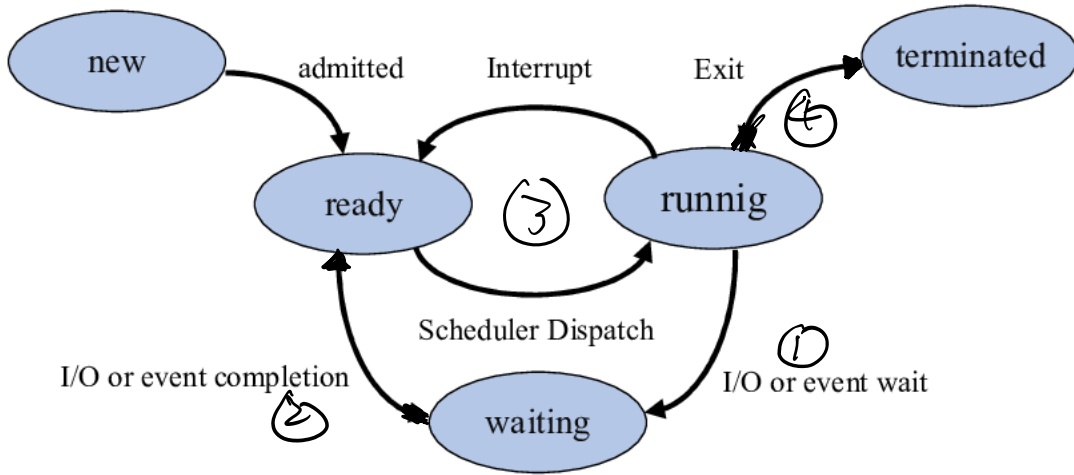


Fig13 from "Priority based round robin (PBRR) CPU scheduling algorithm"

Preemptive : (1) - (4)

non-preemptive : (1), (4)

```

36
37 // Per-process state
38 struct proc {
39     uint sz; // Size of process memory (bytes)
40     pde_t* pgdir; // Page table
41     char *kstack; // Bottom of kernel stack for this process
42     enum procstate state; // Process state
43     int pid; // Process ID
44     struct proc *parent; // Parent process
45     struct trapframe *tf; // Trap frame for current syscall
46     struct context *context; // switch() here to run process
47     void *chan; // If non-zero, sleeping on chan
48     int killed; // If non-zero, have been killed
49     struct file *file[NOFILE]; // Open files
50     struct inode *cwd; // Current directory
51     char name[16]; // Process name (debugging)
52 };
53

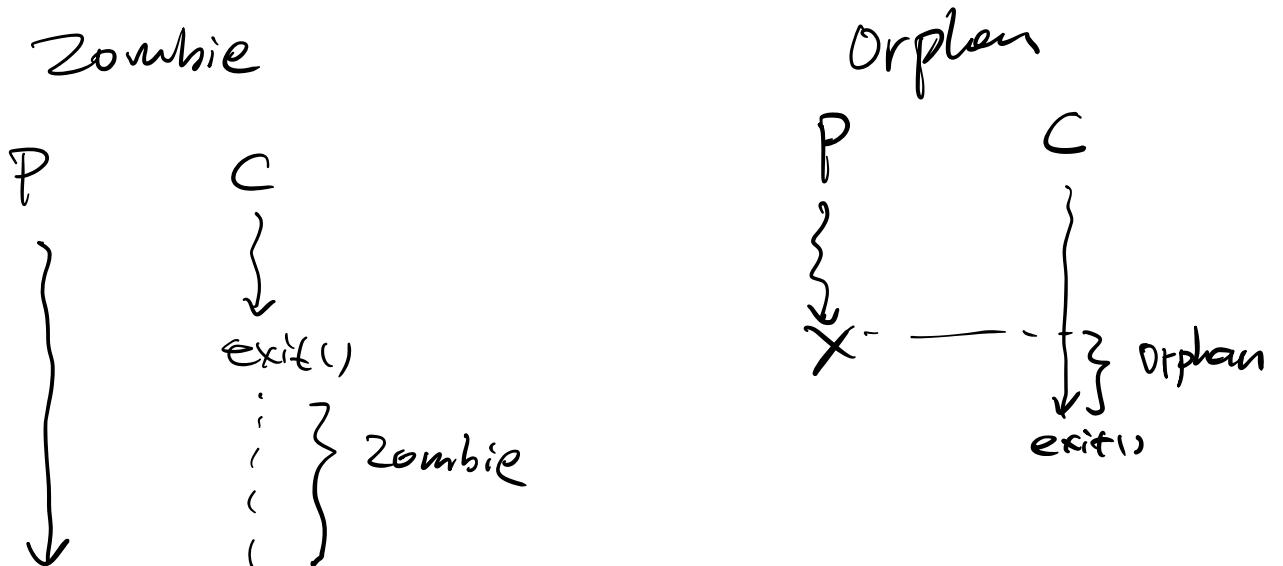
```

running/waiting/ready

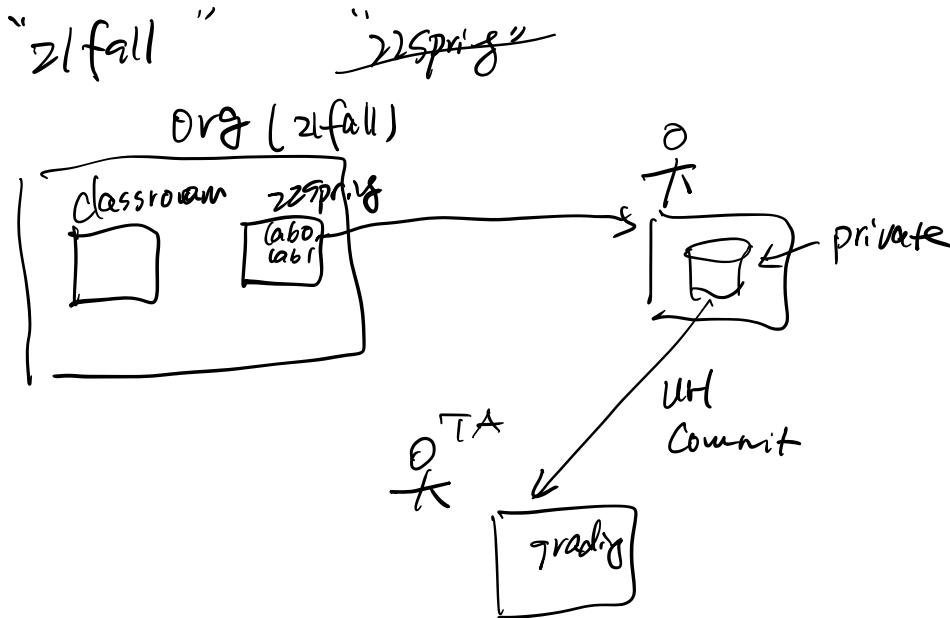
↓
zombie

Q: orphan?

Borrowed from xv6: <https://github.com/mit-pdos/xv6-public/blob/eeb7b415dbcb12cc362d0783e41c3d1f44066b17/proc.h>



1. Shell internal continued & discussions
2. Implementation of processes
3. Context switch intro
4. Scheduling intro
5. Scheduling disciplines (part I)
FIFO/SJF



Lab1


fork bomb

fork() / wait() / exec() execve()

file descriptor

0/1/2

TA →



ls > file.txt cat ~~grades.txt~~ | shuf -a 1

```

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
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 }
27

```

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 &
```

```

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             if (foreground_process) {
50                 wait(0); //wait for child
51             }
52         } else {
53             perror("failed to fork");
54         }
55     }
56 }

```

*fork/exec
separation*

```

56 3. Another syscall example: pipe()
57
58 The pipe() syscall is used by the shell to implement pipelines, such as
59 $ ls | sort | head -4
60 We will see this in a moment; for now, here is an example use of
61 pipes.
62

```

```

63 // C fragment with simple use of pipes
64
65 int fdarray[2];
66 char buf[512];
67 int n;
68
69 pipe(fdarray);
70 write(fdarray[1], "hello", 5);
71 n = read(fdarray[0], buf, sizeof(buf));
72 // buf[] now contains 'h', 'e', 'l', 'l', 'o'
73

```




4. File descriptors are inherited across fork

```

74
75 // C fragment showing how two processes can communicate over a pipe
76
77 // C fragment showing how two processes can communicate over a pipe
78 int fdarray[2];
79 char buf[512];
80 int n, pid;
81
82 pipe(fdarray);
83 pid = fork();
84 if (pid > 0) { // parent
85     write(fdarray[1], "hello", 5);
86 } else { // child
87     n = read(fdarray[0], buf, sizeof(buf));
88 }
89

```



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 void main_loop() {

```

98 while (1) {
99     write(1, "$ ", 2);
100     → readcommand(command, args); // parse input
101     if ((pid = fork()) == 0) { // child?
102         if (pipeline_requested) {
103             handle_pipeline(left_command, right_command)
104         } else {
105             if (output_redirected) {
106                 close(1);
107                 open(redirect_file, O_CREAT | O_TRUNC | O_WRONLY, 0666);
108             }
109             exec(command, args, 0);
110         }
111     } else if (pid > 0) { // parent?
112         if (foreground_process) {
113             wait(0); // wait for child
114         } else {
115             perror("failed to fork");
116         }
117     }
118 }
119 }
120 }
121 }

```



122 void handle_pipeline(left_command, right_command) {

```

123 int fdarray[2];
124
125 if (pipe(fdarray) < 0) panic("error");
126 if ((pid = fork()) == 0) { // child (left end of pipe)
127     dup2(fdarray[1], 1); // make fd 1 the same as fdarray[1],
128                          // which is the write end of the
129                          // pipe. implies close (1).
130     [ close(fdarray[0]);
131       close(fdarray[1]);
132       parse(command1, args1, left_command);
133       exec(command1, args1, 0);
134     ]
135 } else if (pid > 0) { // parent (right end of pipe)
136     dup2(fdarray[0], 0); // make fd 0 the same as fdarray[0],
137                          // which is the read end of the pipe.
138                          // implies close (0).
139     close(fdarray[0]);
140     close(fdarray[1]);
141     parse(command2, args2, right_command);
142     exec(command2, args2, 0);
143 } else {
144     printf("Unable to fork\n");
145 }
146 }
147 }
148 }
149 }
150 }

```

left | right



1 → [] → 0
 cat student.txt | shuf

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.

- * good abstraction
 - file descriptor
 - 0/1/2
 - fork/exec sep
- fork() today

cow

↗
"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."

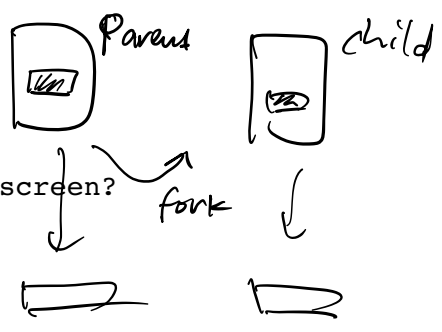
—"A fork() in the road"

<https://www.microsoft.com/en-us/research/uploads/prod/2019/04/fork-hotos19.pdf>

```

print("hello world");
fork();
print("\n"); ←

```



QUESTION: what do you expect to see on screen?

A: hello world\n
hello world\n

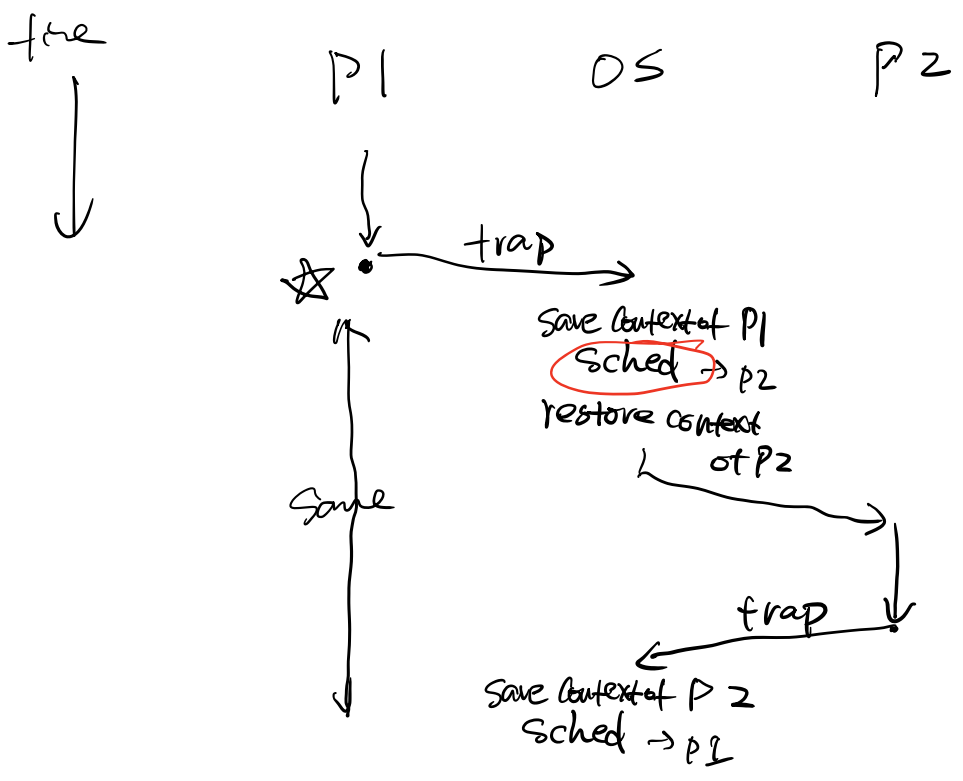


2

B: hello world\n
\n

1

onlinegdb.com

OS.



OS Scheduling



Scheduling game.

- 1 CPU
- multiple process: $P_1 P_2 \dots P_n$
 - arrival time
 - time to completion time



- turnaround time: $(c - a)$
- response time: $(1st - a)$
- throughput:
- fairness: no starvation
- Assumptions:
 - I/O
 - ignore context switch cost