

Jan 17, 13 15:49

I02-handout.txt

Page 1/4

```

1 Handout for CS 439
2 Class 2
3 17 January 2013
4
5 This handout is meant to:
6
7 --communicate the power of the fork()/exec() separation
8
9 --illustrate how the shell itself uses syscalls
10
11 --give an example of how small, modular pieces (file descriptors,
12 pipes, fork(), exec()) can be combined to achieve complex behavior
13 far beyond what any single application designer could or would have
14 specified at design time. (We will not cover pipes in lecture today.)
15
16 1. Pseudocode for a very simple shell
17
18     while (1) {
19         write(1, "$ ", 2);
20         readcommand(command, args); // parse input
21         if ((pid = fork()) == 0) // child?
22             exec(command, args, 0);
23         else if (pid > 0) // parent?
24             wait(0); //wait for child
25         else
26             perror("failed to fork");
27     }
28
29 2. Now add two features to this simple shell: output redirection and
30    backgrounding
31
32    By output redirection, we mean, for example:
33        $ ls > list.txt
34
35    By backgrounding, we mean, for example:
36        $ myprog &
37        $
38
39     while (1) {
40         write(1, "$ ", 2);
41         readcommand(command, args); // parse input
42         if ((pid = fork()) == 0) { // child?
43             if (output_redirected) {
44                 close(1);
45                 open(redirect_file, O_CREAT | O_TRUNC | O_WRONLY, 0666);
46             }
47             // when command runs, fd 1 will refer to the redirected file
48             exec(command, args, 0);
49         } else if (pid > 0) { // parent?
50             if (foreground_process) {
51                 wait(0); //wait for child
52             }
53             } else {
54                 perror("failed to fork");
55             }
56         }

```

Jan 17, 13 15:49

I02-handout.txt

Page 2/4

```

57 3. Another syscall example: pipe()
58
59 The pipe() syscall is used by the shell to implement pipelines, such as
60 $ ls | sort | head -4
61 We will see this in a moment; for now, here is an example use of
62 pipes.
63
64 // C fragment with simple use of pipes
65
66     int fdarray[2];
67     char buf[512];
68     int n;
69
70     pipe(fdarray);
71     write(fdarray[1], "hello", 5);
72     n = read(fdarray[0], buf, sizeof(buf));
73     // buf[] now contains 'h', 'e', 'l', 'l', 'o'
74
75 4. File descriptors are inherited across fork
76
77 // C fragment showing how two processes can communicate over a pipe
78
79     int fdarray[2];
80     char buf[512];
81     int n, pid;
82
83     pipe(fdarray);
84     pid = fork();
85     if(pid > 0){
86         write(fdarray[1], "hello", 5);
87     } else {
88         n = read(fdarray[0], buf, sizeof(buf));
89     }
90

```

Jan 17, 13 15:49

I02-handout.txt

Page 3/4

```

91 5. Putting it all together: implementing shell pipelines using
92 fork(), exec(), and pipe(). (See pipesh.c at the back of the
93 handout for a non-pseudocode version of the pipeline handling.)
94
95
96 // Pseudocode for a Unix shell that can run processes in the
97 // background, redirect the output of commands, and implement
98 // two element pipelines, such as "ls | sort"
99
100 void main_loop() {
101
102     while (1) {
103         write(1, "$ ", 2);
104         readcommand(command, args); // parse input
105         if ((pid = fork()) == 0) { // child?
106             if (pipeline_requested) {
107                 handle_pipeline(left_command, right_command)
108             } else {
109                 if (output_redirected) {
110                     close(1);
111                     open(redirect_file, O_CREAT | O_TRUNC | O_WRONLY, 0666);
112                 }
113                 exec(command, args, 0);
114             }
115         } else if (pid > 0) { // parent?
116             if (foreground_process) {
117                 wait(0); // wait for child
118             } else {
119                 perror("failed to fork");
120             }
121         }
122     }
123 }
124
125 void handle_pipeline(left_command, right_command) {
126
127     int fdarray[2];
128
129     if (pipe(fdarray) < 0) panic ("error");
130     if ((pid = fork ()) == 0) { // child (left end of pipe)
131
132         close (1);
133         dup2 (fdarray[1], 1); // make fd 1 the same as fdarray[1],
134                             // which is the write end of the pipe
135
136         close (fdarray[0]);
137         close (fdarray[1]);
138         parse(command1, args1, left_command);
139         exec (command1, args1, 0);
140
141     } else if (pid > 0) { // parent (right end of pipe)
142
143         close (0);
144         dup2 (fdarray[0], 0); // make fd 0 the same as fdarray[0],
145                             // which is the read end of the pipe
146
147         close (fdarray[0]);
148         close (fdarray[1]);
149         parse(command2, args2, right_command);
150         exec (command2, args2, 0);
151
152     } else {
153         printf ("Unable to fork\n");
154     }
155 }

```

Jan 17, 13 15:49

I02-handout.txt

Page 4/4

6. Commentary

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

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

What makes it work is that the author of ls expressed his or her code in terms of a file descriptor:

```
write(1, "some output", byte_count);
```

This author does not, and cannot, know what the file descriptor will represent at runtime. Meanwhile, the shell has the opportunity, *in between fork() and exec()* , to arrange to have that file descriptor represent a pipe, a file to write to, the console, etc.