

```

1 Handout for CS 372H
2 Class 5
3 31 January 2011
4
5 The first four panels, taken together, are 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.

```

1. Pseudocode for a very simple shell

```

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

```

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         exec(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 }
55 }
56

```

3. Another syscall example: pipe()

```

57
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 // 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
74

```

4. File descriptors are inherited across fork

```
75 // C fragment showing how two processes can communicate over a pipe
```

```

76
77     int fdarray[2];
78     char buf[512];
79     int n, pid;
80
81     pipe(fdarray);
82     pid = fork();
83     if(pid > 0){
84         write(fdarray[1], "hello", 5);
85     } else {
86         n = read(fdarray[0], buf, sizeof(buf));
87     }
88
89
90

```

Jan 31, 12 14:26

shell.txt

Page 3/4

```

91 5. Putting it all together: implementing shell pipelines using
92 fork(), exec(), and pipe().
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         close (1);
131         dup2 (fdarray[1], 1); // make fd 1 the same as fdarray[1],
132                               // which is the write end of the pipe
133
134         close (fdarray[0]);
135         close (fdarray[1]);
136         parse(command1, args1, left_command);
137         exec (command1, args1, 0);
138
139     } else if (pid > 0) { // parent (right end of pipe)
140
141         close (0);
142         dup2 (fdarray[0], 0); // make fd 0 the same as fdarray[0],
143                               // which is the read end of the pipe
144
145         close (fdarray[0]);
146         close (fdarray[1]);
147         parse(command2, args2, right_command);
148         exec (command2, args2, 0);
149
150     } else {
151         printf ("Unable to fork\n");
152     }
153 }

```

Jan 31, 12 14:26

shell.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.

Jan 31, 12 13:57

our_head.c

Page 1/1

```

1  /*
2  * our_head.c -- a C program that prints the first L lines of its input,
3  *   where L defaults to 10 but can be specified by the caller of the
4  *   program.
5  *
6  *   (This program is inefficient and does not check its error
7  *   conditions. It is meant to illustrate filters.)
8  */
9  #include <stdlib.h>
10 #include <unistd.h>
11 #include <stdio.h>
12
13 int main(int argc, char** argv)
14 {
15     int i = 0;
16     int nlines;
17     char ch;
18     int ret;
19
20     if (argc == 2) {
21         nlines = atoi(argv[1]);
22     } else if (argc == 1) {
23         nlines = 10;
24     } else {
25         fprintf(stderr, "usage: our_head [nlines]\n");
26         exit(1);
27     }
28
29     for (i = 0; i < nlines; i++) {
30
31         do {
32
33             /* read in the first character from fd 0 */
34             ret = read(0, &ch, 1);
35
36             /* if there are no more characters to read, then exit */
37             if (ret == 0) exit(0);
38
39             write(1, &ch, 1);
40
41         } while (ch != '\n');
42     }
43
44     exit(0);
45 }
46

```

Jan 31, 12 13:57

our_yes.c

Page 1/1

```

1  /*
2  * our_yes.c -- a C program that prints its argument to the screen on a
3  *   new line every second.
4  *
5  */
6  #include <stdlib.h>
7  #include <string.h>
8  #include <unistd.h>
9  #include <stdio.h>
10
11 int main(int argc, char** argv)
12 {
13     char* repeated;
14     int len;
15
16     /* check to make sure the user gave us one argument */
17     if (argc != 2) {
18         fprintf(stderr, "usage: our_yes string_to_repeat\n");
19         exit(1);
20     }
21
22     repeated = argv[1];
23
24     len = strlen(repeated);
25
26     /* loop forever */
27     while (1) {
28
29         write(1, repeated, len);
30
31         write(1, "\n", 1);
32
33         sleep(1);
34     }
35 }
36

```