

```

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.
14
15 1. Pseudocode for a very simple shell
16
17 while (1) {
18     write(1, "$ ", 2);
19     readcommand(command, args); // parse input
20     if ((pid = fork()) == 0) { // child?
21         execve(command, args, 0);
22     } else if (pid > 0) { // parent?
23         wait(0); //wait for child
24     } else {
25         perror("failed to fork");
26     }
27 }
28
29 2. Now add two features to this simple shell: output redirection
30
31 By output redirection, we mean, for example:
32 $ ls > list.txt
33
34 while (1) {
35     write(1, "$ ", 2);
36     readcommand(command, args); // parse input
37     if ((pid = fork()) == 0) { // child?
38         if (output_redirected) {
39             close(1);
40             open(redirect_file, O_CREAT | O_TRUNC | O_WRONLY, 0666);
41         }
42         // when command runs, fd 1 will refer to the redirected file
43         execve(command, args, 0);
44     } else if (pid > 0) { // parent?
45         wait(0); //wait for child
46     } else {
47         perror("failed to fork");
48     }
49 }

```

```

50 3. Another syscall example: pipe()
51
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
57 // C fragment with simple use of pipes
58
59 int fdarray[2];
60 char buf[512];
61 int n;
62
63 pipe(fdarray);
64 write(fdarray[1], "hello", 5);
65 n = read(fdarray[0], buf, sizeof(buf));
66 // buf[] now contains 'h', 'e', 'l', 'l', 'o'
67
68 4. File descriptors are inherited across fork
69
70 // C fragment showing how two processes can communicate over a pipe
71
72 int fdarray[2];
73 char buf[512];
74 int n, pid;
75
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
85
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.

```