

CSCI2467: Systems Programming Concepts

Slideset 8: System System Level I/O

Source: CS:APP Chapter 10, Bryant & O'Hallaron

Instructor: M. Toups

Spring 2020




THE UNIVERSITY of
NEW ORLEANS

DEPARTMENT OF
COMPUTER SCIENCE

- the name `waitpid` can be misleading
 - it does more than just wait, also reaps
- You should use `waitpid` in your `sigchld_handler`
- `waitfg` can be pretty simple
 - just checks the jobs list repeatedly
 - returns once there is no more foreground job
 - but it only works if your jobs list is correctly updated!
 - (which depends on signal handlers)
- You may want to use `fgpid()` and/or `getjobpid()`
putting a while loop around `sleep(1)` is fine
- Much more in the writeup's **Hints** section - use it!

Directory Hierarchy

- All files are organized as a hierarchy anchored by root directory named / (slash)



- Kernel maintains *current working directory (cwd)* for each process
 - Modified using the `cd` command

- **Opening a file informs the kernel that you are getting ready to access that file**

```
int fd;    /* file descriptor */

if ((fd = open("/etc/hosts", O_RDONLY)) < 0) {
    perror("open");
    exit(1);
}
```

- **Returns a small identifying integer *file descriptor***
 - `fd == -1` indicates that an error occurred
- **Each process created by a Linux shell begins life with three open files associated with a terminal:**
 - 0: standard input (stdin)
 - 1: standard output (stdout)
 - 2: standard error (stderr)

- Closing a file informs the kernel that you are finished accessing that file

```
int fd;      /* file descriptor */
int retval; /* return value */

if ((retval = close(fd)) < 0) {
    perror("close");
    exit(1);
}
```

- Closing an already closed file is a recipe for disaster in threaded programs (more on this later)
- Moral: Always check return codes, even for seemingly benign functions such as `close()`

- Writing a file copies bytes from memory to the current file position, and then updates current file position

```
char buf[512];
int fd;      /* file descriptor */
int nbytes; /* number of bytes read */

/* Open the file fd ... */
/* Then write up to 512 bytes from buf to file fd */
if ((nbytes = write(fd, buf, sizeof(buf))) < 0) {
    perror("write");
    exit(1);
}
```

- Returns number of bytes written from `buf` to file `fd`
 - `nbytes < 0` indicates that an error occurred
 - As with reads, short counts are possible and are not errors!

■ Copying stdin to stdout, one byte at a time

```
#include "csapp.h"

int main(void)
{
    char c;

    while(Read(STDIN_FILENO, &c, 1) != 0)
        Write(STDOUT_FILENO, &c, 1);
    exit(0);
}
```

● Shell lab note

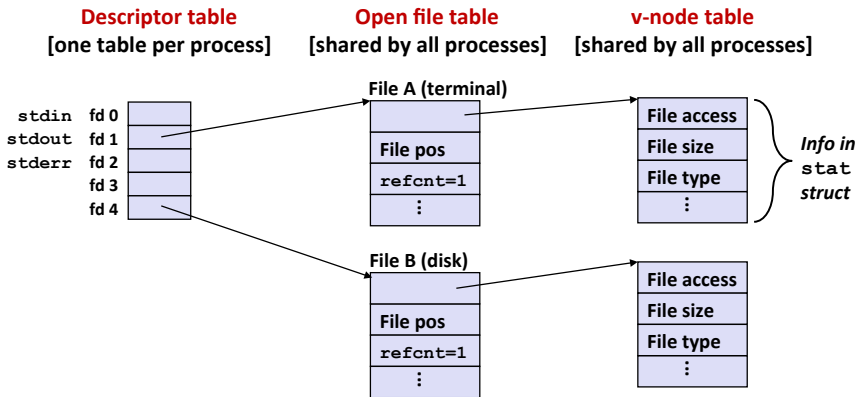
① Race conditions

② System Level I/O

- Unix I/O
- Files
- Opening and closing files
- Metadata, sharing, and redirection
- Standard I/O Functions
- Applying what we know to Shell lab


How the Unix kernel represents open files

- Two descriptors referencing two distinct open files. Descriptor 1 (stdout) points to terminal, and descriptor 4 points to open disk file



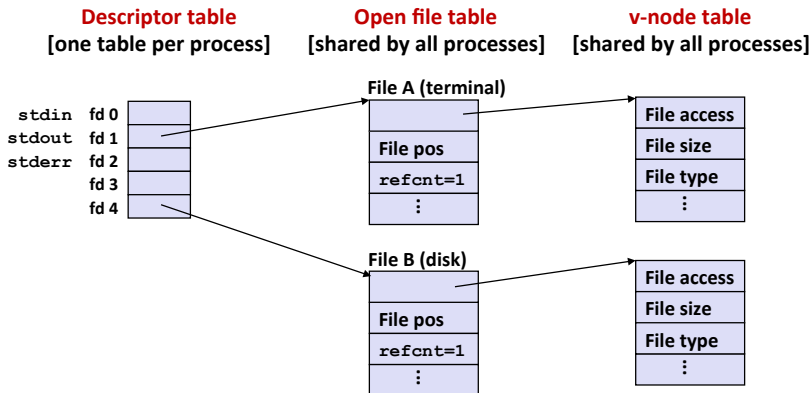
File sharing

- **Two distinct descriptors sharing the same disk file through two distinct open file table entries**
 - E.g., Calling `open` twice with the same `filename` argument



How processes share files: fork

- A child process inherits its parent's open files
 - Note: situation unchanged by `exec` functions (use `fcntl` to change)
- **Before** `fork` call:



- A child process inherits its parent's open files

- Question: How does a shell implement I/O redirection?

```
linux> ls > foo.txt
```

- Answer: By calling the `dup2 (oldfd, newfd)` function
 - Copies (per-process) descriptor table entry `oldfd` to entry `newfd`

Descriptor table
before `dup2 (4, 1)`

fd 0	
fd 1	a
fd 2	
fd 3	
fd 4	b




Descriptor table
after `dup2 (4, 1)`

fd 0	
fd 1	b
fd 2	
fd 3	
fd 4	b

I/O redirection example

- **Step #1: open file to which stdout should be redirected**
 - Happens in child executing shell code, before `exec`



Step #2: call `dup(1, 4)`

● Shell lab note

① Race conditions

② System Level I/O

- Unix I/O
- Files
- Opening and closing files
- Metadata, sharing, and redirection
- Standard I/O Functions
- Applying what we know to Shell lab

Standard I/O Functions

- The C standard library (`libc.so`) contains a collection of higher-level *standard I/O* functions
 - Documented in Appendix B of K&R
- Examples of standard I/O functions:
 - Opening and closing files (`fopen` and `fclose`)
 - Reading and writing bytes (`fread` and `fwrite`)
 - Reading and writing text lines (`fgets` and `fputs`)
 - Formatted reading and writing (`fscanf` and `fprintf`)

Standard I/O Streams

- **Standard I/O models open files as *streams***
 - Abstraction for a file descriptor and a buffer in memory
- **C programs begin life with three open streams (defined in `stdio.h`)**
 - `stdin` (standard input)
 - `stdout` (standard output)
 - `stderr` (standard error)

```
#include <stdio.h>
extern FILE *stdin; /* standard input (descriptor 0) */
extern FILE *stdout; /* standard output (descriptor 1) */
extern FILE *stderr; /* standard error (descriptor 2) */

int main() {
    fprintf(stdout, "Hello, world\n");
}
```