

CIT 5950 Recitation 4 **Solutions** - POSIX I/O Functions

Welcome back to recitation! We're glad that you're here :)

POSIX

Posix is a family of standards specified by the IEEE. These standards maintain compatibility across variants of Unix-like operating systems by defining APIs and standards for basic I/O (file, terminal, and network) and for threading.

1) What does POSIX stand for?

Portable Operating System Interface

2) Why might a POSIX standard be beneficial? From an application perspective? Versus using the C stdio library?

List of answers:

- **More explicit control since read and write functions are system calls and you can directly access system resources.**
- **POSIX calls are unbuffered so you can implement your own buffer strategy on top of read()/write().**
- **There is no standard higher level API for network and other I/O devices**

POSIX and Files

File I/O using POSIX is similar to file I/O using the C stdio library. Some of the operations that can be performed on files using Posix systems calls are: opening a file, reading from a file, writing to a file, closing a file.

```
int open(char* name, int flags, mode_t mode);
```

→ name is a string representing the name of the file. Can be relative or absolute.

→ flags is an integer code describing the access. Some common flags are listed below:

- ◆ O_RDONLY - Open the file in read-only mode.
- ◆ O_WRONLY - Open the file in write-only mode.
- ◆ O_RDWR - Open the file in read-write mode.
- ◆ O_APPEND - Append new information to the end of the file.

★ Returns an integer which is the file descriptor. Returns -1 if there is a failure.

```
int close(int fd);
```

→ fd is the file descriptor (as returned by open()).

★ Returns 0 on success, -1 on failure.

```
ssize_t read(int fd, void *buf, size_t count);
```

```
ssize_t write(int fd, const void *buf, size_t count);
```

→ fd is the file descriptor (as returned by open()).

→ buf is the address of a memory area into which the data is read or written.

→ count is the maximum amount of data to read from or write to the stream.

★ Returns the actual amount of data read from or written to the file.

Exercises:

- 3) A common use of the POSIX I/O function is to **write** to a file; fill in the code skeleton below that writes all of the contents of a string `buf` to the file `595.txt`. *You must use a different method than the "bytes_left" method shown in lecture.*

```
// **NOTE: This is one way to solve this exercise.  
// There exist other correct solutions to this exercise.  
  
int fd = open("595.txt", O_WRONLY); // open 595.txt  
int n = ....;  
char *buf = ..... ; // Assume buf initialized with size n  
int result;  
  
char *ptr = buf; // initialize variable for loop  
  
... // code that populates buf happens here  
  
while (ptr < buf + n) {  
    result = write(fd, ptr, buf + n - ptr);  
  
    if (result == -1) {  
        if (errno != EINTR) {  
            // a real error happened, return an error result  
            close(fd); // cleanup  
            perror("Write failed");  
            return -1;  
        }  
        continue; // EINTR happened, so loop around and try again  
    }  
    ptr += result; // update loop variable  
}  
close(fd); // cleanup
```

- 4) Why is it important to store the return value from the `write()` function? Why do we not check for a return value of 0 like we do for `read()`?

**write() may not actually write all the bytes specified in count.
Writing adds length to your file, so you don't need to check for end of file.**

- 5) Why is it important to remember to call the `close()` function once you have finished working on a file?

In order to free resources i.e. other processes can acquire locks on those files.

POSIX and Errors

Unfortunately, errors that occur when using POSIX system calls are not handled for the user as they are with C standard library functions. So it is important thing is to make sure your code handles errors gracefully.

Note that:

- When an error occurs, the error number is stored in `errno`, which is defined under `<errno.h>`.
- You can use `perror()` to print out a message based on `errno`.
- Remember that `errno` is shared by all library functions and overwritten frequently, so you must read it *right* after an error to be sure of getting the right code.

POSIX functions have a variety of error codes to represent different errors. Some common error conditions:

- ◆ `EBADF` – `fd` is not a valid file descriptor or is not open for reading.
- ◆ `EFAULT` – `buf` is outside your accessible address space.
- ◆ `EINTR` – The call was interrupted by a signal before any data was read. This error, unlike others, is recoverable.
- ◆ `EISDIR` – `fd` refers to a directory.

(Extra Practice) Exercise 6:

6) Given the name of a file as a command-line argument, write a C program that is analogous to `cat`, *i.e.* one that prints the contents of the file to `stdout`. Handle any errors!

```
int main(int argc, char** argv) {
    /* 1. Check to make sure we have a valid command line arguments */
    if (argc != 2) {
        fprintf(stderr, "Usage: ./filedump <filename>\n");
        return EXIT_FAILURE;
    }
    /* 2. Open the file, use O_RDONLY flag */
    int fd = open(argv[1], O_RDONLY);
    if (fd == -1) {
        fprintf(stderr, "Could not open file for reading\n");
        return EXIT_FAILURE;
    }
    /* 3. Read from the file and write it to standard out.*/
    char buf[SIZE];
    ssize_t len;
    do {
        len = read(fd, buf, SIZE);
        if (len == -1) {
            if (errno != EINTR) {
                close(fd);
                perror(NULL);
                return EXIT_FAILURE;
            }
            continue;
        }
        size_t total = 0;
        ssize_t wlen;
        while (total < len) {
            wlen = write(1, buf + total, len - total);
            if (wlen == -1) {
                if (errno != EINTR) {
                    close(fd);
                    perror(NULL);
                    return EXIT_FAILURE;
                }
                continue;
            }
            total += wlen;
        }
    } while (len != 0);
    /*4. Clean up */
    close(fd);
    return EXIT_SUCCESS;
}
```