

**COMP1521 26T1**

**Week 7 Lecture 2**

**File Systems**

# Announcements

**Test 5** (MIPS Strings) and **Test 6** (C bitwise) are due **Thursday 9pm**

**Test 7** is out tomorrow. **Topic** is: C bitwise Operators and Floating Point Representation (I recommend doing after you have completed lab 7 and not before).

# Announcements

**Assignment 1** automarks are available (except for some students with extensions)

Manual style marking for Assignment 1 coming soon...

**Assignment 2** out today: **Files**, bitwise, unicode, processes

Assignment 2 runthrough video coming by Monday

# Announcements

Friday is a **public holiday**:

If you are usually in a Friday `tut_lab`


- Please attend another one this week.
- Either a [make-up tut\\_lab](#) or any regular [tut\\_lab](#)

**Monday next week is a public holiday:**

- No live lecture on Monday Week 8. Pre-recorded lecture instead
- **Lab 7**: is Due Tuesday Week 8 midday

# Today's Lecture

- Recap
  - Stdio basics and buffering
- Stdio buffering and fflush
- Seeking
  - lseek and fseek
- File systems
  - Permissions
  - Hard Links and Symbolic Links
  - Stat syscall
- Working with directories



# Recap Files

**Question 1:** What is better to use to read in a file? `fgetc` or `fgets` or `fscanf`?

**Question 2:** If I successfully open a file using `FILE *f = fopen("data", "w");`

- A. What will happen if the file already exists? What if it doesn't?
- B. What is the difference between mode "a" and "w"?

**Question 3:** How many bytes would the following print to the file `f`:

- A. `fprintf(f, "%d", 255);`
- B. `fputc(255, f);`

# Comparing Libc Wrappers vs stdio

# IO Performance & Buffering libc vs stdio

Let's compare our implementations of cp!

```
$ clang -O3 cp_x.c -o cp_x
```

```
$ dd bs=1M count=10 </dev/urandom >random_file
```

```
10485760 bytes (10 MB, 10 MiB) copied, 0.183075 s, 57.3 MB/s
```

```
$ time ./cp_x random_file random_file_copy
```

Can we get any insights from strace?

```
$ strace ./cp_x random_file random_file_copy
```

Compare:

Linux cp command, cp\_fgetc\_one\_byte.c, cp\_libc\_one\_byte.c, cp\_libc.c

# stdio.h buffering for efficiency

- **Goal:** reduce number of system calls (expensive)
- **Reading:**
  - Uses a **read** system call to fill whole buffer
  - Subsequent reads get bytes from the buffer
  - Does not do another **read** system call till it runs out of data in the buffer
- **Writing:**
  - Delays calls to **write** system call by storing data in buffer (array) instead
  - Calls **write** system call only when
    - Buffer is full
    - a newline is encountered for line buffered output (e.g. terminal)
    - `fflush` is called
    - file is `fclose`d
    - Note: files are closed and buffers are flushed on exit

# fflush stdio buffers

You can manually flush stdio buffers by using:

```
int fflush(FILE *stream);
```

For example

- this would force a write system call to stdout and empty the output buffer  
`fflush(stdout);`
- Can also be used for files that have been opened for writing.
- Should not be used for stdin or files opened for read only.

Demos: `fflush_line.c` `fflush_full.c`

# Assn2 warmup: basic\_stdio.c

Create a type called **basic\_FILE**:

- that contains an integer file descriptor

Implement a basic minimal version of fopen called **basic\_fopen** that

- works for read only
- does not consider buffering
- does no error checking

Implement a basic minimal version of fgetc called **basic\_fgetc** that

- does no error checking
- does not consider EOF
- does not consider buffering

# More File syscalls: Seeking

# Seeking

- So we can now read and write files sequentially... but
  - How do I know which position in the file I am at?
  - How can I skip to the end of the file?
  - How can I go back and read earlier data again?

# Seeking with libc system call wrapper

```
off_t lseek(int fd, off_t offset, int whence);
```

- change the **current position** in given stream
- **offset** is in bytes, and can be negative
- **whence** can be one of
  - SEEK\_SET : set **offset** from start of file
  - SEEK\_CUR: set file **offset** from current position
  - SEEK\_END: set file **offset** from end of file
- returns the new position (or -1 on error with errno set)
- seeking beyond end of file leaves a gap which reads as 0's
- seeking back beyond start of file is an error

# Seeking with stdio.h

```
int fseek(FILE *stream, long offset, int whence);
```

- is stdio equivalent to `lseek()` except:

- requires a `FILE *` input instead of `int` file descriptor
- influences stdio buffers
- returns 0 or -1 for error

```
fseek(stream, 4, SEEK_SET); // move to after 4th byte
```

```
fseek(stream, 2, SEEK_CUR); // 2 bytes forward from current position
```

```
fseek(stream, -3, SEEK_CUR); // 3 bytes backward from current position
```

```
fseek(stream, -1, SEEK_END); // move to before last byte in file
```

```
long ftell(FILE *stream); //return current file position
```

Demo code `fseek.c` and `fuzz.c` and advanced example - `create_gigantic_file.c`

# File Systems

# File Systems

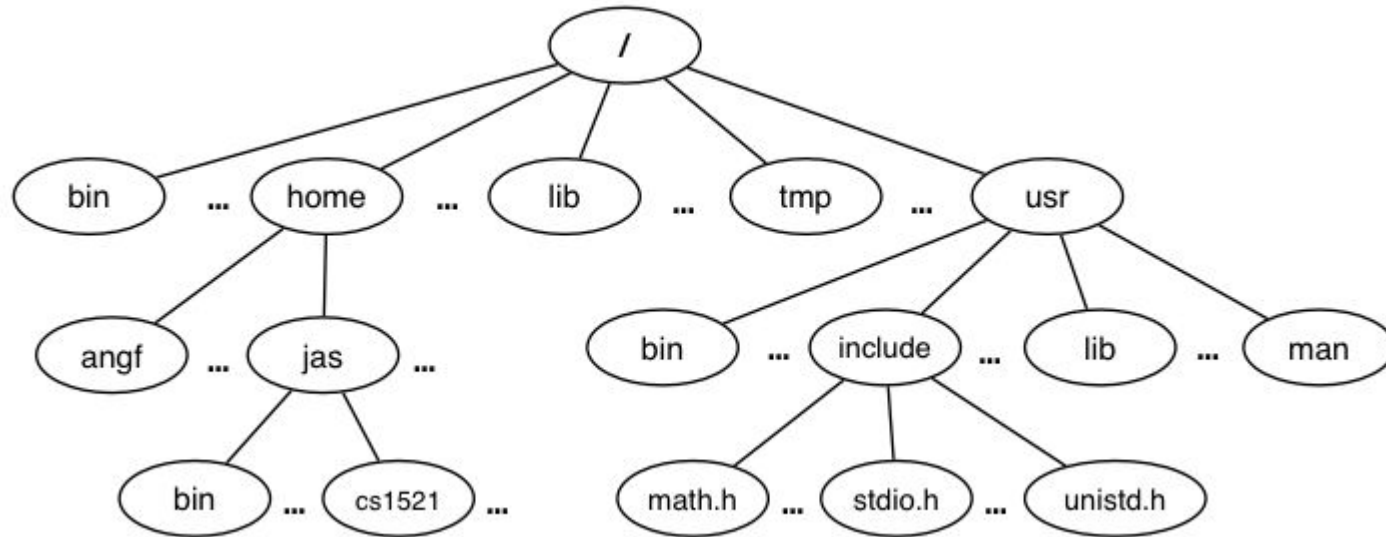
**File systems** manage stored data (e.g. on disk, SSD)

**File** = named sequence of bytes, stored on device

**Directory** = sets of files or directories

- File system maps name to location on device
- File system maintains meta-data (e.g. permissions, time stamps)

# Unix/Linux File System



Unix/Linux file system is tree-like

- symlinks actually make it into a graph
- if traversing you may infinitely loop if following them

# Unix-like File Names

Sequences of 1 or more bytes

- filenames can contain any byte except
- **0x00** bytes (ASCII '\0') used to terminate filenames
- **0x2F** bytes (ASCII '/') used to separate components of pathnames.
- maximum filename length, depends on file system, typically 255

Two filenames have a special meaning:

- . current directory
- .. parent directory

Some programs (shell, ls) treat filenames starting with . specially.

# Paths and directories

**Absolute** pathnames start with a leading / and give full path from root e.g.

- `/usr/include/stdio.h`

**Relative** pathnames do **not** start with a leading / e.g.

- `../..../another/path/prog.c`

- `./a.out`

- `main.c`

# Current Working Directory

Every process (running program) has a **current working directory (CWD)**

- this is an absolute pathname
- this is the directory from where the process was run from
- shell command **pwd** prints the **CWD**

Relative pathnames appended to CWD of process e.g.

- if CWD is **/home/z5555555/lab07/**
- and relative path is **main.c**
- absolute path would be **/home/z5555555/lab07/main.c**

# Unix-like File Metadata

Metadata for file system objects is stored in **inodes**

They hold:

- **location of file contents** in file systems
- file **type** (regular file, directory, ...)
- file **size** in bytes
- file **ownership**
- file **access permissions** - who can read, write, execute the file
- **timestamps** - times of file was created, last accessed, last updated

# Inodes

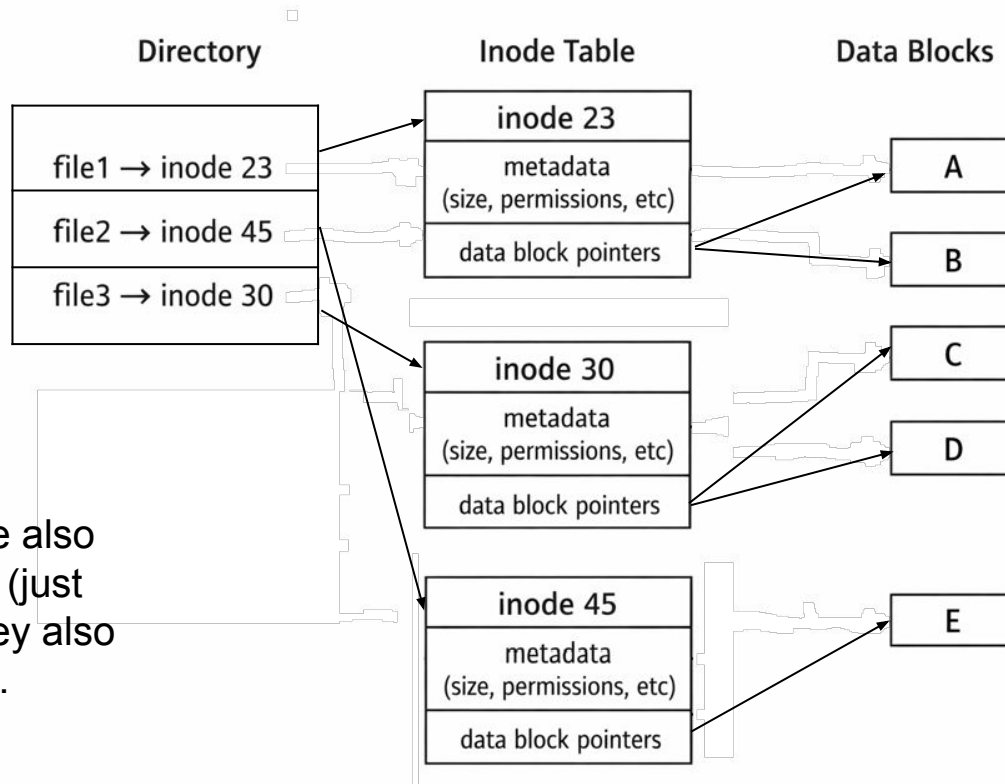
Files system has large table of inodes containing metadata

- Inode-number is the inodes id
  - Unique for file system like zid within UNSW

Directories are effectively a collection of (name, inode-number) pairs

- **ls -i** prints **inode-numbers**

# Directories, Inode Table and Data Blocks



**Note:** Directories are also stored in data blocks (just like regular files). They also have their own inode.

# File Access: Behind the scenes

Access to files by name proceeds (roughly) as:

- **open directory** and scan for **name**
  - if not found, "No such file or directory"
- if found as (**name,inumber**), access **inode table** `inodes[inumber]`
  - collect **file metadata** and:
    - check file access permissions given current user/group
    - if don't have required access, "Permission denied"
    - update access timestamp
- use data in inode (size location) to **access file** contents

# File Permissions

Every file and directory in linux has read, write and execute permissions (access rights) for each of the following user groups:

- **user**: the file's owner
- **group**: the members of the file's group
- **other**: everyone else

Type `ls -l` on command line to see

-

**rwx**

**r--**

**r--**

File Type:

- Normal file

**d** Directory

**l** Symbolic Link

read, write, execute  
permissions for the  
owner of the file

read, write, execute  
permissions for the  
members of the  
**group** owning the file

read, write, execute  
permissions for **other**  
users

# File Permissions: read, write, execute

|                  | <b>Read</b>                    | <b>Write</b>                 | <b>Execute</b>  |
|------------------|--------------------------------|------------------------------|---|
| <b>File</b>      | View contents of file          | Modify file                  | Run as executable   |
| <b>Directory</b> | View names of file e.g. use ls | Create, delete, rename files | Can cd into it. Also needed to access (read, write, execute) items in directory |

# Modifying Permissions

You can think of permissions as a set of bits and then each 3 bits as an octal digit. e.g.

```
rwX  r-x  r-x
111  101  101
   7   5   5
```

You can use the **chmod** command to set the permissions of a file or directory using the desired 3 digit octal code. e.g.

```
$ chmod 700 f.txt
```

# Hard Links and Symbolic Links

File system **links** allow multiple paths to access the same file

## Hard links

- multiple names referencing the **same file (inode)**
- the two entries must be on the same filesystem
- can not create a (extra) hard link to directories
- all hard links to a file have equal status
- file destroyed when last hard link removed
- e.g. Assuming 'fileA' already exists:

**In fileA fileB**

would create a hard link named 'fileB'

# Hard Links and Symbolic Links

File system **links** allow multiple paths to access the same file


## Symbolic links (symlinks)

- point to another **path name**
- accessing the symlink (by default) accesses the file being pointed to
- symbolic link can point to a directory
- symbolic link can point to a pathname on another filesystems
- symbolic links don't have permissions (not needed - they are just a pointer)
- e.g. Assuming 'fileA' already exists:

```
ln -s fileA fileB
```

would create a symbolic link named 'fileB'

# Show me the code!



# C library wrapper for stat system call

```
int stat(const char *pathname, struct stat *statbuf);
```

- returns metadata associated with **pathname** in **statbuf**
- metadata returned includes:
  - inode number
  - type (file, directory, symbolic link, device)
  - size of file in bytes (if it is a file)
  - permissions (read, write, execute)
  - times of last access/modification/status-change
- returns **-1** and sets **errno** if metadata not accessible

# C library wrapper for stat system call

```
int lstat(const char *pathname, struct stat *statbuf);
```

- same as stat() but doesn't follow symbolic links
  - in other words gives you metadata about the symbolic link and not the file it links to
- important not to get stuck in infinite loops

```
int fstat(int fd, struct stat *statbuf);
```

- same as stat() but gets data via an open file descriptor

See **man 2 stat** # For stat syscall and example

**man 3 stat** # For struct stat

**man 7 inode** # For masks, macros and inodes in general

# definition of struct stat

## man 3 stat

```
struct stat {
    dev_t      st_dev;          /* ID of device containing file */
    ino_t      st_ino;         /* Inode number */
    mode_t     st_mode;        /* File type and mode */
    nlink_t    st_nlink;       /* Number of hard links */
    uid_t      st_uid;         /* User ID of owner */
    gid_t      st_gid;         /* Group ID of owner */
    dev_t      st_rdev;        /* Device ID (if special file) */
    off_t      st_size;        /* Total size, in bytes */
    ...
};
```

# st\_mode field of struct stat

## man 7 inode

**st\_mode** is a bitwise-or of these values (& others):

|         |         |                                |
|---------|---------|--------------------------------|
| S_IFLNK | 0120000 | symbolic link                  |
| S_IFREG | 0100000 | regular file                   |
| S_IFDIR | 0040000 | directory                      |
| S_IRUSR | 0000400 | owner has read permission      |
| S_IWUSR | 0000200 | owner has write permission     |
| S_IXUSR | 0000100 | owner has execute permission   |
| S_IRGRP | 0000040 | group has read permission      |
| S_IWGRP | 0000020 | group has write permission     |
| S_IXGRP | 0000010 | group has execute permission   |
| S_IROTH | 0000004 | others have read permission    |
| S_IWOTH | 0000002 | others have write permission   |
| S_IXOTH | 0000001 | others have execute permission |

# Code demos stat.c

stat0.c

stat.c

Tip: There is a good sample program at bottom of man 2 stat

# Making a directory

```
int mkdir(const char *pathname, mode_t mode);
```

returns 0 if successful, returns -1 and sets **errno** otherwise

- for example: `mkdir("newDir", 0755)`

if **pathname** is e.g. ``a/b/c/d``

- all of the directories ``a``, ``b`` and ``c`` must exist
- directory ``c`` must be writable to the caller
- directory ``d`` must not already exist

the new directory contains two initial entries

- ``.`` is a reference to itself
- ``.`` is a reference to its parent directory

Demo: `mkdir.c`

# Opening and Reading directories

// open a directory stream for directory name

```
DIR *opendir(const char *name);
```

// return a pointer to next directory entry

```
struct dirent *readdir(DIR *dirp);
```

// close a directory stream

```
int closedir(DIR *dirp);
```

Found in man 3

Demo list\_directory.c

# Useful Linux (POSIX) functions

`chmod(char *pathname, mode_t mode)` // change permission of file/...

`unlink(char *pathname)` // remove a file...

`rename(char *oldpath, char *newpath)` // rename a file/directory

`chdir(char *path)` // change current working directory

`getcwd(char *buf, size_t size)` // get current working directory

`link(char *oldpath, char *newpath)` // create hard link to a file

`symlink(char *target, char *linkpath)` // create a symbolic link

Demo: `chmod.c` `rm.c` `rename.c` `my_cd.c` `getcwd.c` `nest_directories.c` `many_links.c`  
`chain_links.c`

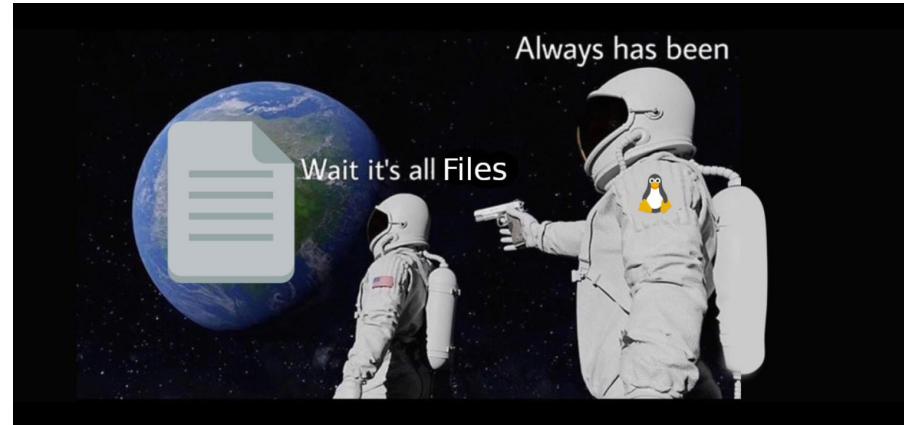
# Everything is a File

Originally files only managed data stored on a magnetic disk.

Unix philosophy is: **Everything is a File**

File system used to access:

- files
- directories (folders)
- storage devices (disks, SSD, ...)
- peripherals (keyboard, mouse, USB, ...)
- system information
- inter-process communication
- network



# Coming up next

Unicode!

Processes!

# Feedback Please!

Your feedback is valuable!

If you have any feedback from today's lecture, please follow the link below or use the QR Code.

Please remember to keep your feedback constructive, so I can action it and improve your learning experience.



<https://forms.office.com/r/hAWAuuhWvh>

# Reach Out

Content Related Questions:

[Forum](#)

Admin related Questions email:

[cs1521@cse.unsw.edu.au](mailto:cs1521@cse.unsw.edu.au)



# Student Support | I Need Help With...

## My Feelings and Mental Health

Managing Low Mood, Unusual Feelings & Depression



**Mental Health Connect**

[student.unsw.edu.au/counselling](https://student.unsw.edu.au/counselling)  
Telehealth



**In Australia Call Afterhours  
UNSW Mental Health Support  
Line**

1300 787 026  
5pm-9am



**Mind HUB**

[student.unsw.edu.au/mind-hub](https://student.unsw.edu.au/mind-hub)  
Online Self-Help Resources



**Outside Australia  
Afterhours 24-hour  
Medibank Hotline**

+61 (2) 8905 0307

## Uni and Life Pressures

Stress, Financial, Visas, Accommodation & More



**Student Support  
Indigenous Student  
Support**

[student.unsw.edu.au/advisors](https://student.unsw.edu.au/advisors)

## Reporting Sexual Assault/Harassment



**Equity Diversity and Inclusion  
(EDI)**

[edi.unsw.edu.au/sexual-misconduct](https://edi.unsw.edu.au/sexual-misconduct)

## Educational Adjustments

To Manage my Studies and Disability / Health Condition



**Equitable Learning Service  
(ELS)**

[student.unsw.edu.au/els](https://student.unsw.edu.au/els)

## Academic and Study Skills



**Academic Language  
Skills**

[student.unsw.edu.au/skills](https://student.unsw.edu.au/skills)

## Special Consideration

Because Life Impacts our Studies and Exams



**Special Consideration**

[student.unsw.edu.au/special-consideration](https://student.unsw.edu.au/special-consideration)