

CS631 - Advanced Programming in the UNIX Environment

—

File Systems, System Data Files, Time & Date

Department of Computer Science
Stevens Institute of Technology
Jan Schaumann

`jschauma@stevens.edu`

`https://stevens.netmeister.org/631/`

HW#1 Notes: Avoid useless comments

```
/* open the source file */
if((fd_src = open(argv[1],O_RDONLY)) == -1){

//check the number of commandline argument
if(argc != 3)

//output the usage of this app.^M
if(argc != 3){^M
    (void)fprintf(stderr,^M
        "Usage: %s [file1] [file2]\n\^M

/* Buffer size */
#define BUFSIZE 4096

/* Check if the destination is a directory */
if(stat(argv[2], &s) != -1 && S_ISDIR(s.st_mode))
```

HW#1 Notes

Errors go to stderr, please! Use `strerror(3)`/`perror(3)`.

```
printf("malloc failed\n");
```

```
printf("Cannot read your source file.\n");
```

```
if(!S_ISREG(src.st_mode) )  
{  
    printf("Source file is not a regular file\n");^M
```

```
    printf("Unable to read %s\n",input_src  
        );
```

HW#1 Notes: Check all return codes!

```
file_in_dir = calloc(len, 1);
snprintf(file_in_dir, len, "%s/%s", argv[2], basename);

int isDirectory(char* testString) {^M
    struct stat dir;^M
    if(stat(testString, &dir) >= 0) {^M
        if(S_ISDIR(dir.st_mode))^M
            return 1;^M
    }^M
    return 0;^M
}^M

dstPath = (char*)malloc(strlen(argv[2]) + strlen(filename) + 2);
dstPath = strcpy(dstPath, argv[2]);

dir_fd=dirfd(dir);
dest_file=openat(dir_fd,argv[1],O_WRONLY | O_CREAT , 0641);
```

HW#1 Notes: Don't overflow your buffers!

Use of `strcat(3)` etc. considered harmful. Use `strlcat(3)/strlcpy(3)` etc. instead. Know your buffer boundaries!

```
char *writePath = argv[2]
bname = basename(readPath);
strcat(writePath, "/");
strcat(writePath, bname);
```

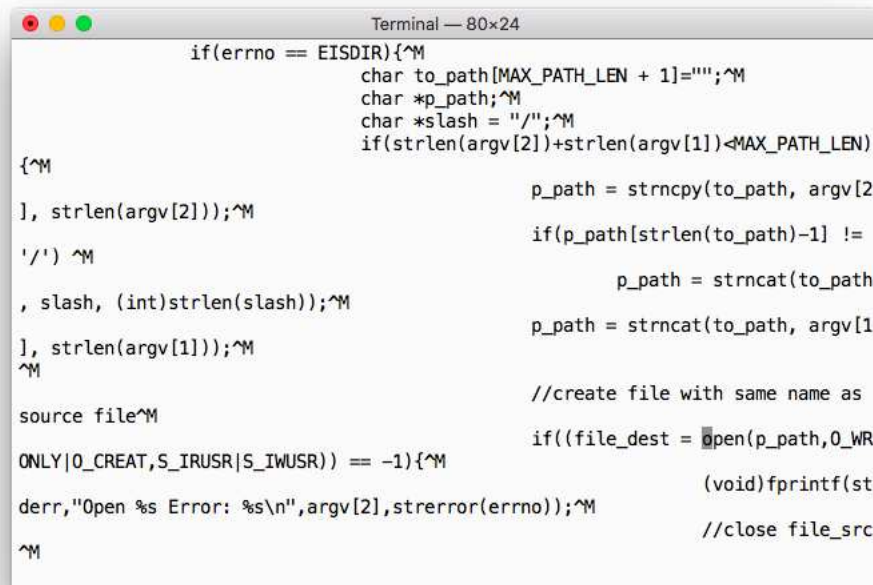
```
char filename[100] = "";
strcat(filename, argv[2]);
strcat(filename, "/");
strcat(filename, getFilename(argv[1]));
```

```
target_path = (char *)malloc(sizeof(argv[2]) + 2 +
    sizeof(basename(argv[1])));
(void)strcat(target_path, basename(argv[1]));
```

But beware: `sizeof` != `strlen(3)` See: `sizeof.c`

HW#1 Notes: Goldilocks Whitespace

```
if((sourceFile=(char*)calloc(strlen(argv[1])+1,sizeof(char)))!=NULL)^M
    strcpy(sourceFile,argv[1]);^M
```



```
Terminal — 80x24
    if(errno == EISDIR){^M
        char to_path[MAX_PATH_LEN + 1]="";^M
        char *p_path;^M
        char *slash = "/";^M
        if(strlen(argv[2])+strlen(argv[1])<MAX_PATH_LEN)
    {^M
        p_path = strncpy(to_path, argv[2]
    ], strlen(argv[2]));^M
        if(p_path[strlen(to_path)-1] !=
        '/') ^M
        p_path = strncat(to_path
    , slash, (int)strlen(slash));^M
        p_path = strncat(to_path, argv[1]
    ], strlen(argv[1]));^M
    ^M
    //create file with same name as
    source file^M
    if((file_dest = open(p_path,O_WR
    ONLY|O_CREAT,S_IRUSR|S_IWUSR)) == -1){^M
        (void)fprintf(st
    derr,"Open %s Error: %s\n",argv[2],strerror(errno));^M
    //close file_src
    ^M
```

HW#1 Notes: Flow

```
if (condition) {
    some
    long
    block
    here
    with
    if (condition) {
        nested
    } else {
        conditions
    }
} else {
    error
}
```

```
if (!condition) {
    error
}

some
long
block
here
with
if (condition) {
    nested
} else {
    conditions
}
```

File Systems

System Data Files, Time & Date

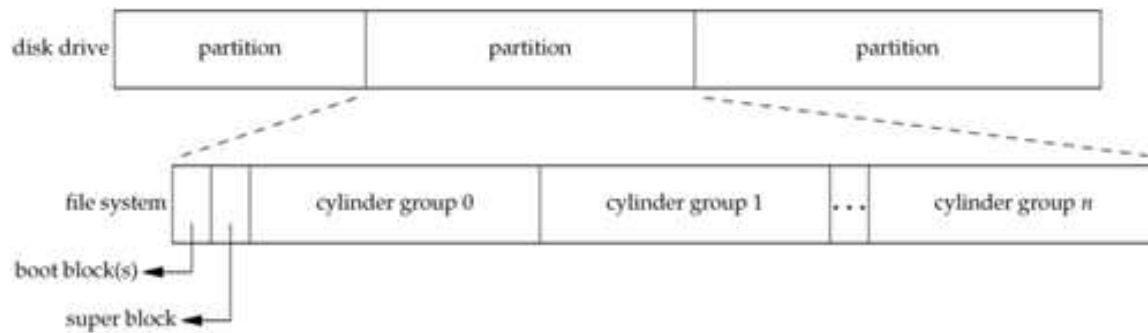
File Systems

- a disk can be divided into logical *partitions*



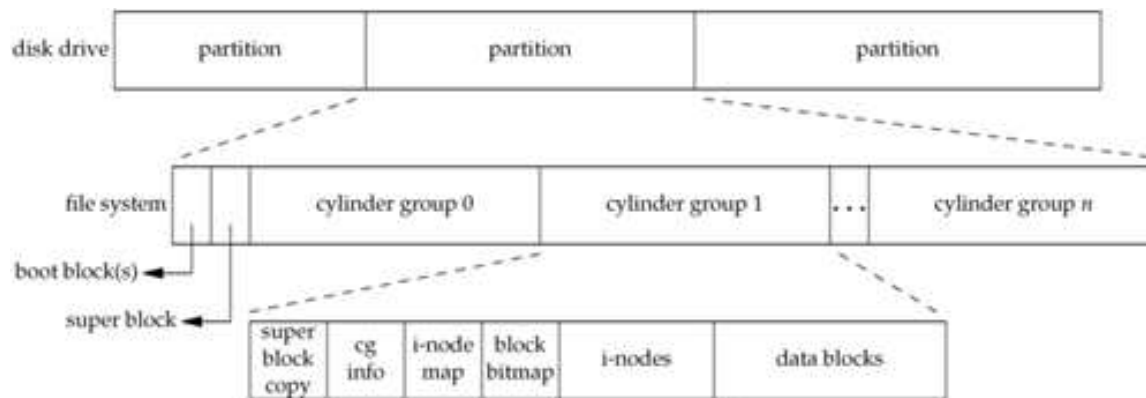
File Systems

- a disk can be divided into logical *partitions*
- each logical *partition* may be further divided into *file systems* containing *cylinder groups*



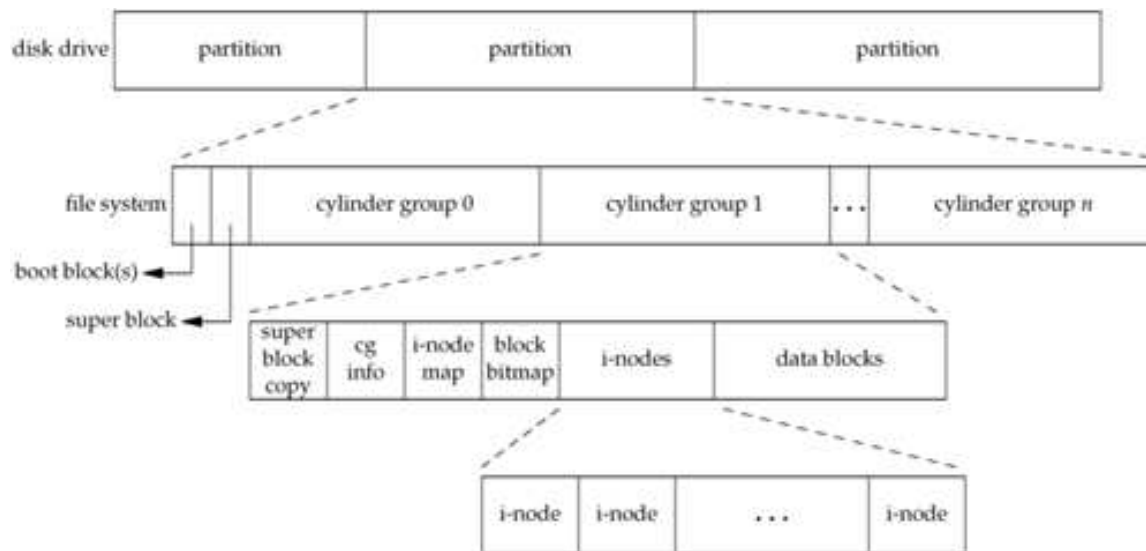
File Systems

- a disk can be divided into logical *partitions*
- each logical *partition* may be further divided into *file systems* containing *cylinder groups*
- each *cylinder group* contains a list of *inodes (i-list)* as well as the actual *directory-* and *data blocks*



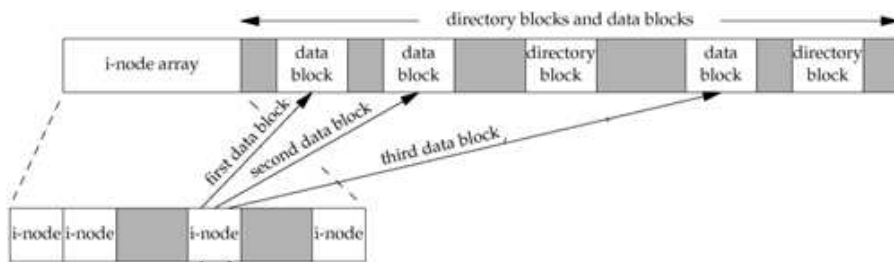
File Systems

- a disk can be divided into logical *partitions*
- each logical *partition* may be further divided into *file systems* containing *cylinder groups*
- each *cylinder group* contains a list of *inodes* (*i-list*) as well as the actual *directory-* and *data blocks*



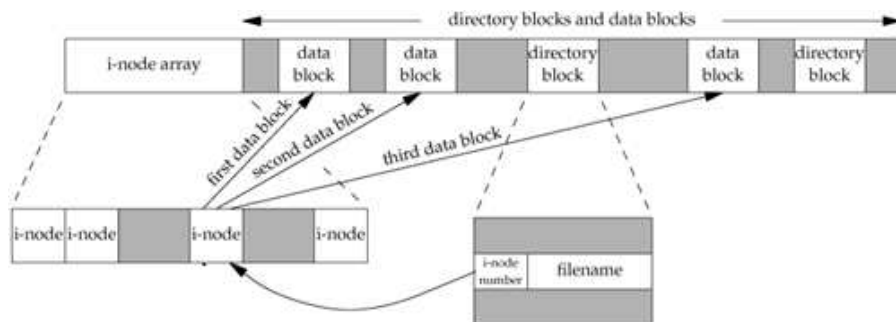
File Systems

- a disk can be divided into logical *partitions*
- each logical *partition* may be further divided into *file systems* containing *cylinder groups*
- each *cylinder group* contains a list of *inodes* (*i-list*) as well as the actual *directory-* and *data blocks*



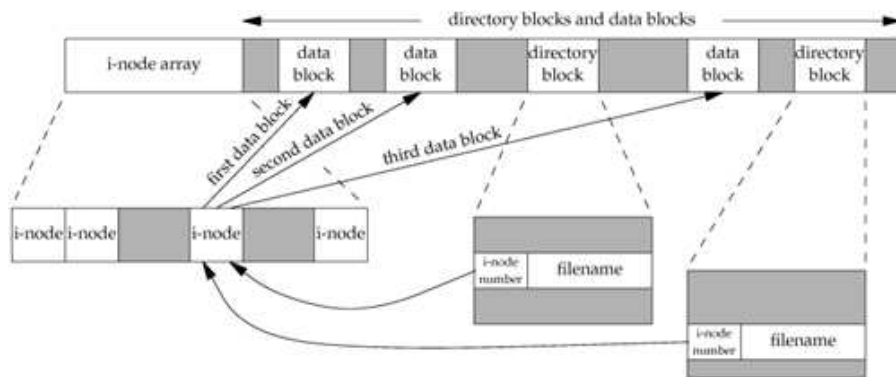
File Systems

- a disk can be divided into logical *partitions*
- each logical *partition* may be further divided into *file systems* containing *cylinder groups*
- each *cylinder group* contains a list of *inodes* (*i-list*) as well as the actual *directory-* and *data blocks*
- a directory entry is really just a *hard link* mapping a “filename” to an inode



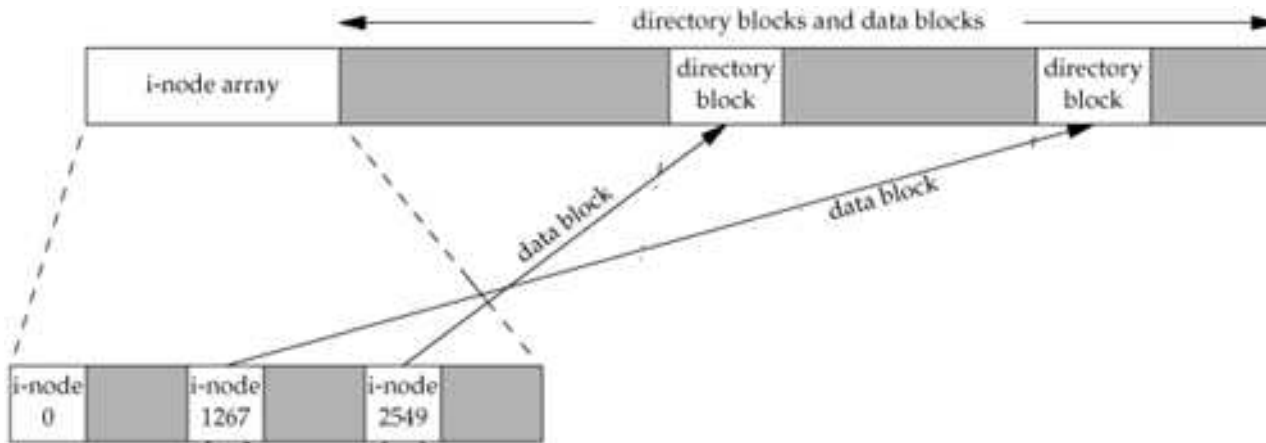
File Systems

- a disk can be divided into logical *partitions*
- each logical *partition* may be further divided into *file systems* containing *cylinder groups*
- each *cylinder group* contains a list of *inodes* (*i-list*) as well as the actual *directory-* and *data blocks*
- a directory entry is really just a *hard link* mapping a “filename” to an inode
- you can have many such mappings to the same file



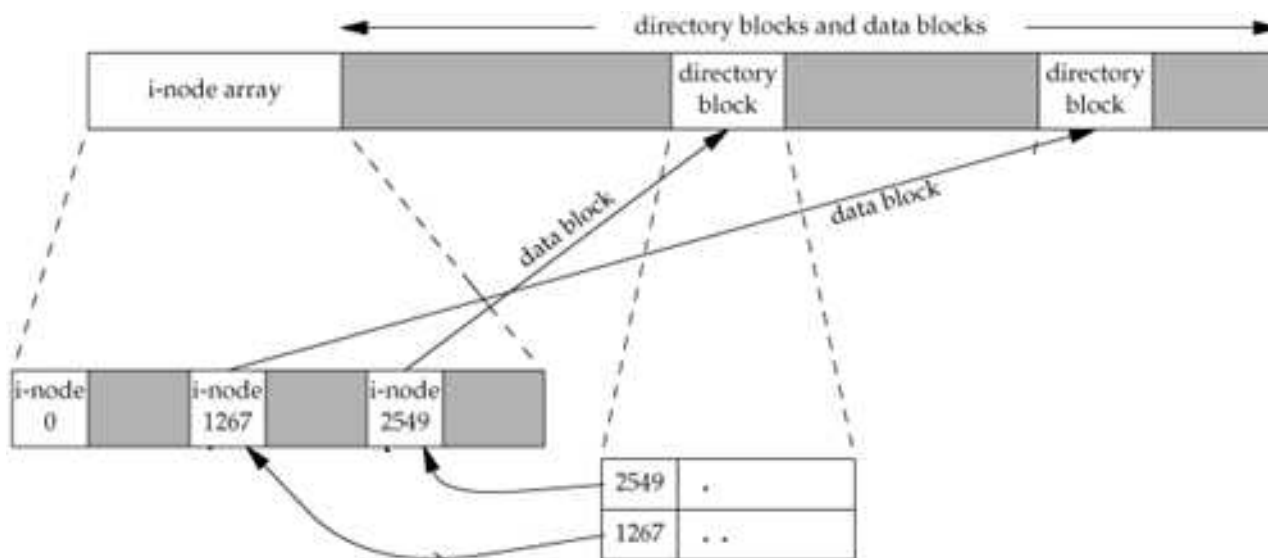
Directories

- directories are special "files" containing hardlinks



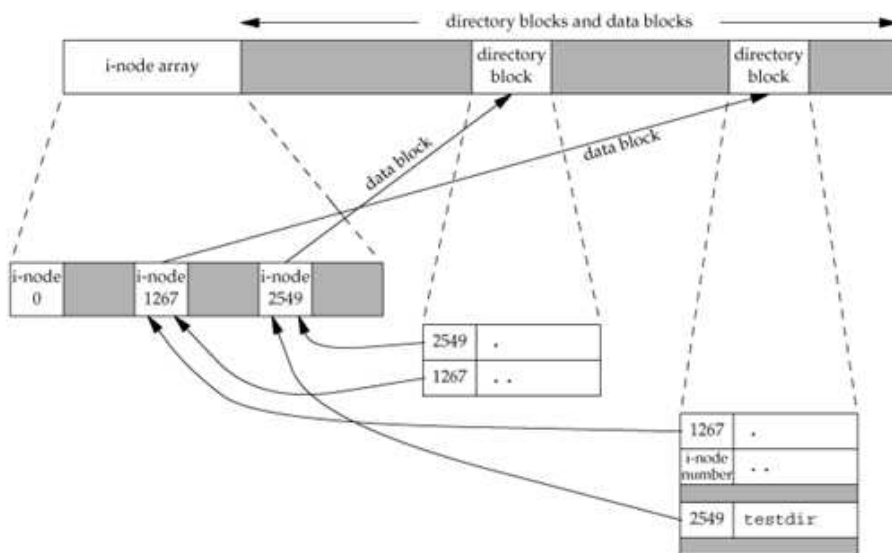
Directories

- directories are special "files" containing hardlinks
- each directory contains at least two entries:
 - . (*this* directory)
 - .. (the parent directory)



Directories

- directories are special "files" containing hardlinks
- each directory contains at least two entries:
 - . (*this* directory)
 - .. (the parent directory)
- the link count (`st_nlink`) of a directory is at least 2



Inodes

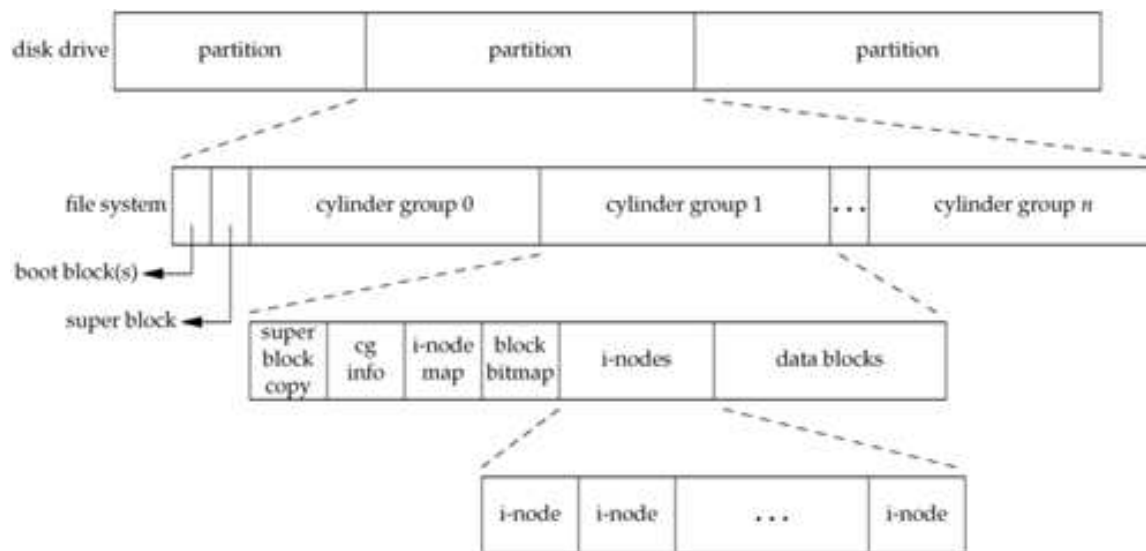
- the *inode* contains most of the information found in the `stat` structure.

Inodes

- the *inode* contains most of the information found in the `stat` structure.
- every *inode* has a *link count* (`st_nlink`): it shows how many “things” point to this inode. Only if this *link count* is 0 (and no process has the file open) are the *data blocks* freed.

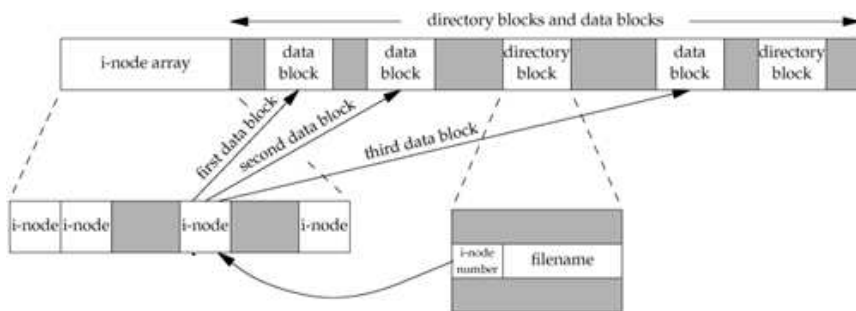
Inodes

- the *inode* contains most of the information found in the `stat` structure.
- every *inode* has a *link count* (`st_nlink`): it shows how many “things” point to this inode. Only if this *link count* is 0 (and no process has the file open) are the *data blocks* freed.
- *inode* number in a directory entry must point to an *inode* on the same file system (no hardlinks across filesystems)



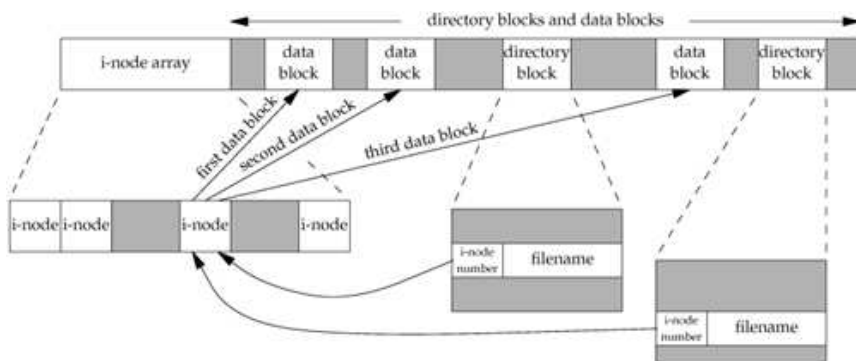
Inodes

- the *inode* contains most of the information found in the `stat` structure.
- every *inode* has a *link count* (`st_nlink`): it shows how many “things” point to this inode. Only if this *link count* is 0 (and no process has the file open) are the *data blocks* freed.
- *inode* number in a directory entry must point to an *inode* on the same file system (no hardlinks across filesystems)
- to move a file within a single filesystem, we can just “move” the directory entry (actually done by creating a new entry, and deleting the old one).



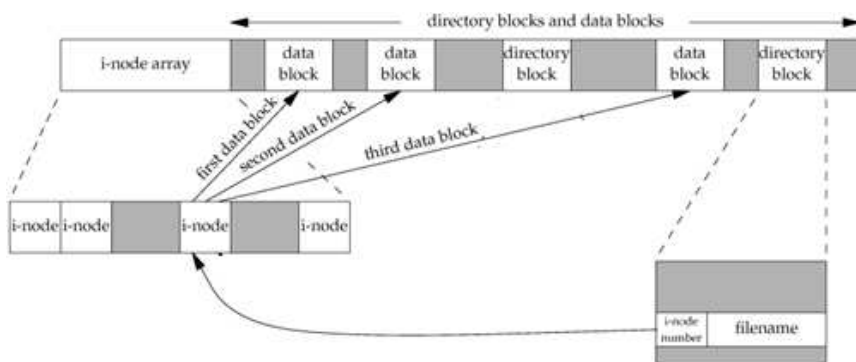
Inodes

- the *inode* contains most of the information found in the `stat` structure.
- every *inode* has a *link count* (`st_nlink`): it shows how many “things” point to this inode. Only if this *link count* is 0 (and no process has the file open) are the *data blocks* freed.
- *inode* number in a directory entry must point to an *inode* on the same file system (no hardlinks across filesystems)
- to move a file within a single filesystem, we can just “move” the directory entry (actually done by creating a new entry, and deleting the old one).



Inodes

- the *inode* contains most of the information found in the `stat` structure.
- every *inode* has a *link count* (`st_nlink`): it shows how many “things” point to this inode. Only if this *link count* is 0 (and no process has the file open) are the *data blocks* freed.
- *inode* number in a directory entry must point to an *inode* on the same file system (no hardlinks across filesystems)
- to move a file within a single filesystem, we can just “move” the directory entry (actually done by creating a new entry, and deleting the old one).



link(2)

```
#include <unistd.h>
```

```
int link(const char *name1, const char *name2);
```

Returns: 0 if OK, -1 on error

- Creates a link to an existing file (hard link).
- POSIX.1 allows links to cross filesystems, most implementations (SVR4, BSD) don't.
- only uid(0) can create links to directories (loops in filesystem are bad)

link(2) and unlink(2)

```
#include <unistd.h>

int link(const char *name1, const char *name2);

Returns: 0 if OK, -1 on error
```

- Creates a link to an existing file (hard link).
- POSIX.1 allows links to cross filesystems, most implementations (SVR4, BSD) don't.
- only uid(0) can create links to directories (loops in filesystem are bad)

```
#include <unistd.h>

int unlink(const char *path);

Returns: 0 if OK, -1 on error
```

- removes directory entry and decrements link count of file
- if file link count == 0, free data blocks associated with file (...unless processes have the file open)

link(2) and unlink(2)

```
$ cc -Wall wait-unlink.c  
$ ./a.out  
$ df .
```

rename(2)

```
#include <stdio.h>
```

```
int rename(const char *from, const char *to);
```

Returns: 0 if OK, -1 on error

If *oldname* refers to a file:

- if *newname* exists and it is not a directory, it's removed and *oldname* is renamed *newname*

rename(2)

```
#include <stdio.h>
```

```
int rename(const char *from, const char *to);
```

Returns: 0 if OK, -1 on error

If *oldname* refers to a file:

- if *newname* exists and it is not a directory, it's removed and *oldname* is renamed *newname*
- if *newname* exists and it is a directory, an error results

rename(2)

```
#include <stdio.h>
```

```
int rename(const char *from, const char *to);
```

Returns: 0 if OK, -1 on error

If *oldname* refers to a file:

- if *newname* exists and it is not a directory, it's removed and *oldname* is renamed *newname*
- if *newname* exists and it is a directory, an error results
- must have w+x perms for the directories containing *old/newname*

rename(2)

```
#include <stdio.h>
```

```
int rename(const char *from, const char *to);
```

Returns: 0 if OK, -1 on error

If *oldname* refers to a file:

- if *newname* exists and it is not a directory, it's removed and *oldname* is renamed *newname*
- if *newname* exists and it is a directory, an error results
- must have w+x perms for the directories containing *old/newname*

If *oldname* refers to a directory:

- if *newname* exists and is an empty directory (contains only . and ..), it is removed; *oldname* is renamed *newname*

rename(2)

```
#include <stdio.h>
```

```
int rename(const char *from, const char *to);
```

Returns: 0 if OK, -1 on error

If *oldname* refers to a file:

- if *newname* exists and it is not a directory, it's removed and *oldname* is renamed *newname*
- if *newname* exists and it is a directory, an error results
- must have w+x perms for the directories containing *old/newname*

If *oldname* refers to a directory:

- if *newname* exists and is an empty directory (contains only . and ..), it is removed; *oldname* is renamed *newname*
- if *newname* exists and is a file, an error results

rename(2)

```
#include <stdio.h>
```

```
int rename(const char *from, const char *to);
```

Returns: 0 if OK, -1 on error

If *oldname* refers to a file:

- if *newname* exists and it is not a directory, it's removed and *oldname* is renamed *newname*
- if *newname* exists and it is a directory, an error results
- must have w+x perms for the directories containing *old/newname*

If *oldname* refers to a directory:

- if *newname* exists and is an empty directory (contains only . and ..), it is removed; *oldname* is renamed *newname*
- if *newname* exists and is a file, an error results
- if *oldname* is a prefix of *newname* an error results

rename(2)

```
#include <stdio.h>
```

```
int rename(const char *from, const char *to);
```

Returns: 0 if OK, -1 on error

If *oldname* refers to a file:

- if *newname* exists and it is not a directory, it's removed and *oldname* is renamed *newname*
- if *newname* exists and it is a directory, an error results
- must have w+x perms for the directories containing *old/newname*

If *oldname* refers to a directory:

- if *newname* exists and is an empty directory (contains only . and ..), it is removed; *oldname* is renamed *newname*
- if *newname* exists and is a file, an error results
- if *oldname* is a prefix of *newname* an error results
- must have w+x perms for the directories containing *old/newname*

Symbolic Links

```
#include <unistd.h>
```

```
int symlink(const char *name1, const char *name2);
```

Returns: 0 if OK, -1 on error

- file whose "data" is a path to another file
- anyone can create symlinks to directories or files
- certain functions dereference the link, others operate on the link

How do we get the contents of a symlink? `open(2)` and `read(2)`?

Symbolic Links

```
#include <unistd.h>

int symlink(const char *name1, const char *name2);

Returns: 0 if OK, -1 on error
```

- file whose "data" is a path to another file
- anyone can create symlinks to directories or files
- certain functions dereference the link, others operate on the link

```
#include <unistd.h>

int readlink(const char *path, char *buf, size_t bufsize);

Returns: number of bytes placed into buffer if OK, -1 on error
```

This function combines the actions of `open`, `read`, and `close`.
Note: *buf* is not NUL terminated.

File Times

```
#include <sys/types.h>

int utimes(const char *path, const struct timeval times[2]);
int lutimes(const char *path, const struct timeval times[2]);
int futimes(int fd, const struct timeval times[2]);
Returns: 0 if OK, -1 on error
```

If *times* is NULL, access time and modification time are set to the current time (must be owner of file or have write permission). If *times* is non-NULL, then times are set according to the `timeval` struct array. For this, you must be the owner of the file (write permission not enough).

Note that `st_ctime` is set to the current time in both cases.

For the effect of various functions on the access, modification and changes-status times see Stevens, p. 117.

Note: some systems implement `lutimes(3)` (library call) via `utimes(2)` syscalls.

mkdir(2) and rmdir(2)

```
#include <sys/types.h>
#include <sys/stat.h>

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

Returns: 0 if OK, -1 on error

Creates a new, empty (except for . and .. entries) directory. Access permissions specified by *mode* and restricted by the `umask(2)` of the calling process.

```
#include <unistd.h>

int rmdir(const char *path);
```

Returns: 0 if OK, -1 on error

If the link count is 0 (after this call), and no other process has the directory open, directory is removed. Directory must be empty (only . and .. remaining)

Reading Directories

```
#include <sys/types.h>
#include <dirent.h>

DIR *opendir(const char *filename);
                                Returns: pointer if OK, NULL on error

struct dirent *readdir(DIR *dp);
                                Returns: pointer if OK, NULL at end of dir or on error

void rewinddir(DIR *dp);
int closedir(DIR *dp);
                                Returns: 0 if OK, -1 on error
```

- read by anyone with read permission on the directory
- format of directory is implementation dependent (always use `readdir` and friends)

`opendir`, `readdir` and `closedir` should be familiar from our small `ls` clone. `rewinddir` resets an open directory to the beginning so `readdir` will again return the first entry.

For directory traversal, consider `fts(3)` (not available on all UNIX versions).

Moving around directories

```
#include <unistd.h>

char *getcwd(char *buf, size_t size);
```

Returns: *buf* if OK, NULL on error

Get the kernel's idea of our process's current working directory.

```
#include <unistd.h>

int chdir(const char *path);
int fchdir(int fd);
```

Returns: 0 if OK, -1 on error

Allows a process to change its current working directory. Note that `chdir` and `fchdir` affect only the current process.

```
$ cc -Wall cd.c
```

```
$ ./a.out /tmp
```


Password File

Called a *user database* by POSIX and usually found in `/etc/passwd`, the password file contains the following fields:

Description	struct passwd member	POSIX.1
username	char *pw_name	X
hashed passwd	char *pw_passwd	
numerical user id	uid_t pw_uid	X
numerical group id	gid_t pw_gid	X
comment field	char *pw_gecos	
initial working directory	char *pw_dir	X
initial shell	char *pw_shell	X

Password field is a one-way hash of the users password.
Some fields can be empty:

- password empty implies no password
- shell empty implies `/bin/sh`

Password File

```
#include <sys/types.h>
#include <pwd.h>

struct passwd *getpwuid(uid_t uid);
struct passwd *getpwnam(const char *name);
```

Returns: pointer if OK, NULL on error

```
#include <sys/types.h>
#include <pwd.h>

struct passwd *getpwent(void);

void setpwent(void);
void endpwent(void);
```

Returns: pointer if OK, NULL on error

- `getpwent` returns next password entry in file each time it's called, no order
- `setpwent` rewinds to "beginning" of entries
- `endpwent` closes the file(s)

See also: `getspnam(3)`/`getspent(3)` (where available)

Group File

Called a *group database* by POSIX and usually found in `/etc/group`, the group file contains the following fields:

Description	struct group member	POSIX.1
groupname	char *gr_name	X
hashed passwd	char *gr_passwd	
numerical group id	uid_t gr_uid	X
array of pointers to user names	char **gr_mem	X

The `gr_mem` array is terminated by a NULL pointer.

Group File

```
#include <sys/types.h>
#include <grp.h>

struct group *getgrgid(gid_t gid);
struct group *getgrnam(const char *name);
```

Returns: pointer if OK, NULL on error

These allow us to look up an entry given a user's group name or numerical GID. What if we need to go through the group file entry by entry? Nothing in POSIX.1, but SVR4 and BSD give us:

```
#include <sys/types.h>
#include <grp.h>

struct group *getgrent(void);

void setgrent(void);
void endgrent(void);
```

Returns: pointer if OK, NULL on error

- `getgrent` returns next group entry in file each time it's called, no order
- `setgrent` rewinds to "beginning" of entries
- `endgrent` closes the file(s)

Supplementary Groups and other data files

```
#include <sys/types.h>
#include <unistd.h>

int getgroups(int gidsetsize, gid_t *grouplist);
    Returns: returns number of suppl. groups if OK, -1 on error
```

Note: if `gidsetsize == 0`, `getgroups(2)` returns number of groups without modifying `grouplist`.

Other system databases

Similar routines as for password/group for accessing system data files:

Description	Data file	Header	Structure	Additional lookup functions
hosts	/etc/hosts	<netdb.h>	hostent	gethostbyname gethostbyaddr
networks	/etc/networks	<netdb.h>	netent	genetbyname getnetbyaddr
protocols	/etc/protocols	<netdb.h>	protoent	getprotobyname getprotobynumber
services	/etc/services	<netdb.h>	servent	getservbyname getservbyport

System Identification

```
#include <sys/utsname.h>

int uname(struct utsname *name);
    Returns: nonnegative value if OK, -1 on error
```

- Pass a pointer to a `utsname` struct. This struct contains fields like `opsys` name, version, release, architecture, etc.
- This function used by the `uname(1)` command (try `uname -a`)
- Not that the size of the fields in the `utsname` struct may not be large enough to id a host on a network

To get just a hostname that will identify you on a TCP/IP network, use the Berkeley-derived:

```
#include <unistd.h>

int gethostname(char *name, int namelen);
    Returns: 0 if OK, -1 on error
```

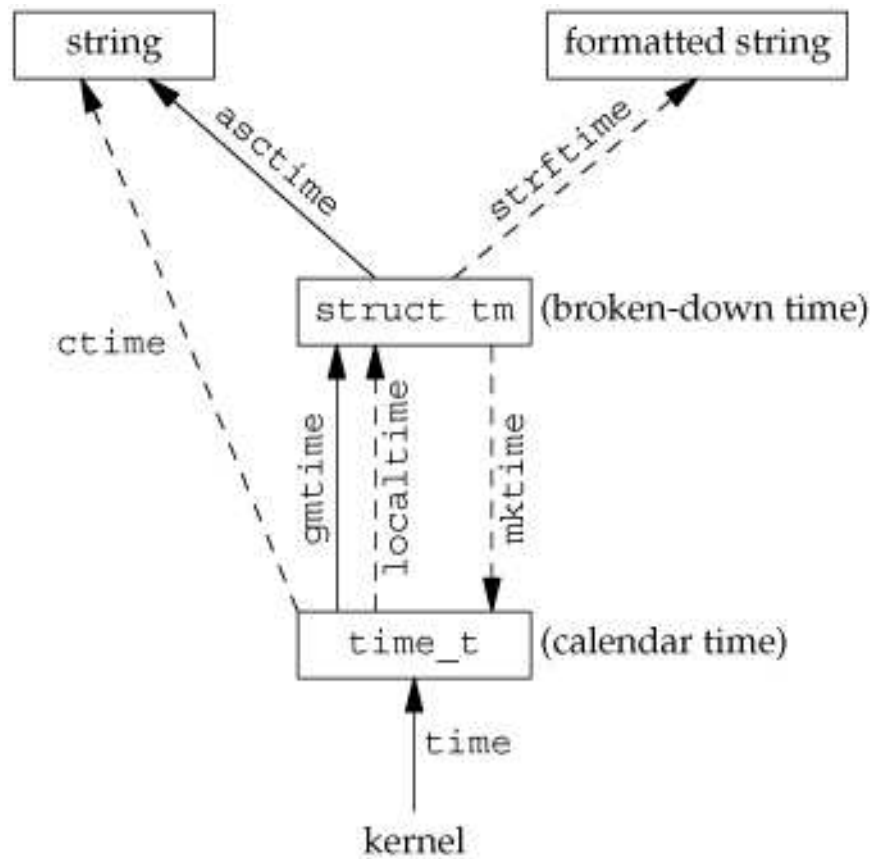
Time and Date

```
#include <time.h>

time_t time(time_t *tloc);
    Returns: value of time if OK, -1 on error
```

- Time is kept in UTC
- Time conversions (timezone, daylight savings time) handled "automatically"
- Time and date kept in a single quantity (`time_t`)

Time and Date



Time and Date

We can break this `time_t` value into its components with either of the following:

```
#include <time.h>

struct tm *gmtime(const time_t *calptr);
struct tm *localtime(const time_t *calptr);
                Returns: pointer to broken down time
```

`localtime(3)` takes into account daylight savings time and the `TZ` environment variable.

```
#include <time.h>

time_t mktime(struct tm *tmpr);
                Returns: calendar time if OK, -1 on error
```

The `mktime(3)` function operates in the reverse direction.

Time and Date

To output human readable results, use:

```
#include <time.h>

char *asctime(const struct tm *tm_ptr);
char *ctime(const struct tm *tm_ptr);
        Returns: pointer to NULL terminated string
```

Lastly, there is a printf(3) like function for times:

```
#include <time.h>

size_t strftime(char *buf, size_t maxsize, const char *restricted format, const struct tm *timeptr);
        Returns: number of characters stored in array if room, else 0
```

Homework

Reading:

- Stevens, Chapter 4 and 6
- Falsehoods Programmers believe about time: <http://is.gd/yFSYR0>

Think about code. You should now be able to implement:

- `ln(1)`, `mv(1)`, `rm(1)`, `rmdir(1)`, `stat(1)`
- `date(1)`, `id(1)`, `touch(1)`, `uname(1)`

Other:

- work on your midterm project!