Operating Systems

File Systems
(Select parts of Ch 6)

Outline

- Files
- Directories
- Partitions

File Systems

- Abstraction to disk (convenience)
  - "The only thing friendly about a disk is that it has persistent storage."
  - Devices may be different: tape, IDE/SCSI, NFS
- Users
  - don't care about detail
  - care about interface (won't cover, assumed knowledge)
- OS
  - cares about implementation (efficiency)

File System Structures

- *Files* - store the data
- *Directories* - organize files
- *Partitions* - separate collections of directories (also called "volumes")
  - all directory information kept in partition
  - `mount` file system to access

Example: Unix `open()`

```
int open(char *path, int flags [, int mode])
```

- `path` is name of file
- `flags` is bitmap to set switch
  - O_RDONLY, O_WRONLY…
  - O_CREATE then use `mode` for perms
- success, returns index

Unix `open()` - Under the Hood

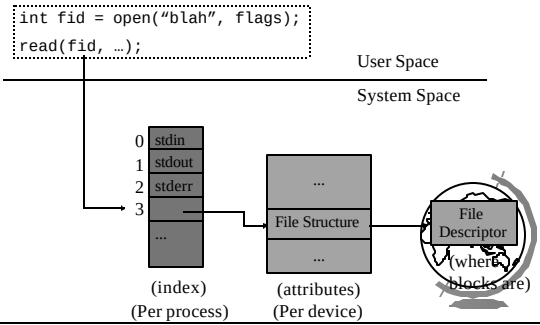

File System Implementation

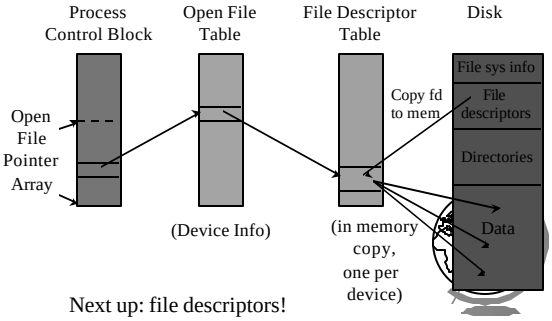

Next up: file descriptors!

File System Implementation

- Which blocks with which file?
- File descriptor implementations:
  - Contiguous
  - Linked List
  - Linked List with Index
  - I-nodes

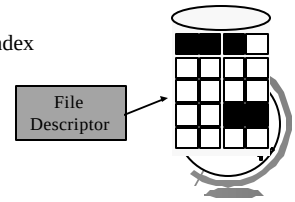

I-nodes

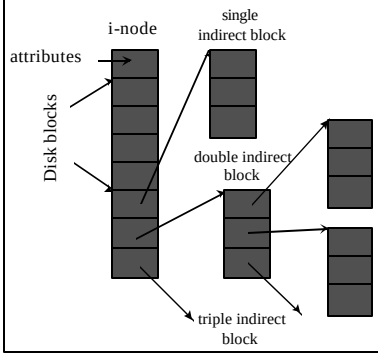

- Fast for small files
- Can hold big files
- Size?
  - 4 kbyte block

Outline

- Files (done)
- Directories ¬
- Partitions

Directories

- Before reading file, must be opened
- Directory entry provides information to get blocks
  - disk location (block, address)
  - i-node number
- Map `ascii` name to the *file descriptor*

Hierarchical Directory (Unix)

- Tree
- Entry:
  - name
  - inode number (try "ls –I" or "ls –iad .")
- example:

```
60  /usr/bob/mbox
```

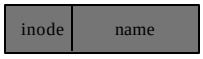

Unix Directory Example

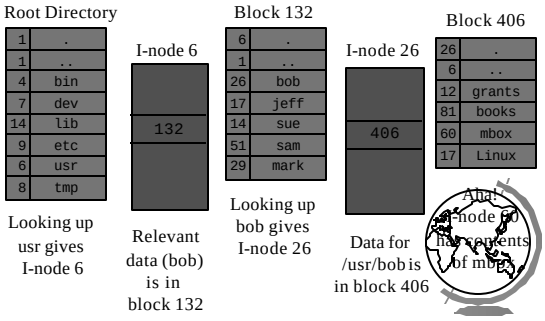

Outline

- Files (done)
- Directories (done)
- Partitions ¬

Outline

- Files (done)
- Directories (done)
- Disk space management ¬
- Misc

Partitions

- `mount`, `unmount`
  - load "super-block" from disk
  - pick "access point" in file-system
- Super-block
  - file system type
  - block size
  - free blocks
  - free I-nodes

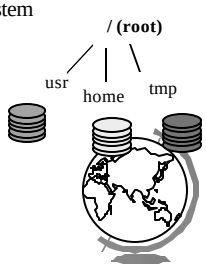

Partitions: `fdisk`

- Partition is large group of sectors allocated for a specific purpose
  - IDE disks limited to 4 physical partitions
  - logical (extended) partition inside physical partition
- Specify number of cylinders to use
- Specify type
  - magic number recognized by OS

(Hey, show example)

File System Maintenance

- Format:
  - create file system structure: super block, I-nodes
  - `format` (Win), `mke2fs` (Linux)
- "Bad blocks"
  - most disks have some
  - `scandisk` (Win) or `badblocks` (Linux)
  - add to "bad-blocks" list (file system can ignore)
- Defragment
  - arrange blocks efficiently
- Scanning (when system crashes)
  - lost+found, correcting file descriptors...