

Lecture 6: File Systems

Last Day:

- Space Management
- File Directories

Today: File Directory Structures

- Tree structures
- Generalized structures/shared files
- Reference Counters
- Garbage Collection

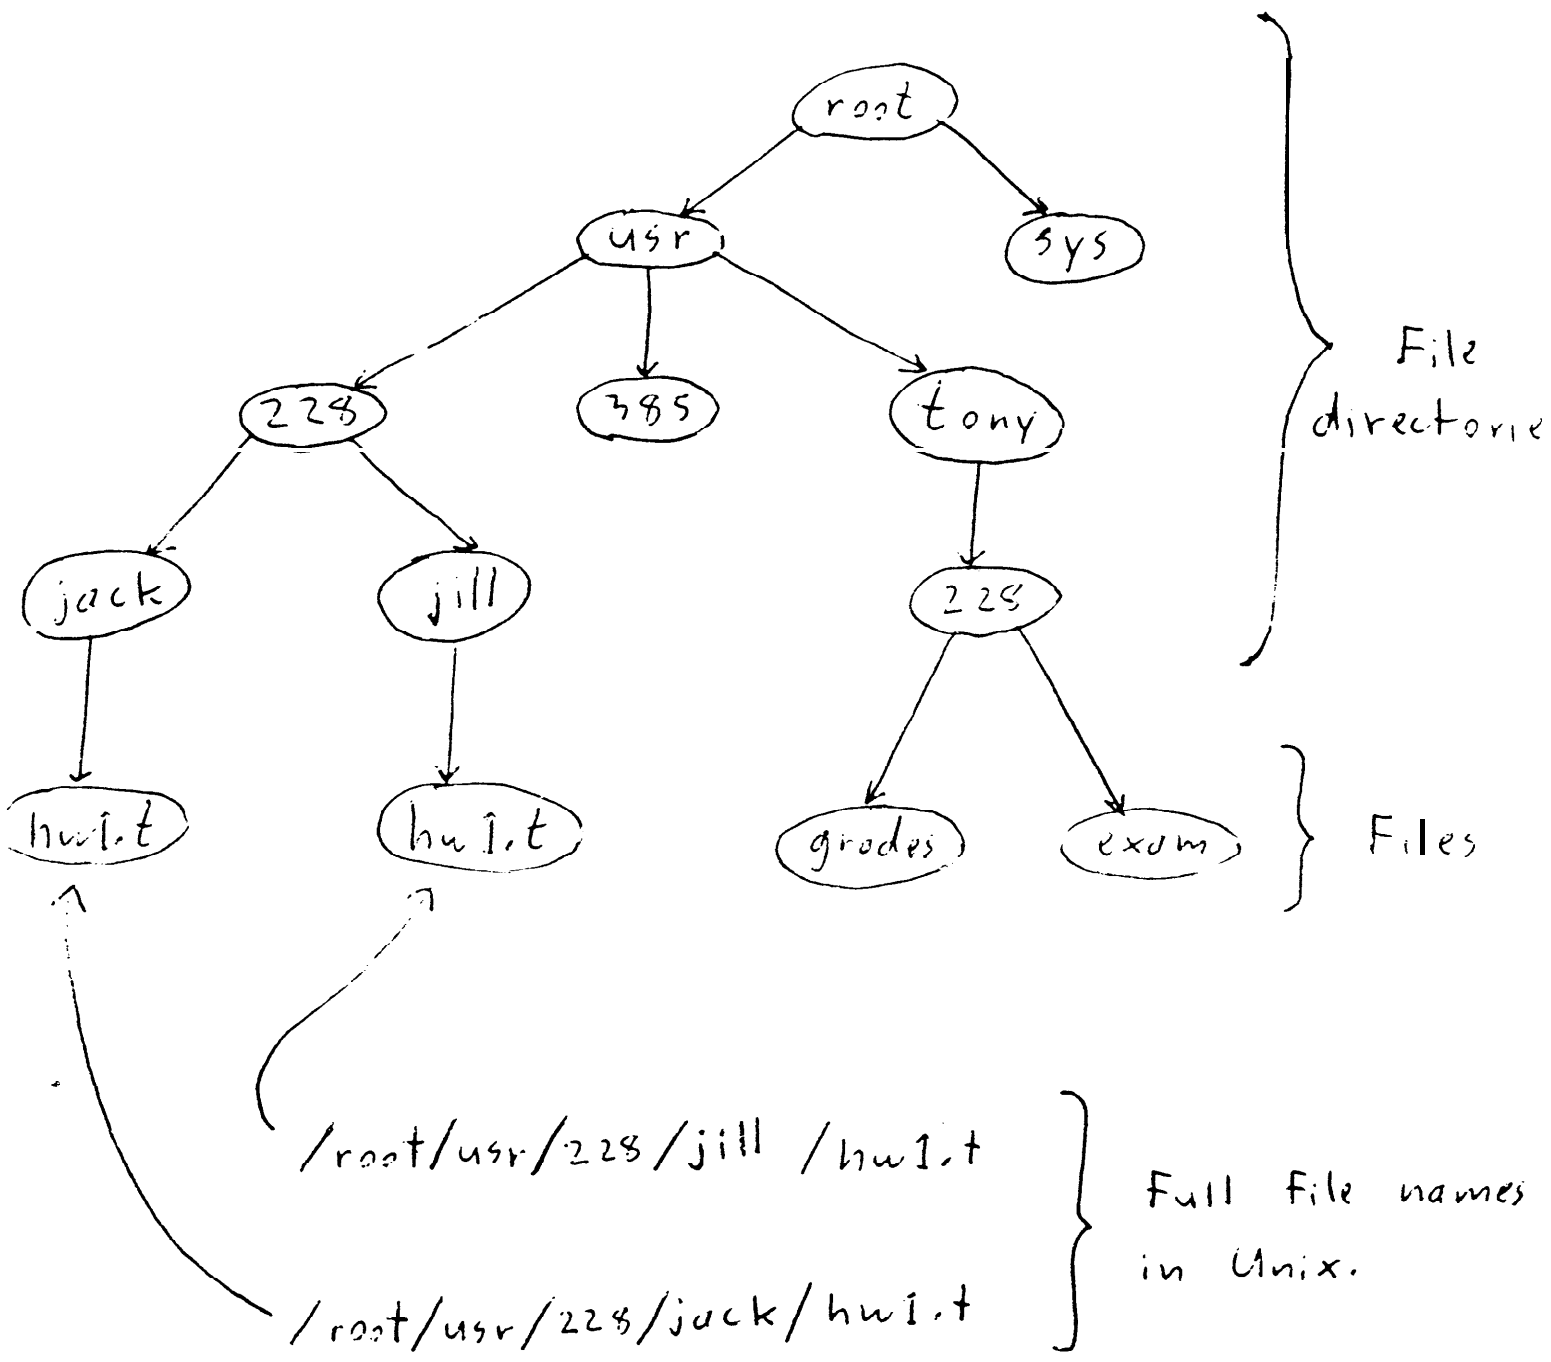
See handout on File systems.

• Folk & Zoelick § 2.7

File Directory Structures

- A single file directory does not allow different users to have files of the same name (eg. temp.c, hw1.t).
- ∴ Each user is given a separate file directory
- Tree-Structured Directories:
 - A file directory may have subdirectories, which may have sub-subdirectories, etc.

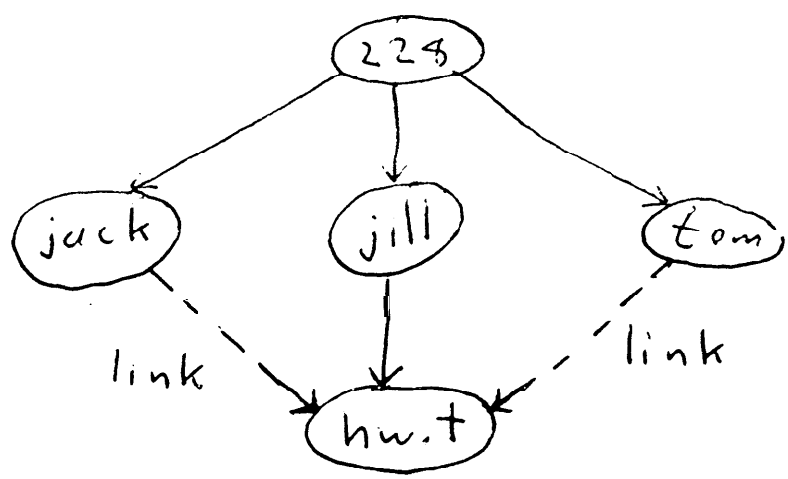
Example



Generalized Directory Structures

- A tree-structured directory does not allow users to share files.
- i.e., A file can be in only one user's directory at a time.
- Links allow users to share files.

eg.



It appears to jack that hw.t is a file in his directory. Likewise for Tom.

- In UNIX, the command

In `/228/jill/hw.t`

creates a (hard) link to the file hw.t
in the directory `/228/jill/`

- To see the link, type `"ls -l"` in UNIX.


- Note: This is different from storing separate copies of one file.

- Copied files require extra disk space.


- When one copy is changed, other copies are not changed.

- In contrast, a shared file is stored only once, and all users linked to the file see all changes to it.

Deleting Shared Files



- If user₁ deletes "grades", does user₂ still see the file, or not?
- In most systems, "delete file" means delete the link to the file.



This is the result of user₁ deleting grades.

- When the last link to a file is deleted, the system deletes the file too.
- Problem 1: How to determine when there are no more links to a file.


- Solution 1: Reference Counters.

For each file, maintain a counter of the number of references (links) to the file.


- When a link to the file is created, increment the counter.
- When a link is removed, decrement the counter.
- When the counter is zero, delete the file.

Problem 2: Cycles

eg



- Suppose Fred deletes here and there:



- here and there both have a reference count of 1 but they cannot be accessed, and should be deleted (returned to the free pool).

- Unfortunately, the file system has lost track of the files here and there!

Solution 2: Garbage Collection

- Each disk block is given a special bit.
- The block is said to be "marked" when the bit is 1, and "unmarked" when it is 0.
- When the free pool is low or empty, the following steps are executed:

Garbage Collection

Step 1: Scan the entire disk, and unmark every block (i.e., set its bit to 0).

Step 2: Starting at the root directory, traverse the directory structure, and mark every block of every file (i.e., set its bit to 1).

Note: Any blocks left unmarked are free (i.e., not pointed to by any directory).

Step 3: Scan the entire disk, and add any unmarked blocks to the free pool.

Note: Garbage collection is expensive
(since it scans the entire disk).

A Practical Compromise: Use both techniques.

- Cycles are rare, so reference counters will work for a long time.
- Periodically (eg, once a week, at night), run a garbage collection to sweep any unreferenced cycles into the free pool.