UNIVERSITY OF CALIFORNIA AT BERKELEY
College of Engineering
Department of Electrical Engineering and Computer Sciences
Computer Science Division

CS 162
Fall 1988

John Ousterhout

Topic 17: Directories

- Readings for this topic:
  - Silberschatz/Peterson: Section 9.4-9.5.
  - The Unix implementation paper.

- Naming: how do users refer to their files? How does OS find file, given name?

- The file descriptor information has to be stored on disk, so it will stay around even when the OS doesn't.
  - In Unix, all the descriptors are stored in a fixed- size array on disk. The descriptors also contain protection and accounting information.
  - Special areas of disk are used for this
    - Originally: descriptor array at one side of disk.
    - Then: descriptor array mid-way across disk.
    - Today: many small descriptor arrays spread across disk, so descriptors can be near to file data.
  - The sizes of the descriptor arrays are determined when the disk is initialized, and can't be changed. In Unix, the descriptor is called an *i-node*, and its index in the array is called its *i-number*. Internally, the OS uses the i-number to refer to the file.
  - When a file is open, its descriptor is kept in main memory. When

the file is closed, the descriptor is stored back to disk.

- Users need a way of getting back to files that they leave around on disk. One approach is just to have users remember descriptor indexes. *This is the Social Security approach.*

- Of course, users want to use text names to refer to files. Special disk structures called *directories* are used to tell what descriptor indices correspond to what names.

- Approach #1: have a single directory for the whole disk. Use a special area of disk to hold the directory.
  - Directory contains <name, index> pairs.
  - If one user uses a name, no-one else can.
- Many personal computers work this way.

- Approach #2: have a separate directory for each user (TOPS-10 approach). This is still clumsy: names from a user's different projects get confused.

- Unix approach: generalize the directory structure to a tree.
  - Directories are stored on disk just like regular files (i.e. file descriptor with 14 pointers, etc.) except file descriptor has special flag bit set. User programs can read directories just like any other file (try it!). Only special system programs may write directories.
  - Each directory contains <name, fd index> pairs in no particular order. The file pointed to by the index may be another directory. Hence, get hierarchical tree structure. Names have slashes separating the levels of the tree.

    - There is one special directory, called the *root*. This directory has no name, and is the file pointed to by descriptor 2 (descriptors 0 and 1 have other special purposes).
    - *Go through lookup example: /a/b/c.*

- It is possible for more than one directory entry to refer to a single file (‘‘hard links’’). UNIX uses reference counts in the file descriptors to keep track of the directory entries, only delete file when last directory entry goes away.

- Other things kept in UNIX file descriptors: file size, access times, owner and group id, protection bits.

- It is very nice that directories and file descriptors are separate, and that directories are implemented just like files. This simplifies the implementation and management of the structure (can write ‘‘normal’’ programs to manipulate them as files).

- Working directory: it is cumbersome constantly to have to specify the full path name for all files.
    - In Unix, there is one directory per process, called the working directory, which the system remembers.
    - When it gets a file name, it assumes that the file is in the working directory. ‘‘/’’ is an escape to allow full path names.
    - Many systems allow more than one current directory. For example, check first in A, then in B, then in C. This set of directories is called the *search path* or *search list*. This is convenient when working on large systems with many different programmers in different areas.
    - For example, in Unix the shell will automatically check in

several places for programs. However, this is built into the shell, not into Unix, so if any other program wants to do the same, it has to rebuild the facilities from scratch. Should be in the OS.
  - This is yet another example of locality.

- *Mention that C-shell implements ‘‘~’’?*

- Symbolic links: a file whose contents are just another file name. Also stored on disk just like regular files, but with a special flag set in descriptor.

- *Lecture 34 ended here.*