CSE 421/521 - Operating Systems
Fall 2011

LECTURE - XVIII

FILE SYSTEMS

Tevfik Koşar

University at Buffalo
November 3rd, 2011

File Systems

- Provides organized and efficient access to data on secondary storage:

  1. Organizing data into files and directories and supporting primitives to manipulate them (create, delete, read, write etc)
  2. Improve I/O efficiency between disk and memory (perform I/O in units of blocks rather than bytes)
  3. Ensure confidentiality and integrity of data

  – Contains file structure via a File Control Block (FCB)
    – Ownership, permissions, location..

A Typical File Control Block

| |
|---|
| file permissions |
| file dates (create, access, write) |
| file owner, group, ACL |
| file size |
| file data blocks or pointers to file data blocks |

Directories

➤ <u>Directories are special files that keep track of other files</u>
  - ✓ the collection of files is systematically organized
  - ✓ first, disks are split into partitions that create logical volumes (can be thought of as "virtual disks")
  - ✓ second, each partition contains information about the files within
  - ✓ this information is kept in entries in a **device directory** (or volume table of contents)
  - ✓ the directory is a symbol table that translates file names into their entries in the directory
    - ▪ it has a logical structure
    - ▪ it has an implementation structure (linked list, table, etc.)

Directories

➢ <u>Single-level directory structure</u>

- ✓ simplest form of logical organization: one global or **root** directory containing all the files
- ✓ problems
  - ▪ global namespace: unpractical in multiuser systems
 - ▪ no systematic organization, no groups or logical categories of files that belong together

directory	cat	bo	a	test	data	mail	cont	hex	records

files

Single-level directory

Directories

➢ <u>Two-level directory structure</u>

- ✓ in multiuser systems, the next step is to give each user their own private directory
- ✓ avoids filename confusion
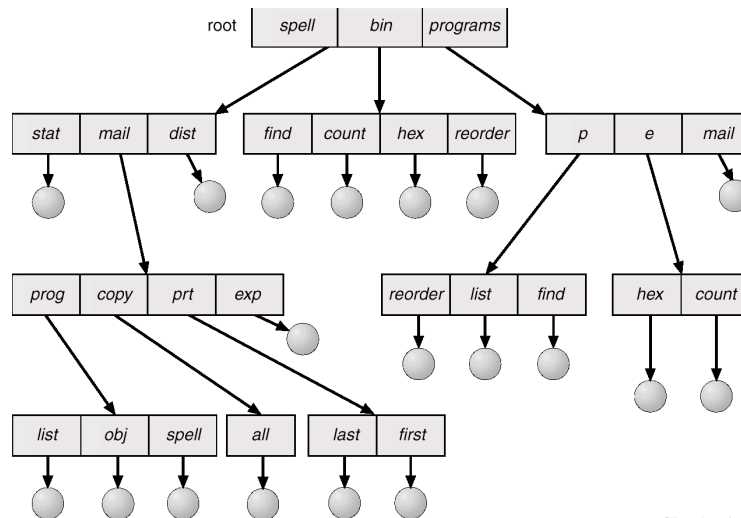- ✓ however, still no grouping: not satisfactory for users with many files

master file directory

user 1	user 2	user 3	user 4

user file directory

cat	bo	a	test	a	data	a	test	x	data	a

Two-level directory

Directories

➢ <u>Tree-structured directory structure</u>

Tree-structured directory

Silberschatz, A., Galvin, P. B. and Gagne. G. (2003)
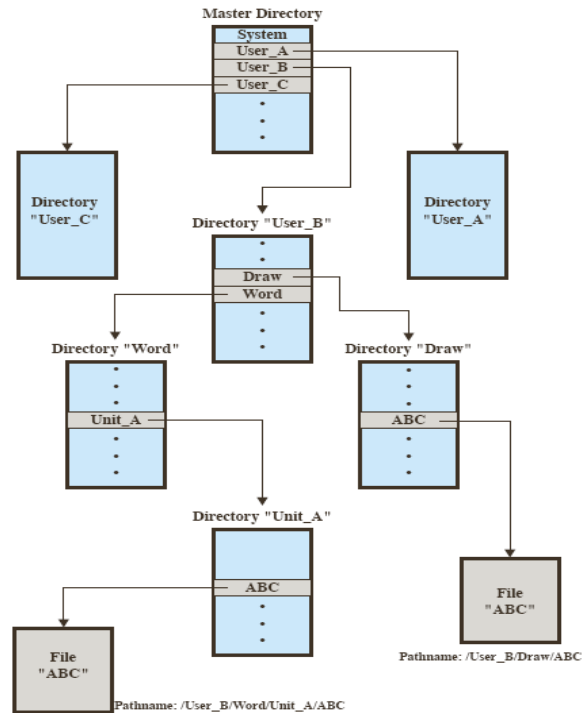*Operating Systems Concepts with Java (6th Edition).*

7

Directories

➢ <u>Tree-structured directory structure</u>

- ✓ natural extension of the two-level scheme
- ✓ provides a general hierarchy, in which files can be grouped in natural ways
- ✓ good match with human cognitive organization: tendency to categorize objects in embedded sets and subsets
- ✓ navigation through the tree relies on **pathnames**
  - ▪ absolute pathnames start from the root, example: /jsmith/ academic/teaching/cs446/assignment4/grades
  - ▪ relative pathnames start at from a current **working directory**, example: assignment4/grades
  - ▪ the current and parent directory are referred to as . and ..

8

Directory Implementation

Master Directory

System
User_A
User_B
User_C

Directory "User_C"

Directory "User_B"
Draw
Word

Directory "User_A"

Directory "Word"
Unit_A

Directory "Draw"
ABC

Directory "Unit_A"
ABC

File "ABC"

File "ABC"

Pathname: /User_B/Word/Unit_A/ABC

Pathname: /User_B/Draw/ABC

Stallings, W. (2004) *Operating Systems: Internals and Design Principles (5th Edition)*.

9

Directory Implementation

- **Linear list** of file names with pointer to the data blocks.
  - simple to program
  - time-consuming to execute

- **Hash Table** – linear list with hash data structure.
  - decreases directory search time
  - **collisions** – situations where two file names hash to the same location
  - fixed size

UNIX Directories

- Directory is a special file that contains list of names of files and their inode numbers
- to see contents of a directory:

```
$ls -1ia .
9535554 .
9535489 ..
9535574 .bash_history
9535555 bin
9535584 .emacs.d
9535560 grading
9535803 hw1
9535571 test
9535801 .viminfo
```

Example inode listing

```
$ ls -iaR demodir/y
865 .      193 ..     277 a      520 c      491 y

demodir/a/y:
277 .      865 ..     402 x

demodir/c:
520 .      865 ..     651 d1         247 d2

demodir/c/d1:
651 .      520 ..     402 xlink

demodir/c/d2:
247 .      520 ..     680 xcopy
```

Directories - System View

- user view vs system view of directory tree
  - representation with "dirlists (directory files)"
- The real meaning of "A file is in a directory"
  - directory has a link to the inode of the file
- The real meaning of "A directory contains a subdirectory"
  - directory has a link to the inode of the subdirectory
- The real meaning of "A directory has a parent directory"
  - ".." entry of the directory has a link to the inode of the parent directory

---

User View vs System View

Consider the following directory structure (user view):

Assume **mydir** (10), **a** (20), and **b** (30) are directories and **x** (40), **y** (50), and **z** (60) are files with inode numbers given in parenthesis. The inode number for mydir's parent directory is 1.

1) Please show the system representation (system view) of this directory tree.

Link Counts

- The kernel records the number of links to any file/directory.

- The *link count* is stored in the inode.

- The *link count* is a member of *struct stat* returned by the *stat* system call.

Change Links

- What will be the resulting changes in directory tree?

```
$ cp mydir/x mydir/b

$ ln mydir/a/z mydir/b/t

$ mv mydir/x mydir/a
```

Implementing "pwd"

1. "." is 247
   chdir ..

2. 247 is called "d2"
   "." is 520
   chdir ..

3. 520 is called "c"
   "." is 865
   chdir ..

4. 865 is called "demodir"
   "." is 193
   chdir ..

17

Allocation Methods

- An allocation method refers to how disk blocks are allocated for files:

- **Contiguous allocation**

- **Linked allocation**

- **Indexed allocation**

Contiguous Allocation

- Each file occupies a set of contiguous blocks on the disk

- + Simple – only starting location (block #) and length (number of blocks) are required

- - Wasteful of space (dynamic storage-allocation problem - fragmentation)

- - Files cannot grow

Contiguous Allocation of Disk Space

directory

file	start	length
count	0	2
tr	14	3
mail	19	6
list	28	4
f	6	2

Linked Allocation

- Each file is a linked list of disk blocks: blocks may be scattered anywhere on the disk.

block = | pointer |

+ Simple – need only starting address
+ Free-space management system – no waste of space
+ Defragmentation not necessary
- No random access
- Extra space required for pointers
- Reliability: what if a pointer gets corrupted?

Linked Allocation

File-Allocation Table

directory entry

test	• • •	217

name   start block

0	
217	618
339	
618	339

FAT

Indexed Allocation

- Brings all pointers together into the *index block,* to allow random access to file blocks.
- Logical view.

index table

+ Supports direct access
+ Prevents external fragmentation
- High pointer overhead --> wasted space

Example of Indexed Allocation

directory

file	index block
jeep	19

19

9
16
1
10
25
−1
−1
−1

# Free Space Management

- Disk space limited
- Need to re-use the space from deleted files
- To keep track of free disk space, the system maintains a free-space list
  - Records all free disk blocks
- Implemented using
  - Bit vectors
  - Linked lists

Free-Space Management (Cont.)

- Bit vector   (*n* blocks)
  - Each block is represented by 1 bit
  - 1: free, 0: allocated

$$0 \quad 1 \quad 2 \qquad\qquad\qquad n-1$$

$$\ldots$$

$$\text{bit}[i] = \begin{cases} 1 \Rightarrow \text{block}[i] \text{ free} \\ 0 \Rightarrow \text{block}[i] \text{ occupied} \end{cases}$$

- e.g. 000011111000100010010000

Free-Space Management (Cont.)

- Linked List Approach

Free-Space Management (Cont.)

- Bit map requires extra space
  - Example:

    block size = 2^{12} bytes

    disk size = 2^{30} bytes (1 gigabyte)

    $n = 2^{30}/2^{12} = 2^{18}$ bits (or 32K bytes)
- Easy to get contiguous files
- Linked list (free list)
  - Cannot get contiguous space easily
  - requires substantial I/O
- Grouping
  - Modification of free-list
  - Store addresses of n free blocks in the first free block
- Counting
  - Rather than keeping list of n free addresses:
    - Keep the address of the first free block
    - And the number n of free contiguous blocks that follow it

---

Free-Space Management (Cont.)

- Linked List

free–list head

grouping ($n=3$)

1	2,3,7
7	8,9,12
12	14,15,16
16	17,18,–1

bit map: 011100011100101111100000

counting: (1,3), (7, 3), (12, 1), (14, 5)

Exercise

directory

file	start	length
count	0	2
tr	14	3
mail	19	6
list	28	4
f	6	2

Acknowledgements

- "Operating Systems Concepts" book and supplementary material by A. Silberschatz, P. Galvin and G. Gagne

- "Operating Systems: Internals and Design Principles" book and supplementary material by W. Stallings

- "Modern Operating Systems" book and supplementary material by A. Tanenbaum

- R. Doursat and M. Yuksel from UNR