

## File Systems

1

## Long-term Information Storage

1. Must store large amounts of data
2. Information stored must survive the termination of the process using it
3. Multiple processes must be able to access the information concurrently

December 10, 2002 CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

2

## File Naming


Extension	Meaning
file.bak	Backup file
file.c	C source program
file.gif	CompuServe Graphical Interchange Format image
file.hlp	Help file
file.html	World Wide Web HyperText Markup Language document
file.jpg	Still picture encoded with the JPEG standard
file.mp3	Music encoded in MPEG layer 3 audio format
file.mpg	Movie encoded with the MPEG standard
file.o	Object file (compiler output, not yet linked)
file.pdf	Portable Document Format file
file.ps	PostScript file
file.tex	Input for the TEX formatting program
file.txt	General text file
file.zip	Compressed archive

### Typical file extensions.

December 10, 2002 CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

3

## File Structure



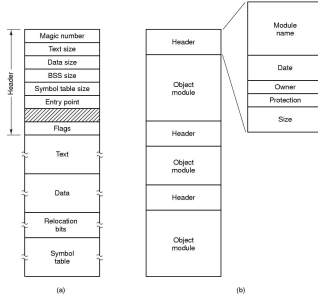
### Three kinds of files

- byte sequence
- record sequence
- tree

December 10, 2002 CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

4

## File Types



(a) An executable file (b) An archive

December 10, 2002 CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

5

## File Access

- Sequential access
  - read all bytes/records from the beginning
  - cannot jump around, could rewind or back up
  - convenient when medium was mag tape
- Random access
  - bytes/records read in any order
  - essential for data base systems
  - read can be ...
    - move file marker (seek), then read or ...
    - read and then move file marker

December 10, 2002 CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

6

## File Attributes

Attribute	Meaning
Protection	Who can access the file and in what way
Password	Password needed to access the file
Creator	ID of the person who created the file
Owner	Current owner
Read-only flag	0 for read/write; 1 for read only
Hidden flag	0 for normal; 1 for do not display in listings
System flag	0 for normal files; 1 for system file
Archive flag	0 for has been backed up; 1 for needs to be backed up
ASCII/binary flag	0 for ASCII file; 1 for binary file
Random access flag	0 for sequential access only; 1 for random access
Temporary flag	0 for normal; 1 for delete file on process exit
Lock flags	0 for unlocked; nonzero for locked
Record length	Number of bytes in a record
Key position	Offset of the key within each record
Key length	Number of bytes in the key field
Creation time	Date and time the file was created
Time of last access	Date and time the file was last accessed
Time of last change	Date and time the file has last changed
Current size	Number of bytes in the file
Maximum size	Number of bytes the file may grow to

### Possible file attributes

December 10, 2002

CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

7

## File Operations

1. Create
2. Delete
3. Open
4. Close
5. Read
6. Write
7. Append
8. Seek
9. Get attributes
10. Set Attributes
11. Rename

December 10, 2002

CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

8

## An Example Program Using File System Calls (1/2)

```

/* File copy program. Error checking and reporting is minimal. */

#include <sys/types.h>          /* include necessary header files */
#include <fcntl.h>
#include <stdlib.h>
#include <unistd.h>

int main(int argc, char *argv[]); /* ANSI prototype */

#define BUF_SIZE 4096           /* use a buffer size of 4096 bytes */
#define OUTPUT_MODE 0700       /* protection bits for output file */

int main(int argc, char *argv[])
{
    int in_fd, out_fd, rd_count, wt_count;
    char buffer[BUF_SIZE];

    if (argc != 3) exit(1);      /* syntax error if argc is not 3 */

```

December 10, 2002

CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

9

## An Example Program Using File System Calls (2/2)

```

/* Open the input file and create the output file */
in_fd = open(argv[1], O_RDONLY); /* open the source file */
if (in_fd < 0) exit(2);          /* if it cannot be opened, exit */
out_fd = creat(argv[2], OUTPUT_MODE); /* create the destination file */
if (out_fd < 0) exit(3);        /* if it cannot be created, exit */

/* Copy loop */
while (TRUE) {
    rd_count = read(in_fd, buffer, BUF_SIZE); /* read a block of data */
    if (rd_count <= 0) break; /* if end of file or error, exit loop */
    wt_count = write(out_fd, buffer, rd_count); /* write data */
    if (wt_count <= 0) exit(4); /* wt_count <= 0 is an error */
}

/* Close the files */
close(in_fd);
close(out_fd);
if (rd_count == 0) /* no error on last read */
    exit(0);
else /* error on last read */
    exit(5);
}


```

December 10, 2002

CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

10

## Memory-Mapped Files



(a) Segmented process before mapping files into its address space


(b) Process after mapping  
existing file *abc* into one segment  
creating new segment for *xyz*

December 10, 2002

CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

11

## Directories: Single-Level Directory Systems



- A single level directory system
  - contains 4 files
  - owned by 3 different people, A, B, and C

December 10, 2002

CS 3204: Operating Systems, Fall 2002  
© Mir Farooq Ali, 2002

12

## Two-level Directory Systems

Letters indicate *owners* of the directories and files

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 13

## Hierarchical Directory Systems

A hierarchical directory system

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 14

## Path Names

A UNIX directory tree

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 15

## Directory Operations

1. Create
2. Delete
3. Opendir
4. Closedir
5. Readdir
6. Rename
7. Link
8. Unlink

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 16

## File System Implementation

A possible file system layout

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 17

## Implementing Files (1)

(a) Contiguous allocation of disk space for 7 files  
 (b) State of the disk after files D and E have been removed

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 18

### Implementing Files (2)

Storing a file as a linked list of disk blocks

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 19

### Implementing Files (3)

Linked list allocation using a file allocation table in RAM

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 20

### Implementing Files (4)

An example i-node

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 21

### Implementing Directories (1)

(a) A simple directory  
fixed size entries  
disk addresses and attributes in directory entry

(b) Directory in which each entry just refers to an i-node

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 22

### Implementing Directories (2)

Two ways of handling long file names in directory

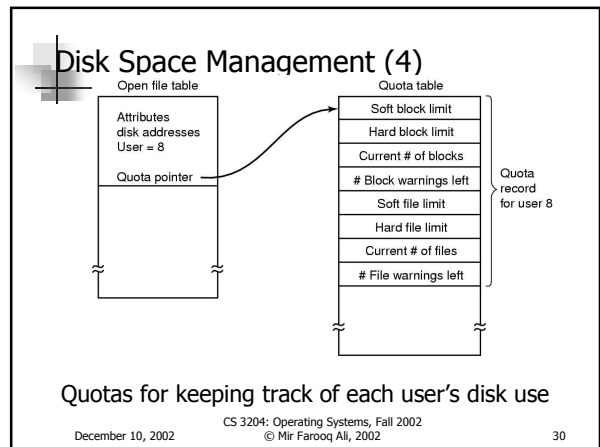
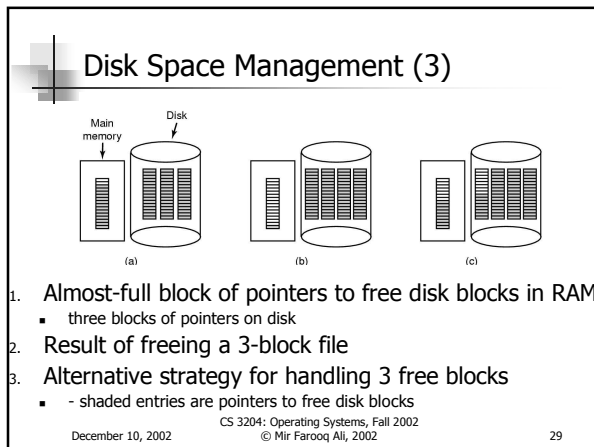

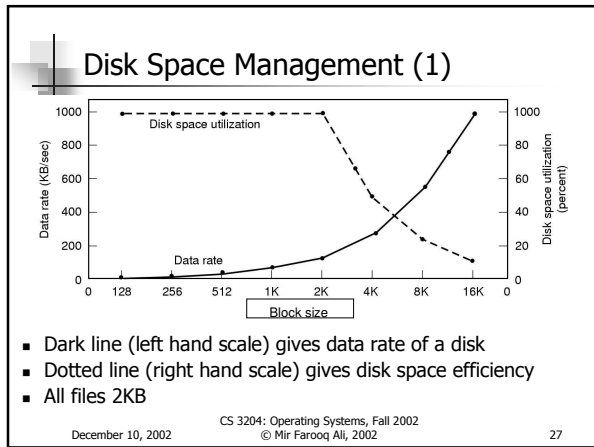

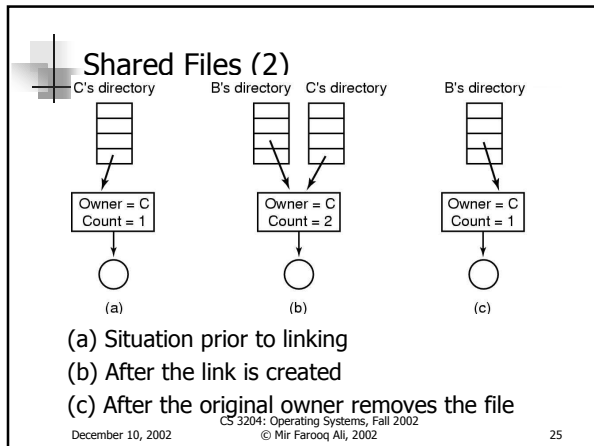
- (a) In-line
- (b) In a heap

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 23

### Shared Files (1)

File system containing a shared file

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 24



### File System Performance (1)

The diagram illustrates the block cache data structures. On the left, a vertical column of boxes represents a hash table. Arrows from these boxes point to a horizontal sequence of boxes representing cache blocks. The first block is labeled 'Front (LRU)' and the last is 'Rear (MRU)'. Curved arrows show the movement of blocks between the hash table and the cache, and between adjacent blocks in the list.

The block cache data structures

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 31

### File System Performance (2)

Diagram (a) shows a disk layout where i-nodes are clustered at the beginning of the disk. Diagram (b) shows a disk divided into concentric cylinder groups, with each group containing its own i-nodes.

- I-nodes placed at the start of the disk
- Disk divided into cylinder groups, each with its own i-nodes

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 32

### Log-Structured File Systems

- With CPUs faster, memory larger
  - disk caches can also be larger
  - increasing number of read requests can come from cache
  - thus, most disk accesses will be writes
- LFS Strategy structures entire disk as a log
  - have all writes initially buffered in memory
  - periodically write these to the end of the disk log
  - when file opened, locate i-node, then find blocks

December 10, 2002 CS 3204: Operating Systems, Fall 2002 © Mir Farooq Ali, 2002 33