



1B11 Operating Systems

Files

Prof. Steve R Wilbur
s.wilbur@cs.ucl.ac.uk




1B11-4 1999



Lecture Objectives

- 1 What properties do we need from files?
- 1 Why does the OS implement files?
- 1 How is protection managed?



1B11-4 1999 Slide 2

Why have Files?

- 1 Up to 1970s disks were rare and small (1.2MB disk satisfied annual teaching needs of UCL CS in 1975)
- 1 Users “owned” disks and directly allocated space (sectors and cylinders) for data storage - primitive *files*
- 1 They kept paper records (*catalogues*) of where data was stored
- 1 As use became more common, catalogues were kept on disk - primitive *directories*

1B11-4 1999 Slide 3

Why have Files - 2?

- 1 Nowadays, disks are large, e.g.:
 - o laptop now has 4-10 Gigabytes
 - o UCL-CS has > 400 Gigabytes of server file-store
- 1 User files are mostly small - statistics vary with business, but:
 - o many files < 1 kbyte,
 - o some at 10-100s kbytes
 - o few of > 10 Mbyte
- 1 But, *file size* needs increase over time:
 - o Reason 1 -
 - o Reason 2 -

1B11-4 1999 Slide 4

Why have Files - 2?

- 1 User needs:
 - o Easy way to store and retrieve data over long periods
 - o Lots of files (Unix philosophy: "files are cheap")
 - o Various sizes and can change during life-time
 - o Virtual files - e.g. "larger" than disk/file system
 - o Ability to control access to files - restrict/share

1B11-4 1999 Slide 5

Why have Files - 3?

- 1 OS needs:
 - o To present a clean, maintained abstraction
 - o To organise disk storage and allocate it efficiently
 - o To track *ownership* and *access rights/permissions*
 - o In context of multi-user environment
 - o These are main reasons why OS has to manage the *file system*
 - o OS is also a user:
 -] long-term tables (e.g. passwords) stored as files
 -] OS needs to load programs

1B11-4 1999 Slide 6

Disk Organisation

- | OS provides access to I/O devices such as disk via *device drivers* (software) which deal with competing I/O demands from the various processes
- | Typically they offer "*raw*" access to an abstract disk that looks like an array of blocks of a fixed size, often 1 or 4 or 16kbytes - "chosen" when disk is *formatted*
- | Access is near random (i.e. similar time to get any given block)

1B11-4 1999 Slide 7

Disk Organisation - 2

- | Users may want to "buy" certain area for own use
- | We may want to divide a disk into logical regions too - called "*partitions*"

- | Single machine may have many disks
- | How do we manage this scale?
- | ... Hierarchy of directories/folders
- | So we have:
 - o Machine:Disk:Partition:Directory:File:Block

1B11-4 1999 Slide 8

Disk/File Organisation

Computer System

- | Partition 1
- | Partition 2
- | Partition 3

- | Partitions are usually arranged as cylinders
- | Minimises head movement when partition being used

1B11-4 1999 Slide 9

Disk/File Organisation

Partition/Disk

- | Partion will have *root* or *master* directory
- | File system will spring from this as hierarchy
- | Directories contain *names* of files (and sub-directories) and other *attributes*
- | File is a *sequence of bytes*
- | Implemented by *mapping* blocks of partition to these byte addresses
- | Disk space *not allocated* if no data at that file address

File
Block 14
Block 897
Block 356
Empty
Block 694
Block 665

1B11-4 1999 Slide 10

File Store Implementation

- 1 We'll treat partitions and disks interchangeably
- 1 Important data structures:
 - o *Free space list* - free blocks - may be **bit map** or free blocks are **chained**
 - o *Bad block list* - must be separate list! - identified when disk formatted, possibly later too
 - o *Directories* - relate naming, protection and other info to files
 - o *File maps* - indicate where blocks of file are on disk
- 1 Directories and file maps often kept together
- 1 Free space list and root directory in fixed place

1B11-4 1999 Slide 11

File Store Implementation - 2

- 1 So, how are files created, extended, etc?
- 1 Unix and DOS separate directory from file map
 - o DOS: File Allocation Table - FAT - also contains free space list
 - o Unix: Free list and I-nodes

1B11-4 1999 Slide 12

File Store Implementation - 3

DOS-like

The diagram illustrates a DOS-like file store implementation. On the left, a **Directory** contains entries for files **A** and **B**. File **A** is associated with block number **6**, and file **B** is associated with block number **5**. In the center, the **FAT** (File Allocation Table) is shown as a vertical array of 9 entries (indices 0-8). Entry 2 is marked **EOF**, entry 4 is marked **EOF**, entry 8 is marked **Bad**, and entry 9 is marked **Free**. Other entries contain block numbers: 3 is **Free**, 5 is **2**, 6 is **8**, and 8 is **4**. On the right, **Files** are shown as green blocks. File **A** consists of three blocks of sizes 6, 8, and 4. File **B** consists of two blocks of sizes 5 and 2. Arrows indicate the mapping from directory entries to FAT entries and then to the physical blocks.

- | Directory contains address (block number) of start of file
- | FAT contains entry for each block of disk
- | Bad blocks flagged
- | Linked list of blocks making up a file
- | End of file marker at end (e.g. FFFF)

1B11-4 1999 Slide 13

File Store Implementation - 4

Unix-like

The diagram illustrates a Unix-like file store implementation. It shows a **Free space list** and **Current Alloc'n block**. The **Free list head** points to a green block. The **Current Alloc'n block** is a yellow block. Below, three orange blocks represent allocated space. The first block starts at address **498** and contains addresses 498, 673, 674, ..., 1052, 1053. Its **link** field contains **499**. The second block starts at address **499** and contains addresses 499, ..., 2076, 2077, 2089. Its **link** field contains **500**. The third block starts at address **500** and contains addresses 500, ..., 2345, 2347. Its **link** field contains **0**. Arrows show the linked list structure: the free list head points to the first block, and each block's link field points to the next block.

- | One example shown
- | Free blocks used to hold free space list
- | Block's own address is last to be freed in its set - i.e. it itself gets used
- | When disk is formatted bad blocks are excluded from the free list

1B11-4 1999 Slide 14

File Store Implementation - 5

Unix-like

- | A contains file attributes, e.g. owner
- | B - access to 10 file blocks
- | 1 - further 256 blocks accessed by indirect block
- | 2 - further 256*256 blocks
- | 3 - further 256*256*256 blocks
- | **Indirect block** scaling related to block size and word size

1B11-4 1999 Slide 15

File Store Implementation - 6

- | DOS
 - o Simpler structure
 - o File size < disk size
 - o Good sequential access
 - o Poor random access
 - o File "holes" take up disk space

- | Unix
 - o Separates free space and file structures
 - o Sequential or random access good performance
 - o Max. 3 disk accesses for address of any part of file
 - o Holes may take no space

1B11-4 1999 Slide 16

User Interface to Files

- | The user has an *Applications Programming Interface (API)* to the file system
 - o Unix and C - standard I/O library
 - o Java VM by a class library
- | API makes *system calls* (traps) to operating system, e.g.:
 - o *open(), close(), read(), write(), seek(), unlink()*

1B11-4 1999 Slide 17

File System Performance

- | Disk access is 1000s times slower than memory
- | OS keeps copies of file blocks in use in its own memory (*buffers* or *cache*)
- | *Principle of Locality* so OS needs to manage read/write of blocks to disk and buffer replacement policy
- | Same principle suggests high likelihood of sequential file access
- | Unix uses this to *read-ahead* next block so ready when (if) user needs it

1B11-4 1999 Slide 18

File System Performance - 2

- | To avoid loss of data due to power glitches OS periodically *flushes* buffers to disk
- | Also need cache write-through policy to ensure file updates not lost

1B11-4 1999 Slide 19

Directory Entries

- | A directory entry on disk might hold:
 - o file name + disk address (I-node in Unix)
 - o attributes of the file:
 -] owner
 -] size, type
 -] protection information, e.g. rwx for owner, group, others
 -] creation time, last access or modified time
- | Few systems (except MACs) put the type in the directory - elsewhere:
 - o implicit in the name/extension, by convention
 - o derivable by looking at part of the file
 - o e.g. .txt, .c, .psd, .doc, ...

1B11-4 1999 Slide 20

Kinds of Files

- | Most OSs nowadays support two kinds of file:
 - o Array of bytes
 - o Indexed sequential
- | Windows and Unix only support the first
- | Programmers need to enhance the “array of byte” model - but this is easy...
- | Just write another file as the index (there are Unix application programmer’s libraries to help you do just this)

1B11-4 1999 Slide 21

File System Differences

- | The basic ideas in MAC-OS, Unix and Windows file systems are very similar nowadays, but there are a lot of detailed differences - some include:
- | maximum file or disk or partition size supported
 - o model of user id, permissions
 - o file type information
 - o file name restrictions.

1B11-4 1999 Slide 22

Backup/Archive of Files...

- 1 Disks are not 100% reliable, nor are people - files get lost
- 1 Backup systems periodically copy all changed files (incremental) to tape (or CD or Zip drive)
- 1 Less frequently, do full “dump” of file system, and then store tapes somewhere different (possibly 2 places)
- 1 UCL CS has tapes going back to early '80s and can in theory retrieve any file from any day since then.....in theory:-)

1B11-4 1999 Slide 23

Summary

- 1 Files implemented by OS as convenient way of centrally managing disk space
- 1 File system needs to deal flexibly with range of file sizes and give good performance
- 1 Also, in memory caching used to improve performance
- 1 Basic byte-array file can be extended by libraries in user space

1B11-4 1999 Slide 24