CSC 828

Lecture 4: File Systems

Last two days: Secondary Storage Devices:

- Magnetic Disk
- Magnetic Tape
- Buffering

Today: File Systems:

- Overview
- Transparency
- Space Management
- Free Pool

Zoellick + Folk § 3.4, 3.5, 2.1

+ handout

Overview of Database Management Systems

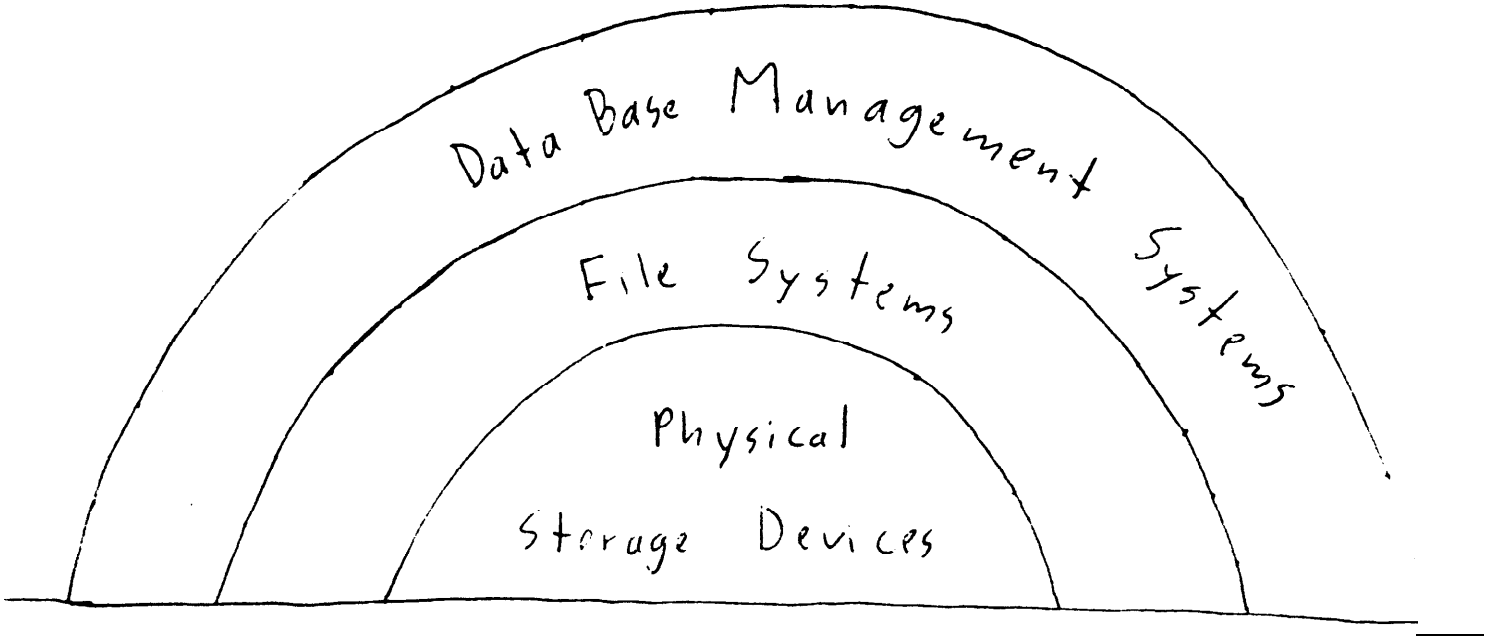

Each level hides details of the lower levels.

Physical Storage Devices (disk, tape)

- Data on disks is addressed by Surface, Sector, & Track.
- One block of data is retrieved at a time.
- eg. Retrieve the block at Surface 3, Sector 7, Track 89.

File Systems (eg. Unix)

- Data is addressed by record number.
- One record is retrieved at a time.
- eg. Retrieve record 276.

Database Management Systems (eg. Oracle, Ingress)

Data is accessed by field values.

Specified fields are returned.

- eg. Retrieve the name & address of each student in CSC 228.

File Systems

- A File system is a service supported by the operating system (O.S.) that allows users to create, read, write and delete Files

- Often, other file operations are also supported eg, copy, edit, rename, ...

- The File system hides (abstracts) the details of secondary storage, so that users need not think about blocks, disk addresses, cylinders, read/write heads, interblock gaps, etc.

File Storage

- Files can be stored on disk or tape.

- These days, tapes are used mostly for archival purposes, ie, for storing files that are rarely (if ever) needed.

- For day-to-day operations, files are stored on disks (or similar devices) since they provide quick, direct access to any block of data.

- Recall: Blocks on disk are addressed by three numbers: Surface, Sector, Track.

- Thus, disk address space can be thought of as a 3-dimensional array.

File Organization

- A file is stored on disk as a list of data blocks.
- These blocks may be scattered randomly all over disk, in no particular order.
- For quick access, though, they should all be stored on one cylinder, or on adjacent cylinders.
- A file system hides all these details from the user.

Transparency

- To a user, a file is just a sequence of records.

- The user is unaware of what blocks his file is stored in. He does not even know that blocks exist!

- The file system manages the details. It divides the file into blocks, and assigns each block to a location on disk.

- Note: Records that are close together in a file may be stored far apart on disk (though the file system tries to avoid this, for efficiency).

Space Management

- Many Files may be stored on disk.
- Each disk block is either allocated to a File, or it is free (not allocated to any File)
- The set of free blocks is called the Free pool
- Initially, all blocks are Free.
- As more & more Files are created, the free pool shrinks.
- As Files are removed (deleted), allocated blocks become free, and the free pool expands

A Snapshot of Disk Allocation

eg.

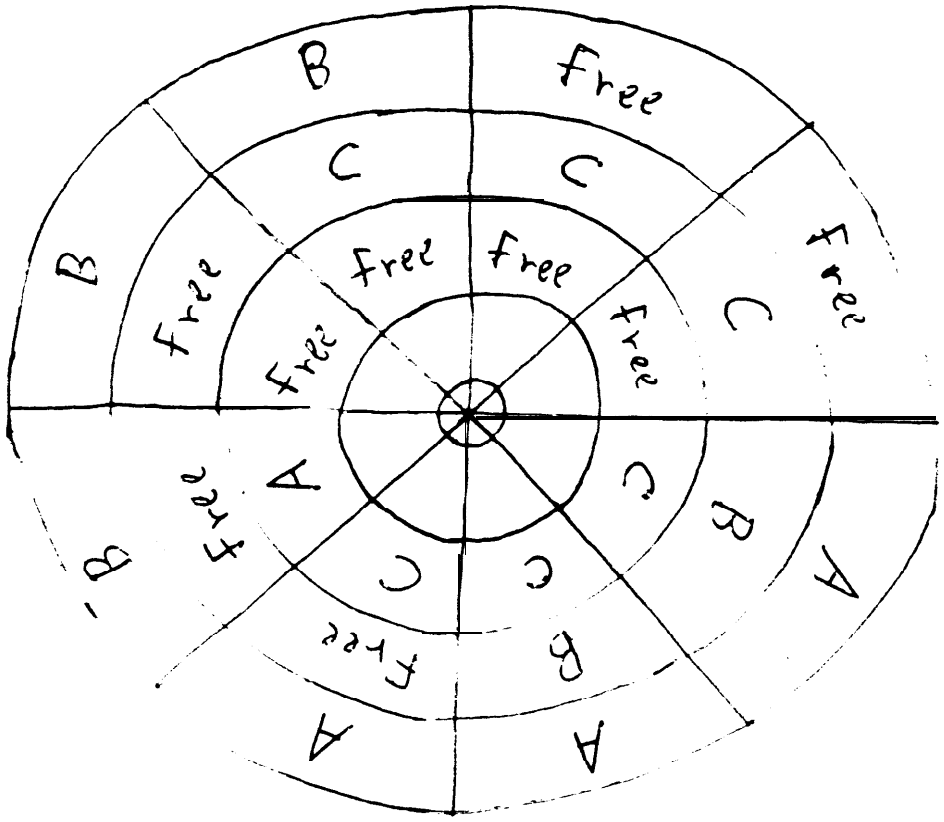

3 Files: A, B, C.

Space Management

The file system must:

(A) keep track of free blocks (the free pool)

(B) Allocate space for new files
(which shrinks the free pool)

(C) Return blocks to the free pool
when users delete files

(A) Managing the Free Pool

There are three main techniques:

1. Linked list of free blocks.
2. Linked list of segments.
3. Bit map

① Linked List of Free Blocks

- Each free block points to the next free blk.

eg.

- Note: A "pointer" is a few bytes at the front or end of a block that gives the disk address of the next block in the list.

- The address of the first block is kept in a special place in main memory.

- A more typical picture:

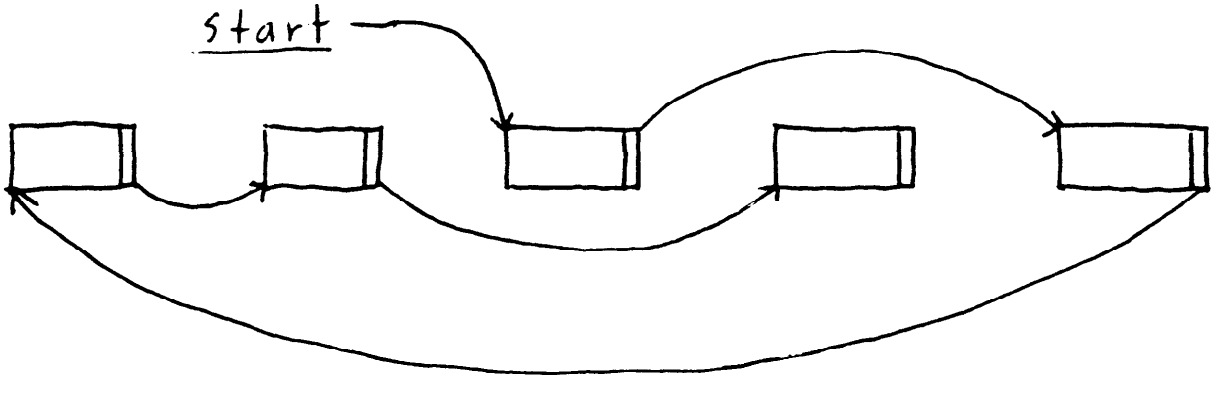

② Linked List of Segments

A segment is a sequence of contiguous (ie, adjacent) blocks on disk.

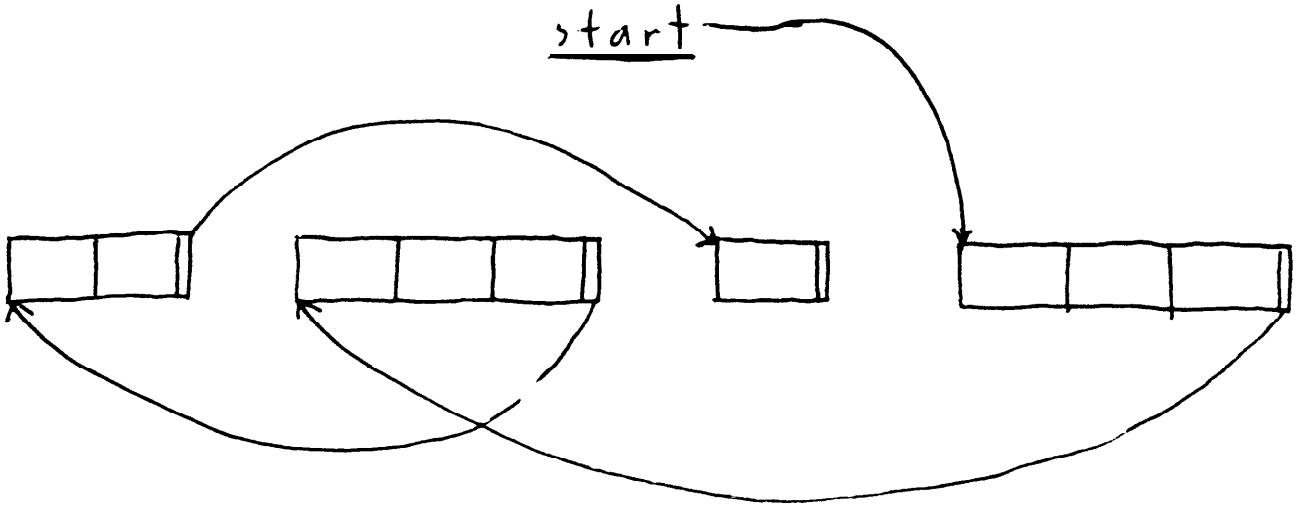

③ Bitmap: Keep an array of bits, one bit for each block.

- Bit (ijk) is 0 iff the block at disk address (ijk) is free.