University of Dublin
Trinity College

Storage Hardware

Owen.Conlan@cs.tcd.ie

Hardware Issues

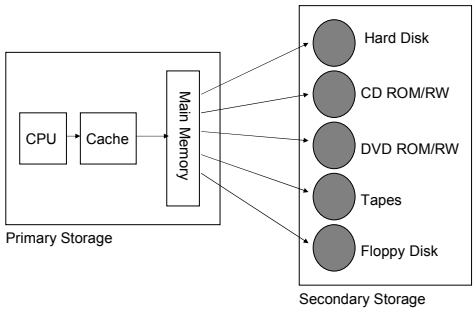

Hardware Issues

Primary Storage is …

- Limited
- Volatile
- Expensive

However, it is also …

- Fast (May be accessed directly from the CPU)

Hardware Issues

Secondary Storage is …

- Extendible
- Persistent
- Cheap

However, it is …

- Relatively slow (must be copied to main memory before being accessed by the CPU)

Hardware Issues

Why do we use secondary storage?

- Primary storage (RAM) costs more than disk space
- We like to switch our computers off and on again

But secondary storage is very slow

- Retrieving a single character from RAM takes about 150 nanoseconds (150 billionths of a second)
- Retrieving the same character from disk takes about 75 milliseconds (thousandths of a second)
- 75 msec is 500,000 times longer than 150 ns.

Disk Storage Devices

Disk Storage Devices …

- Direct Access Storage
  - *As opposed to Tape drives, which are serial devices*
- Offer high storage capacity and low cost
- Data stored as magnetized areas on magnetic platters surfaces
- Each disk has one or more platters
- A disk pack contains several magnetic platters connected to a rotating spindle

Disk Pack with Read/Write H/W

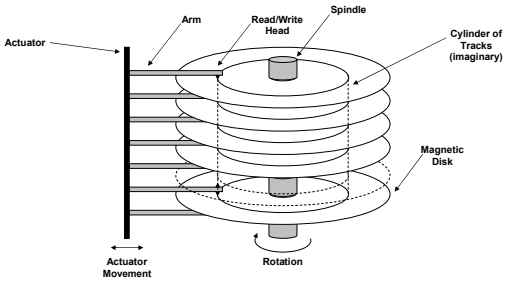

Movable vs Fixed-head Disks

Some disks have fixed-heads

- As many read/write heads as there are tracks on the platter
- Track is selected electronically and is therefore much faster
- Cost of additional read/write heads is the limiting factor to production

Disks with an actuator are called moveable-head disks

- Actuator moves the (single) read/write head per platter to the appropriate track

Tracks and Sectors

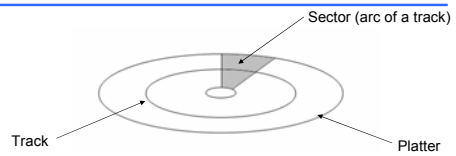

- Disks are divided into concentric circular tracks on each platter surface
- Track capacities vary typically from 4 to 50 Kbytes
- The division of the disk into sectors is hard-coded and may not be changed
- Subtended at fixed angle at the centre of a platter are sectors (figure above)
- Not all disks have their tracks divided into sectors

Cylinder

Tracks directly above and below one another form a cylinder

All information on a cylinder can be accessed without moving the arm (called seeking) that holds the read/write heads

- A cylinder consists of a group of tracks
- A track consists of a group of sectors
- A sector consists of a group of bytes

Estimating Capacities

Track capacity = #sectors per track * bytes per sector
Cylinder capacity = #tracks per cylinder * Track capacity
Drive capacity = #cylinders * Cylinder capacity

Knowing these relationships allows us to compute the amount of disk space a file is likely to require

- How many cylinders to store a file with 20,000 fixed length records of 256 bytes each on a disk with –
  - *512 bytes per sector*
  - *40 sectors per disk*
  - *11 tracks per cylinder*

Estimating Capacities

Answer –

- The file is 20,000 * 256 = 5,120,000 bytes
- 5,120,000 / 512 = 10,000 sectors
- 10,000 / 40 = 250 tracks
- 250 / 11 = 22.72 cylinders

1f 22.72 physically contiguous cylinders are not available then the file will have to be spread out over the disk (fragmentation)

Blocks

What are Blocks?

- A track is divided into blocks or pages
- Block size (generally) fixed for each operating system
- Typical block sizes range from 512 bytes to 4096 bytes
- A disk with hard-coded sectors often has the sectors further subdivided onto blocks

Why are they important?

- Whole blocks are transferred between disk and main memory for processing

Disk I/O

Input/Output (I/O) from/to a Disk …

- A read-write head moves to the track that contains the block to be transferred *(seek)*
- Disk rotation moves the block under the read-write head for reading and writing *(rotational delay or latency)*
- Operating System finds the correct block
  - *A physical disk block address consists of a surface number, track number (within surface), and block number (within track)*
- Entire block read/written from/to an area in RAM called a buffer
- Time taken to transfer the block (block transfer time)

Reading or writing a disk block is time consuming because of the seek times and rotational delay

- Double buffering can be used to speed up the transfer of contiguous disk blocks

Buffers

When several blocks need to be transferred from disk to main memory (and all the block addresses are known beforehand) several buffers can be reserved in main memory to speed the transfer

- CPU processes data in a buffer while another is being read/written

Facilitated by independent Input/Output processor

- Can transfer data block to/from main memory independent of and in parallel to CPU processor
- This technique is referred to as double buffering

Example of Double Buffering

Use of two buffers, A and B, for reading from disk

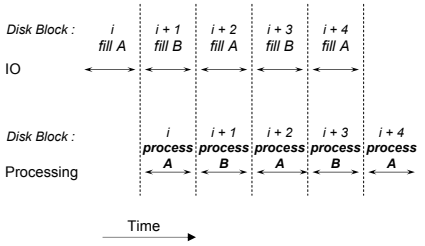

Physical and Logical Files

A disk drive may contain many physical files

- A program uses these files via the operating system, but knows nothing about where the data is on disk

A program sees only *logical files*

- When a program wants to access a file, the operating system provides the hook-up between logical and physical file
- The programmer does not need to know the physical location on disk, just the *logical* name creating

Sector Organisation

There are two basic ways to organise data on a disk – by sector organisation or user-defined block

We will study only sector organisation

The physical placement of sectors –

- You cannot read a series of sectors that are all in the same track one right after another!
- After reading the data, it takes the machine time to process the data before retrieving more
- So, if logically adjacent sectors were placed physically adjacent, retrieval would be slow
- Solution : Interleave the sectors

Sector Organisation - Clusters

A cluster is a fixed number of contiguous sectors

- All clusters on a disk are the same size
- To view a file as a series of clusters and still maintain the sectored view, the file manager ties the logical sectors to the physical clusters using a file allocation table (FAT)

FAT -

- The FAT contains a linked list of all the clusters in a file, ordered according to the logical order of the sectors in a cluster
- With each entry in the FAT is an entry giving the physical location of the cluster

Review

Compared to RAM disk accesses are slow, but cheap!

Disks facilitate persistent storage of data

Disks enable direct access

Anatomy of a Disk

- Platters, tracks, cylinders, sectors, blocks

Physical components of a Disk

- Magnetic disks, actuator, arm, read/write heads, spindle

Estimating Capacities

Disk I/O and Buffering

Sector Organisation and Clustering