Chapter 6 - External Memory

Luis Tarrataca
luis.tarrataca@gmail.com

CEFET-RJ

Table of Contents I

Table of Contents I

Motivation

Computers have a set of internal memory mechanisms

- registers, RAM, etc

Computers also need to interact with peripherals

- Some of these are storage devices;
- *A.k.a* external memory.

How does such an external memory work?

How is such an external memory organized?

Some examples of external memory include:

- Magnetic disks;
- Redundant Array of Independent Disks (RAID);
- Solid State Drives (SSD);
- Optical Memory;
- Magnetic tapes.

Lets take a closer look at some of these.

Magnetic Disks

Circular platter coated with a magnetizable material.

Figure: Interior of a magnetic hard drive

Data are recorded on and later retrieved from the disk via a head:

- Most common design: a read head and a write head;
- During a read or write operation:
  - Head is stationary while the platter rotates beneath it.

Write Mechanism

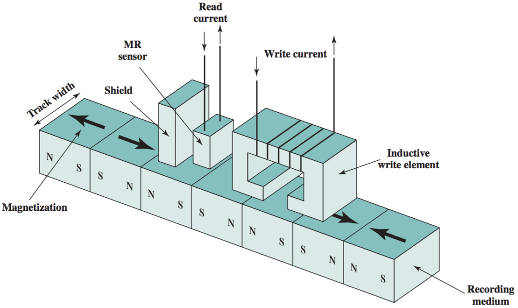

Figure: Inductive Write / Magnetoresistive Read Head. (Source: (Stallings, 2015))

Write Mechanism:

- Electric pulses are sent to the write head;
- Resulting magnetic patterns are recorded on the surface below:

Read Mechanism

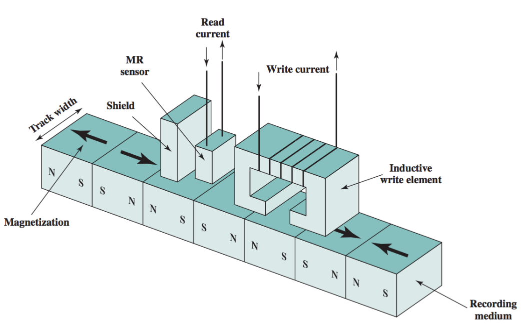

Figure: Inductive Write / Magnetoresistive Read Head. (Source: (Stallings, 2015))

Read Mechanism:

- Disk surface passes under the read head;
- Generating a current of the same polarity as the one recorded.

Contemporary Read Mechanism

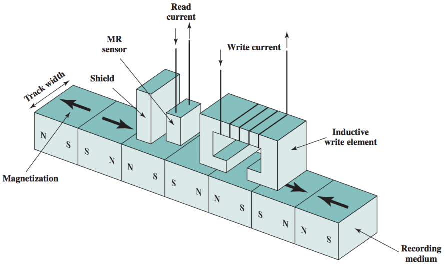

Figure: Inductive Write / Magnetoresistive Read Head. (Source: (Stallings, 2015))

Contemporary disk systems use a separate read head:

- Head consists of a magnetoresistive (MR) sensor;
- Resistance of MR material depends on the direction of the magnetization;
- Sensor detects resistance changes as voltage signals.

Data Organization and Formatting

Based on the read / write head mechanism how are disks organized? Any ideas?

Data Organization and Formatting

Based on the read / write head mechanism how are disks organized? Any ideas?

Impractical to have a large tape moving under the head:

- Old magnetic tapes;
- Substantial seek time times;

Circular Information Storing:

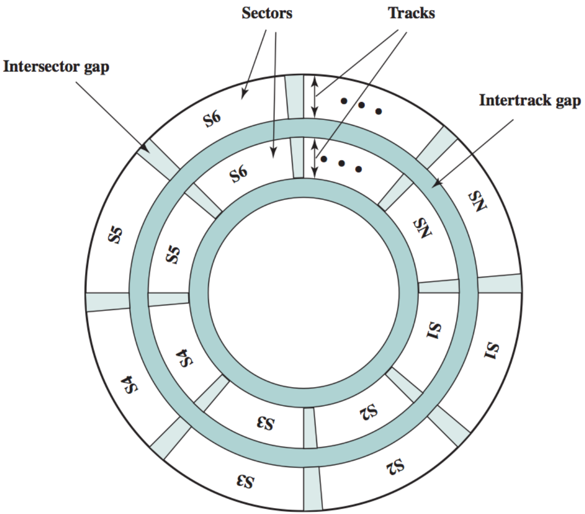

Figure: Disk Data Layout. (Source: (Stallings, 2015))

Each platter has a concentric set of rings, called **tracks**:

- Each track is the same width as the head.
- There are thousands of tracks per surface.
- Adjacent tracks are separated by gaps in order to prevent;
    - Misalignment of the head;
    - Magnetic field interference.

Data are transferred to and from the disk in **sectors**:

- There are typically hundreds of sectors per track;
- These may be of either fixed or variable length
  - Nowadays 512 bytes is the universal sector size.
- Adjacent sectors are separated by intersector gaps.

Now that we know about the existence of tracks and sectors:

How can the head find these elements within the disk?

How can the head find these elements within the disk?

Head needs to locate sector positions within a track, requiring knowing:

- Starting point on the track;
- Start and end of each sector;

How can the head find these elements within the disk?

Head needs to locate sector positions within a track, requiring knowing:

- Starting point on the track;
- Start and end of each sector;

How can this be performed? Any ideas?

How can this be performed? Any ideas?

- Location data needs to be recorded on the disk:
  - Disk is formatted with extra data;
  - This data is used only by the drive
    - Not accessible to the user.
  - Reason why:
    - Space seen as available by OS $\neq$ to physical disk space

Example (1/3)

Lets look at an example:

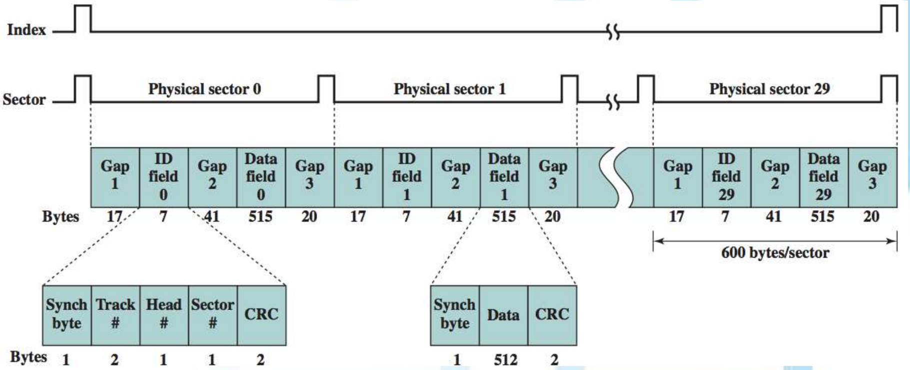

Figure: Winchester Disk Format. (Source: (Stallings, 2015))

Example (2/3)

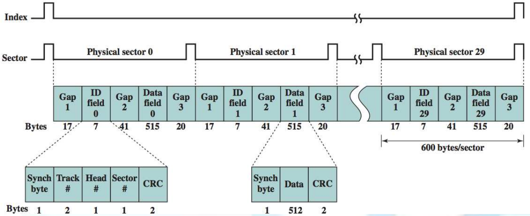

Figure: Winchester Disk Format. (Source: (Stallings, 2015))

Each track contains 30 fixed-length sectors of 600 bytes each.

- Each sector holds 515 bytes of data plus other control information;
- This means that only $515/600 \approx 85\%$ is available for data...

Example (3/3)

Lets look at a specific example:

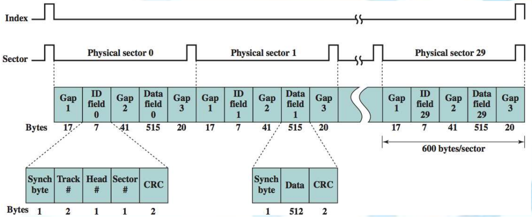

Figure: Winchester Disk Format. (Source: (Stallings, 2015))

The ID field uniquely identifies a sector containing:

- SYNCH byte is a special bit pattern that delimits the beginning of the field;
- Track number
- Head / Surface number for disks with multiple surfaces;
-

Components of a Disk Drive

One way of organizing the components of a disk drive:

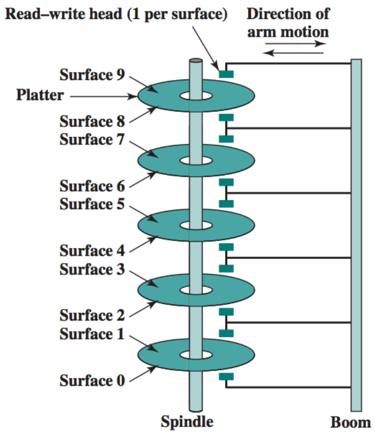

Figure: Components Of A Disk Drive. (Source: (Stallings, 2015))

Figure: Interior of a magnetic hard drive

Disk Performance Parameters

So, what do you think are some variables that influence the performance of a disk?

Disk Performance Parameters (1/4)

- **Seek time** - the time it takes to position the head at the track;
- **Rotational delay** - once the track is selected:
  - The sector still needs to line up with the head;
  - *E.g.:* an hard disk rotating at 20000 rpm:
    - 20000 rpm $\rightarrow$ one revolution per 3 ms;
    - On average we will have to wait for half the plate to spin;
    - rotational delay = 1.5 ms;
- **Access time** = SeekTime + Rotational Delay

Disk Performance Parameters (2/4)

- **Transfer time** - Data transfer portion of the operation;

$$T = \frac{b}{rN}$$

  - **b** - number of bytes to transfer;
  - **r** - rotation speed per second (rps)
  - **N** - number of bytes on a track;

Disk Performance Parameters (3/4)

- **Total average time**:

$$T = T_{\text{seek}} + T_{\text{rotational delay}} + T_{\text{transfer time}}$$

$$T = T_{\text{seek}} + \frac{1}{2r} + \frac{b}{rN}$$

Disk Performance Parameters (4/4)

There are also delays associated with I/O operation:

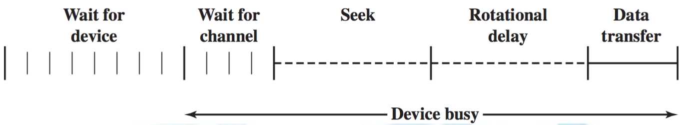

Figure: Timing of a disk I/O transfer (Source: (Stallings, 2015))

1. Wait for device:
    - A process must wait in a queue for a device to be available;
2. Eventually, the device is assigned to the process;
3. Wait for I/O channel:
    - If the device shares the I/O channel it must wait for it to be available;
4. Only after this point do we proceed with the head seek.

Characteristics	Constellation ES.2	Seagate Barracuda XT	Cheetah NS	Momentus
Application	Enterprise	Desktop	Network attached storage, application servers	Laptop
Capacity	3 TB	3 TB	400 GB	640 GB
Average seek time	8.5 ms read 9.5 ms write	N/A	3.9 ms read 4.2 ms write	13 ms
Spindle speed	7200 rpm	7200 rpm	10, 075 rpm	5400 rpm
Average latency	4.16 ms	4.16 ms	2.98	5.6 ms
Maximum sustained transfer rate	155 MB/s	149 MB/s	97 MB/s	300 MB/s
Bytes per sector	512	512	512	4096
Tracks per cylinder (number of platter surfaces)	8	10	8	4
Cache	64 MB	64 MB	16 MB	8 MB

Figure: Typical Hard Disk Driver Parameters (Source: (Stallings, 2015))

Where to focus your study

After this class you should be able to:

- Understand the key properties of magnetic disks.
- Understand the performance issues involved in magnetic disk access.
- Explain the concept of RAID, importance of redundancy and mechanisms for redundancy;
- Compare and contrast hard disk drives and solid disk drives.
- Describe in general terms the operation of flash memory.

Less important to know:

- details of all RAID levels;
- details of specific magnetic disk formats;

Your focus should always be on the building blocks for developing a solution =)

References I

Stallings, W. (2015).

Computer Organization and Architecture: Designing for Performance.

Pearson Education, 10th edition edition.