COMPUTER ARCHITECTURES

Devices, device controllers

Általános
INFORMATIKAI
Tanszékk

Schedule

- Devices, their classes, their architecture
- Controllers, adapters
- IT, their role and management
- Structured devices: disks, CDs, DVDs
  - architecture, basic concepts, access

Role of Devices

- **Connection**
  - with users (**Human Interface Devices**),
  - between machines and systems (**Networking Devices**),
  - to information sources and actuators (**detectors, actuators**).
- **Secondary data storage** (disks),
- **tertiary storage** (tapes, cassettes, CD, DVD),
- and **special devices** (e.g. watch).

Most common architecture

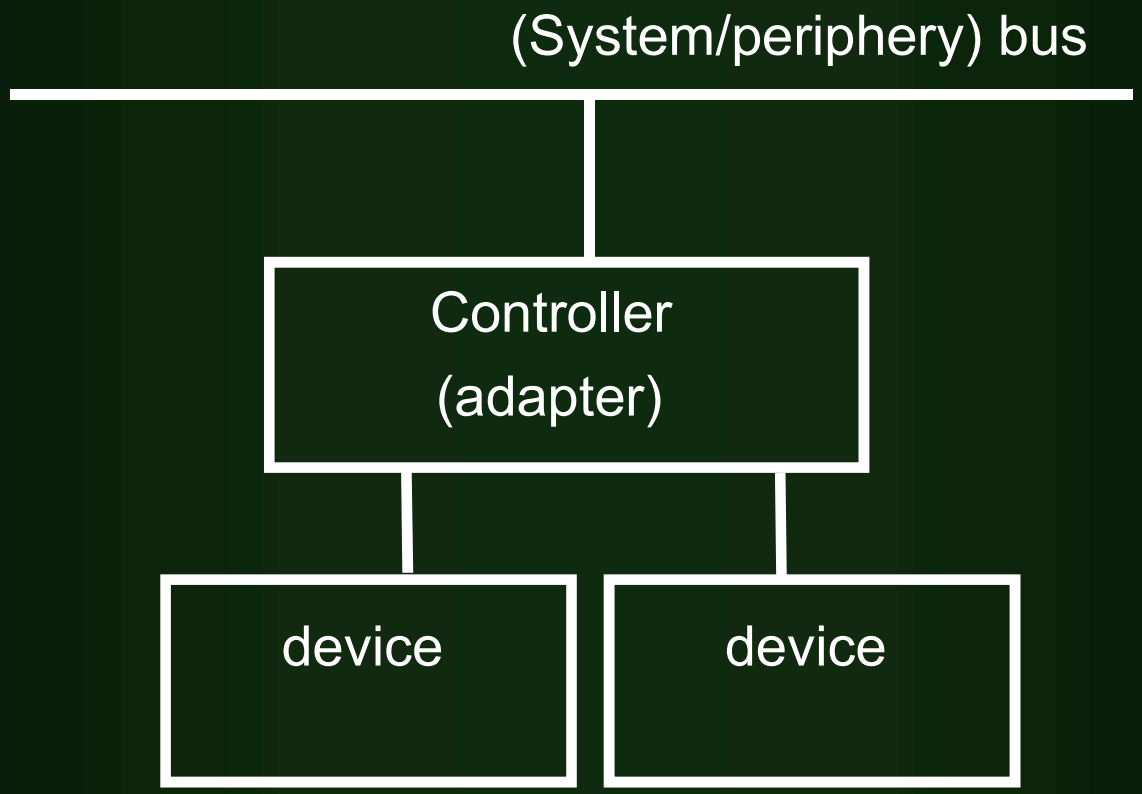

The controllers (adapters)

- Interface through the bus to the rest of the machine,
- sometimes ability to control the bus ,
- synchronization (IT generation).
- It emits signals to “move” the device(s),
- controlled data flow between the controller's buffer(s) and the device, error handling,

- Circuits, they can be on the motherboard or on a separate card. They can have registers and buffers.

"Programming" the controllers

- E.g. an imagined disk controller: at least 2 data (areg1,areg2), 1 control register (creg), block size buffer
- Script to import a block:
  - MOVE lba, areg1 # lba = what
  - MOVE mem, areg2 # mem = where
  - MOVE in, creg # in = the direction
  - controller works independently, # and at the end
  - SWIT # generates an interrupt.
- Where can I find the above code?
  - In BIOS, OS core routine (device driver).

The interrupt

- A signal to the CPU indicating the occurrence of an asynchronous event.
- When discussing the bus, we mentioned the bus that conducts the interrupt ...
- The CPU interrupts the current instruction flow (context saving) and the IT-dependent instruction sequence (the handler) is executed.
- After that, the original instruction flow continues (context recovery).

Frequent "events"

- **Clock device** interrupts (for setting time/date fields, counting time quantums, etc.)
- **Peripheries** (controllers indicating that they have completed some kind of transfer).
- Interrupts **caused by** other processes,
- **CPU mode changes** (trap),
- **error events.**
- (We don't separate IT and exceptions yet!)

Vector ITs, IT levels

- ITs identified by serial numbers and the addresses of their handlers in a vector table .
  - The IT serial number is also sent after/in addition to the IT indication
  - Polling is used to query who sent IT (?)
- IT priority levels: IT with a higher priority may interrupt the handling of a lower one, but
- after the service of the higher one, the serving of the a lower follows: pending interrupts can be queued (they are not lost).
- IT masking, IT blocking .

Asset classes

- **Structured (block-oriented) devices :**
  - disks, CDs, DVDs, cassettes, tapes, etc.
  - Blocks of data transfer, block addresses on the device ,
  - a file system can be created on them.
- **Unstructured (character-oriented) devices :**
  - terminals, printers, serial/parallel ports, etc.
  - Byte/character/line transfer ,
  - (in this case, the "row structure" does not count as “structure”).
- **Special devices** (e.g. the clock device)

Device drivers

- **Service routine package belonging to the core of the operating system (OS kernel), which**
- **makes the device manageable at a higher level.**
- **They implement the MOVEs, they "contain" the IT handlers.**
- **It will be detailed in the OS subject.**

Magnetic disk storage, disks

- Their purpose: secondary storage (file system, virtual memory).
- They are based on a change in magnetization: they do not forget when turned off.
- Magnetic signal recording:
  - the current creates a magnetic field, this can change the magnetization of a magnetizable material (signal recording);
  - a voltage is induced in the conductor in a changing magnetic field (the basis of the reading).
  - Magnetoresistive: resistance changes under the influence of a magnetic field
  - Hall effect: if current flows in a conductor or semiconductor and it is placed in a magnetic field, the electrons are affected by the Lorentz force, which causes a potential difference on both sides of the conductor (Hall voltage)

http://www.usbyte.com/common/HDD.htm

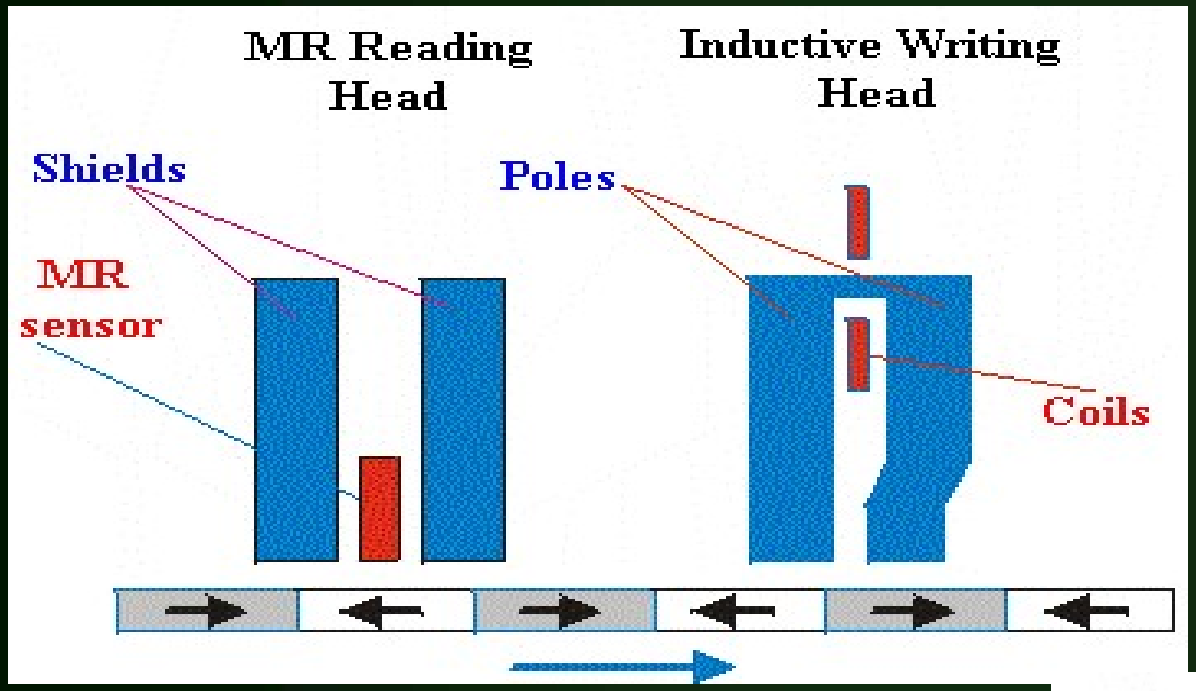

MR: Magneto Resistive

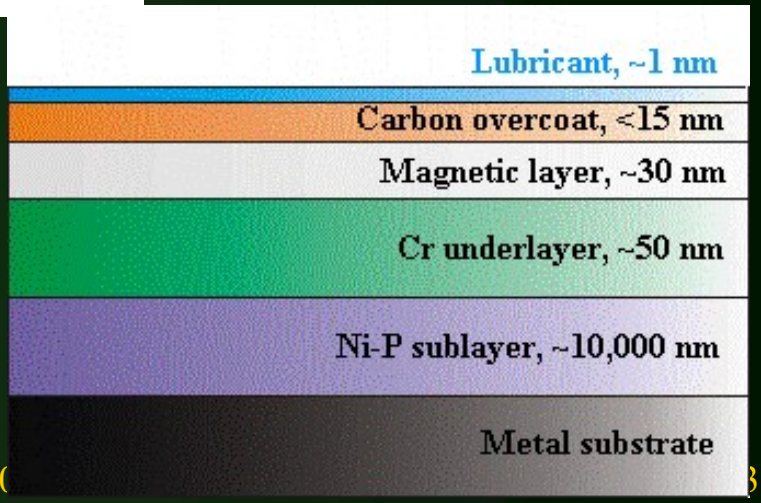

Structure

- **Disk sides - write/reade heads; the address of the sides (head address);**
- **tracks - a concentric circle on one side, at a given head position (radius); their addresses: (track, cylinder address, head location);**
- **sectors: a sector on a circle, with gaps between them; they have addresses.**
- **Cylinder : overlapping tracks nn disk sides, all available with one head position.**

The addresses of the tracks are "written" (as magnetic marks) at the head of the sectors.
And at the end of the sectors there is a checksum (Error Correction Code, ECC).
These aare written during the disc initialization.

Magnetic disk storage

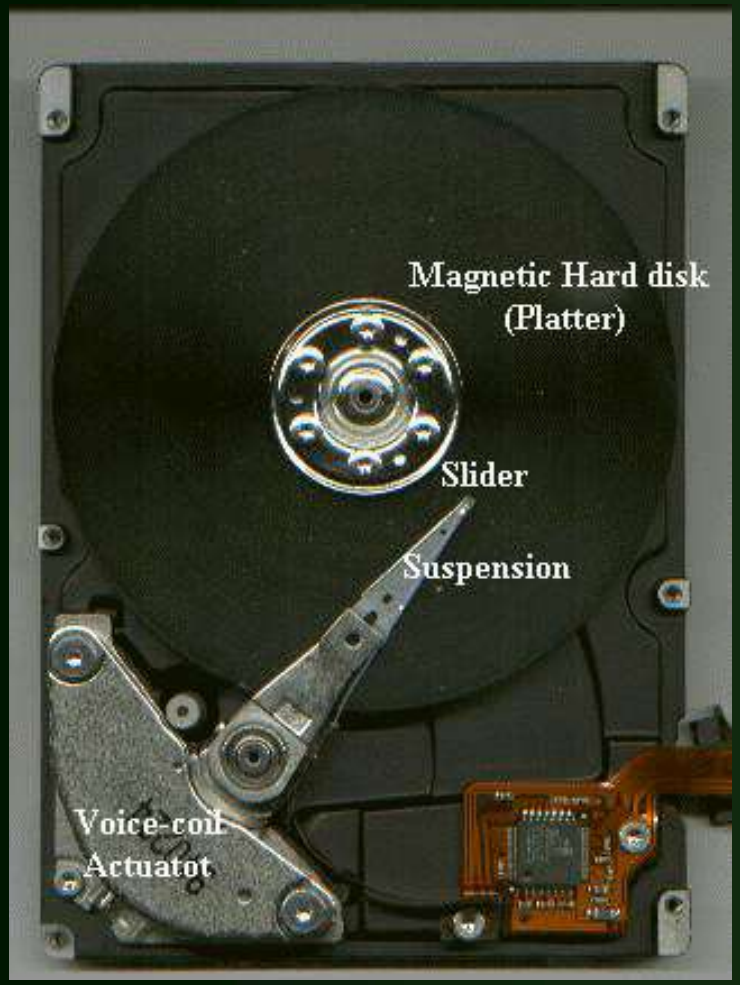

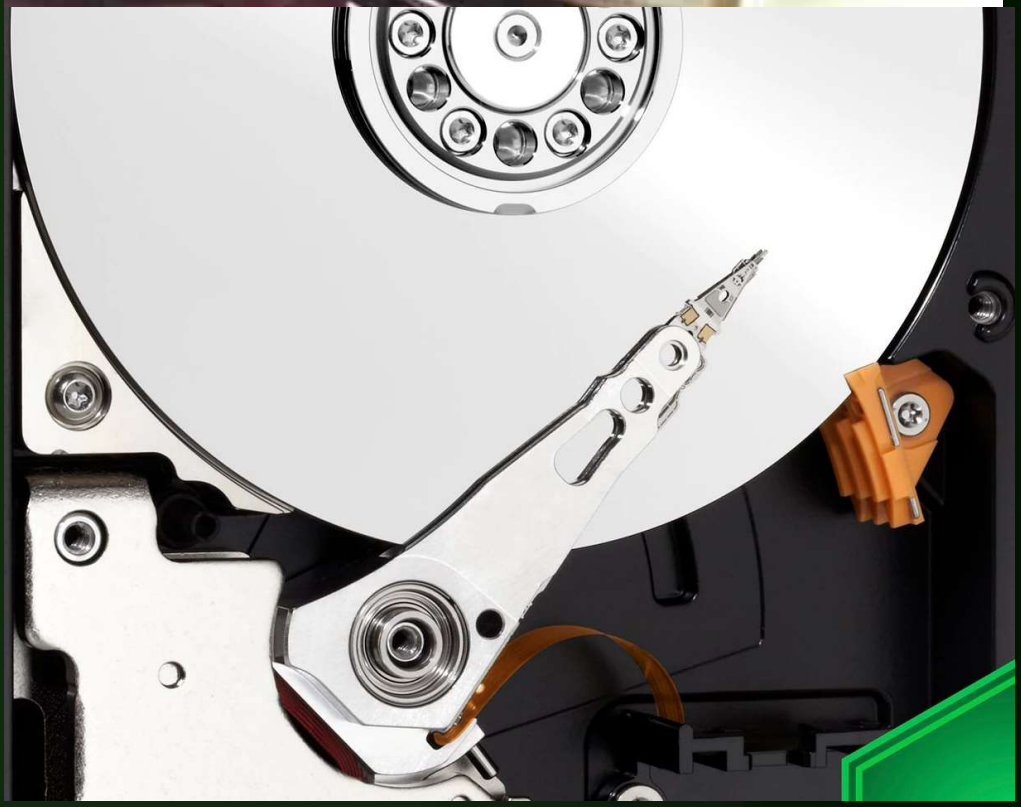

http://www.usbyte.com/common/HDD.htm

The concepts of page, band, sector...

The flight height of the head is a few tens of nanometers.

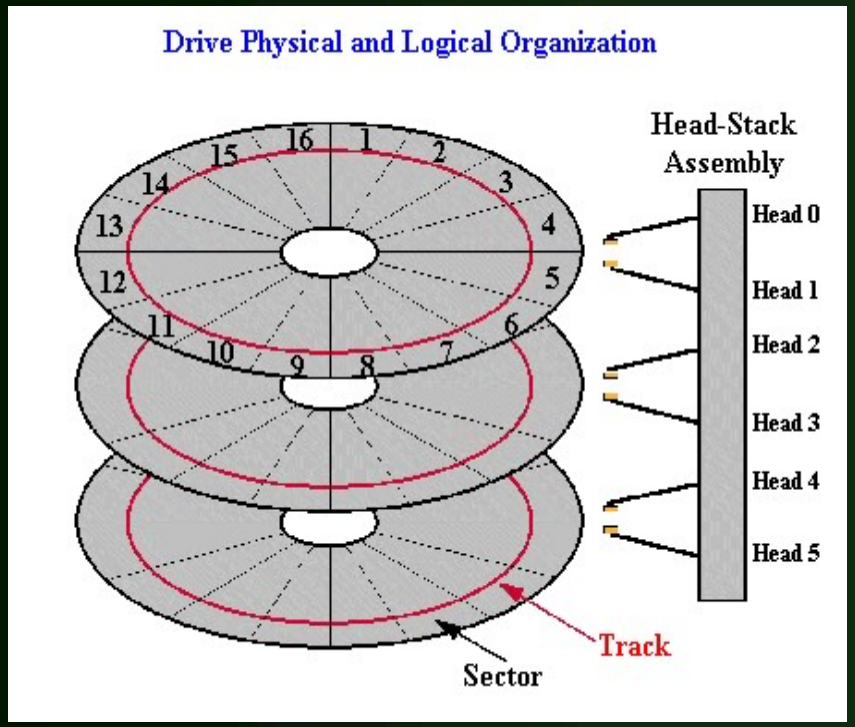

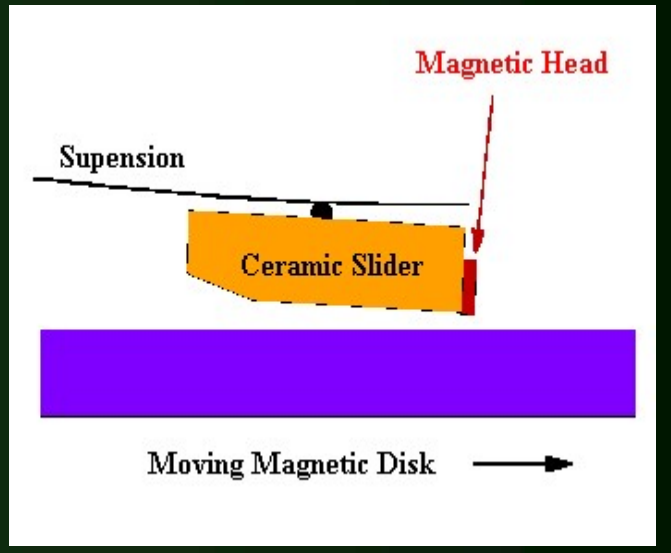

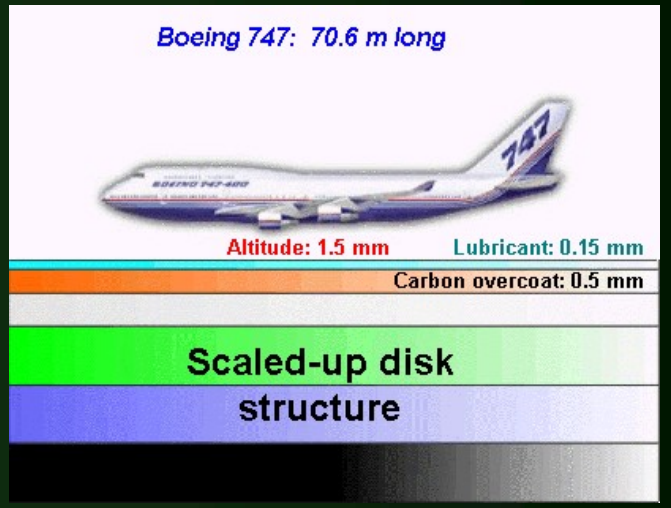

Contact Start-Stop disks: the read/write head of the disk at rest lies on track near the center (Parking track).
Load-unload on the ramp outside the disc.

Additional basic concepts

- **The arc length of the tracks/sectors varies, but the information content is constant. It used to be a fixed sector number. Changing today (10-20 zones: notches)**
- **The peripheral speed of the read-write heads is variable, the angular speed is constant (within limits).**
- **"One channel" is reading and writing. Synchronization?**
- **Writing density: the number of bits that can be placed per unit length. It depends on material quality and flux density.**

Encodings

- **NRZI** (Non Return to Zero Inverting): flux change for bits 1. For tapes.
- **FM** (frequency modulation): synchronous signals with a given frequency, in between another flux change for bit 1. (0: 1 pulse, 1: 2 pulses, average 1.5 pulses)
- **MFM** (modified FM): for bit 1 encoding, current level change between the current and the next sync signal, 0 encoding depends on the previous bit. Before 0: switching at the moment of the sync signal, before 1: no switching at the moment of the sync signal.
  (1 : 1 pulse, 0 : 0, or 1 pulse, average 0.75 pulses)

NRZI, FM and MFM coding

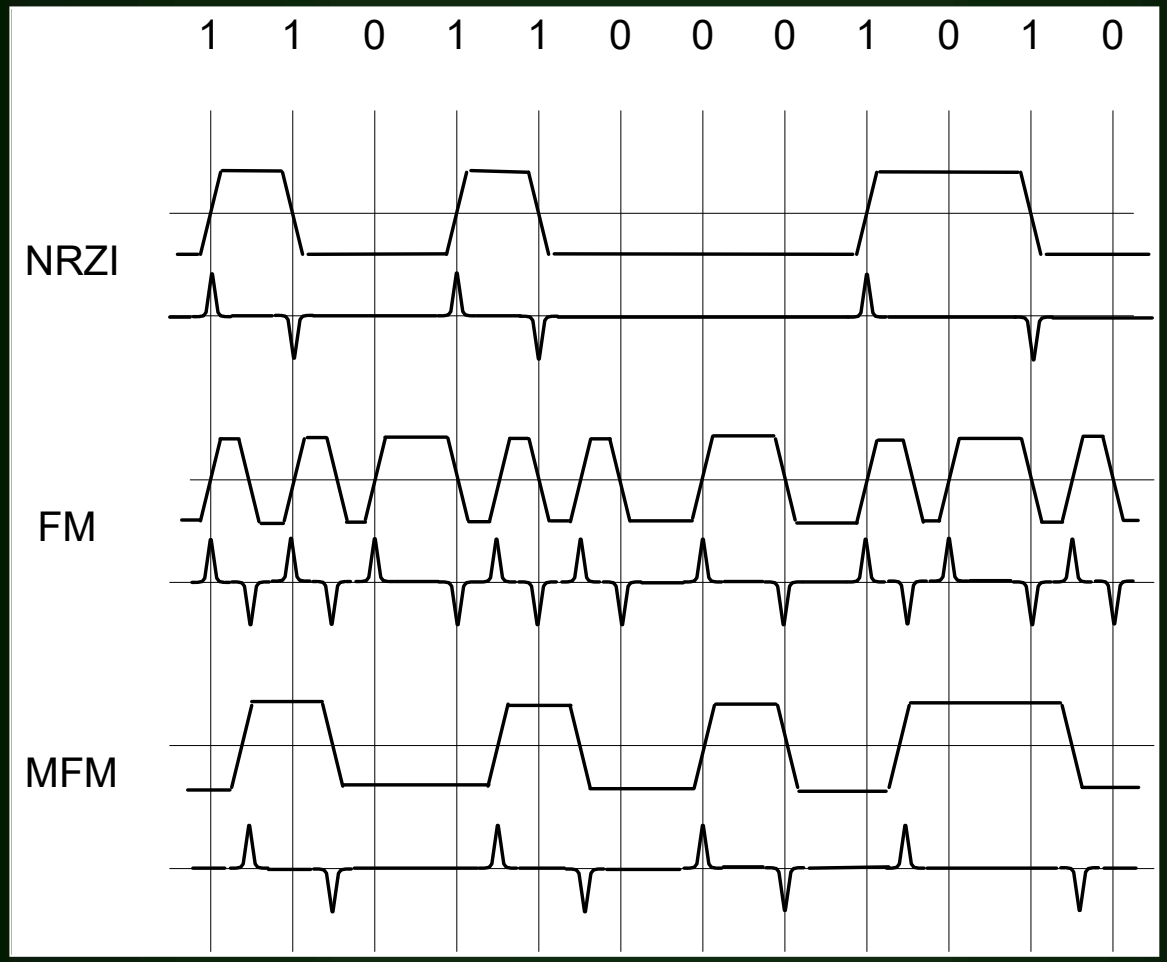

1: flux change

0: flux change,
1: even a flux change
between bit positions

switch between
current and
next level ,
0: depends on previous b
before 0: switch in
bit position,
before 1: no switch

FM, MFM and MMFM coding

FM

```
   Data:  1 1 0 1 0 0 1 0 0 0 1 0 0 0 0 1 0 0 0 0 0 1
  Clock: 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
Encoded: 11111011101011101010111010101011101010101001
```

MFM

```
   Data:  1 1 0 1 0 0 1 0 0 0 1 0 0 0 0 1 0 0 0 0 0 1
  Clock: 0 0 0 0 0 1 0 0 1 1 0 0 1 1 1 0 0 1 1 1 1 0
Encoded: 0101000100100100101001010010101001010101001
```

MMFM

```
   Data:  1 1 0 1 0 0 1 0 0 0 1 0 0 0 0 1 0 0 0 0 0 1
  Clock: 0 0 0 0 0 1 0 0 1 0 0 0 1 0 1 0 0 1 0 1 0 0
Encoded: 0101000100100100100010010010010010010010010001
```

MMFM: Changes every second 1 clock bit to 0

The addresses of the sectors

- **Track-diskside-sector** address triplets (**cyl-head-sec**).
- **One-dimensional logical addresses** (**LBA:** **Logical Block Address**) can be created if
  - the pages are numbered in a given order,
  - the tracks are also numbered in a given order.
- The one-dimensional logical address can be mapped (back) from the address triplet. This mapping can be done by the controller! (Or the disc itself!)
- The disk looks like a 0 ... n sectors (blocks) series.
- Buffering is common, caching is possible.

It affects reading and writing

- The search time (seek time): head movement to track (smaller the closer);
- the rotation latency: while the sector rotates under the head;
  - 5400 – 7200 – 10000 - 15000 rpm;
  - average half turn: 5.56 - 4.17 - 3.00 - 2.00 ms
- data transfer time: the time of the controlled transfer.
- Of these three, the first is the largest, and it is the most decisive. This should be optimal.
- Interleaving concept: non-continuous sector numbering within a track, sector processing during the rotation latency.

Seagate, RPM 15K drives (15000 rpm):
Seek time 3.6 – 4.7 msec
Rot latency 2.00 – 1.99. ms
Disk access time: 5.6 - 6.6 msec

Disk scheduling algorithms

- the **seek-time**: in what order should we "serve" incoming track (cylinder) requests?
- Algorithms:
  - **FCFS** (First Come First Served): no optimization.
  - **SSF** (Shortest Seek First): the smallest head movements.
  - **Elevator** algorithm: one-way collector.
- An example
  Cylinder requests: 11, 1, 25, 20, 28, 9, 12
  - SSF: 11, 12, 9, 1, 20, 25, 28
  - Elevator: 11, 12, 20, 25, 28, 9, 1

Recent magnetic disks

- Winchester disks:
  - in a closed box, protected from dirt and moisture,
  - high speed, many sides (heads),
  - heads "fly" on the surface,
  - linear or circular head movement.
  - Large capacities. More and more intelligence in the device.
  - Caching is now natural. Attention: write caching must be enabled on SCSI!
- Floppies (extinct).

For PCs

- EIDE
- cheaper,
- controller on the motherboard,
- 2 channels (chanel)
  - primary: 2 devices
  - secondary: 2 devices
- device can be: W, CD
- problem: if both W and CD are on a channel, and a CD operation is started, the channel is busy, the performance is bad!

- SCSI
- more expensive, you need an extra controller.
- 7 devices on SCSI,
- 15 on wide SCSI.
- Device can be: W, CD, scanner, etc.
- If an operation is in progress but not currently using the bus, another operation may be in progress.
- It asks for LBA, does the BIOS know?

Furthermore

- EIDE
- 2.1 (ATA) - 16.6 (ATA2) MByte /sec
- ATA/ATAPI-4, Ultra ATA/33, (Attachment Packet Interface: ATAPI) (CD-ROM, tape drives, CompactFlash for solid state drives,
- ATA/ATAPI-5, Ultra ATA/66,
- ATA/ATAPI-6, Ultra ATA/100
- ATA/ATAPI-7, Ultra ATA/133 (UDMA 6)
- SATA150: 150 Mbytes / sec
- SATA300: 300 Mbytes / sec
- SATA 6 00: 6 00 MByte / sec

- SCSI
- SCSI-1: 5MHz, 5 Mbyte/s
- SCSI-2: 10MHz, 10-20 MB/s
- Fast20, Ultra: 20 MHz, 20-40 Mbyte/s
- Fast40, Ultra-2: 40 MHZ, 40-80-160 Mbyte/sec

Furthermore

Name	Raw bandwidth (Mbit/s)	Transfer speed (MByte/s)	Max. cable length (m)	Power provided	Devices per Channel
eSATA	3,000	300	2 with eSATA (1 passive adapter)	Woman	1 (15 with port multiplier)
eSATAp				5V/ 12V	
SATA 600	4,800	600	1	Woman	
SATA 300	2,400	300			
SATA 150	1,200	150			1 per line
HOOF 133	1,064	133.5	0.46 (18in)	Woman	2
SAS 600	6,000	600	10	Woman	1 (>65k with expanders)
SAS 300	3,000	300			
SAS 150	1,500	150			
FireWire 3200	3.144	393	100 (spec . cable)	15W, 12-25V	63 (with hub)
FireWire 800	786	98.25	100		
FireWire 400	393	49.13	4.5		
USB 3.0	3,200	400	3	4.5W, 5V	127 (with hub)
USB 2.0	480	60	5	2.5W, 5V	
USB 1.0	12	1.5	3	Yes	
SCSI Ultra-320	2,560	320	12	Woman	15 (plus HBA)

Interfaces for disks...

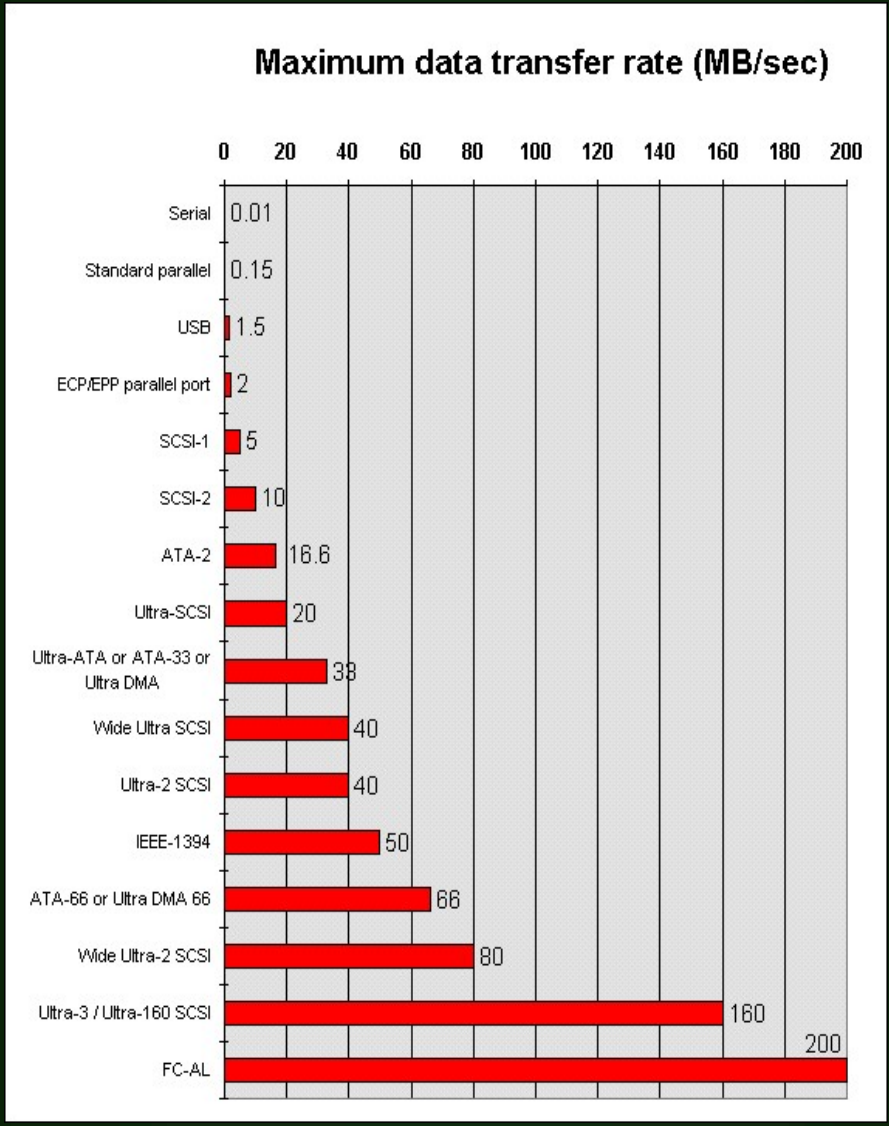

Development of PC systems

	1987	2001	Growth
CPU performance	1 MIPS	~2000 MIPS	2000x
Memory size	64 KB	512 MB	8000x
Memory access	100 µsec	100 nsec	1000x
Disk capacity	20 MB	72 MB	3600x
Disk access	60 msec	6 msec	10x

Some articles...

http://seagate.com/docs/pdf/whitepaper/disc_capacity_performance.pdf

http://www.usbyte.com/common/whitepapers/WDC/IDE_Drive_Installation_Guide_WDC.pdf

http://www.usbyte.com/common/whitepapers/WDC/Quick_Install_For_WDC_FireWire_Drives_WDC.pdf

Here is a collection of articles:

http://www.usbyte.com/common/whitepapers/HDD_WP.htm

CD discs

- Optical technology, laser light...
- The cross section...
- The "strip" of data storage is a "spiral" here...

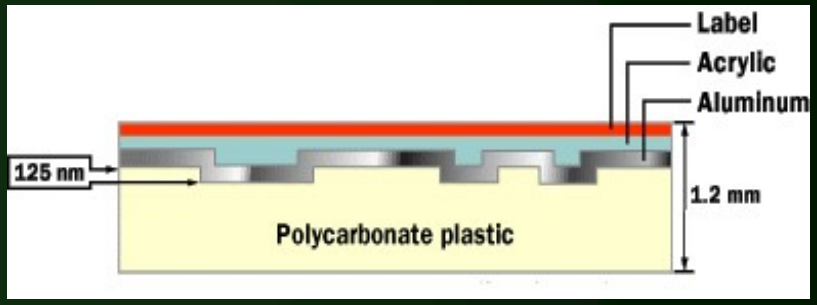

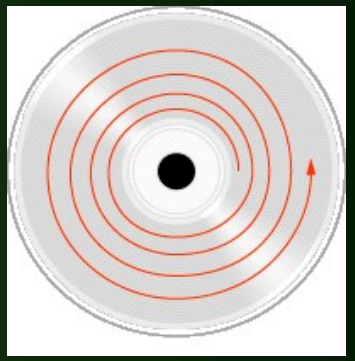

- Storage: reflective base surface (land) and less reflective (bump/pit) surfaces
- When reading, the laser light is focused on the "track" and the intensity of the reflected light changes

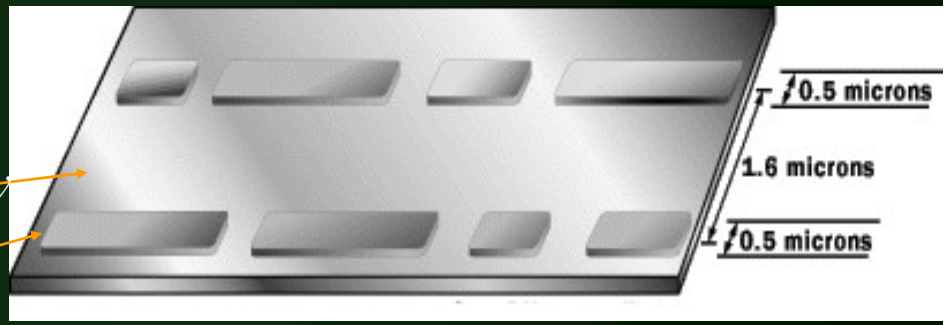

Minimum dimensions of the bump/pit:
length~0.83 μm , width~0.5 μm , height~125nm

Basics

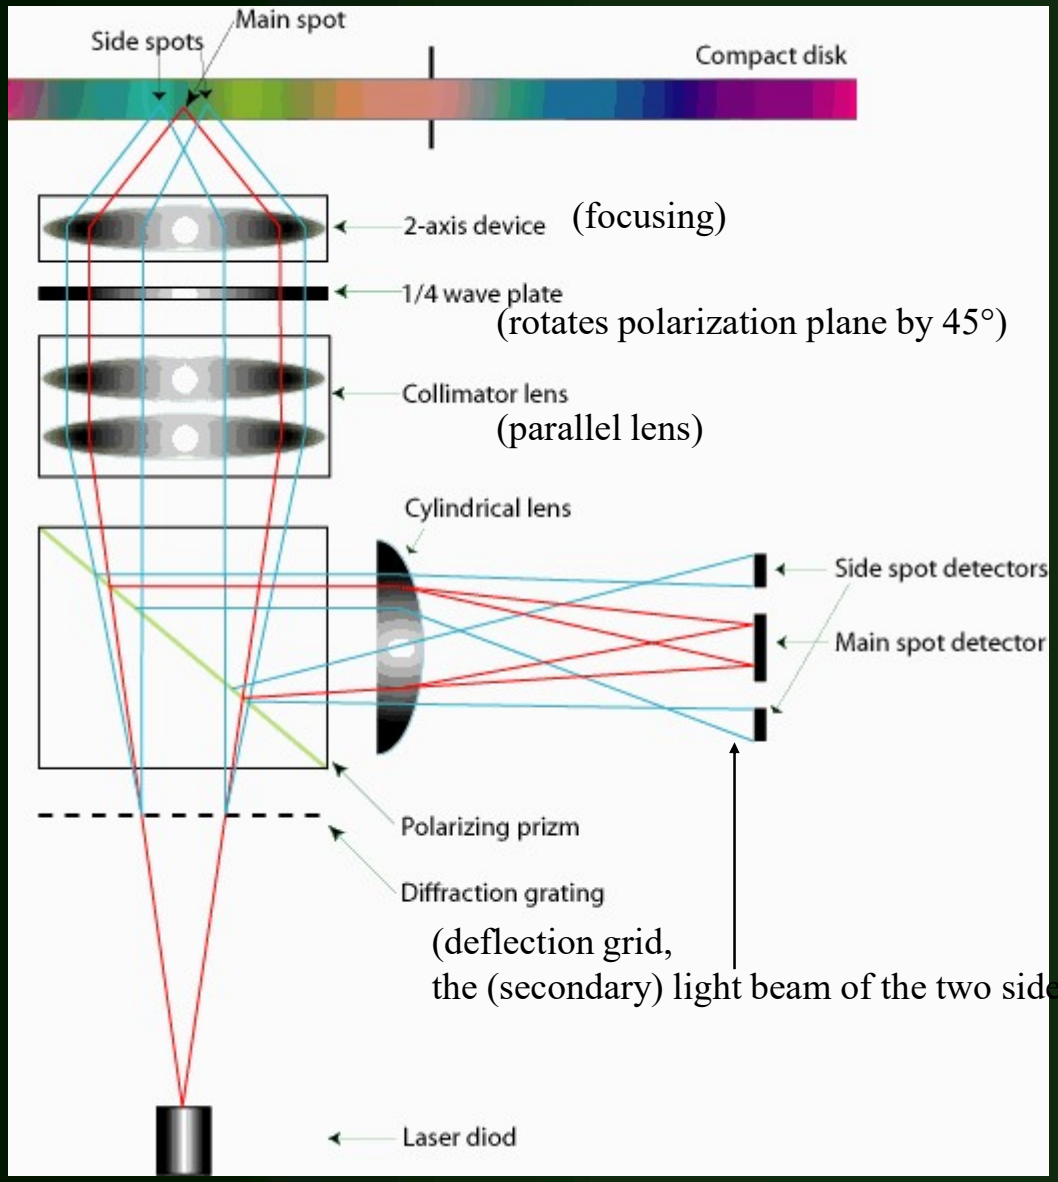

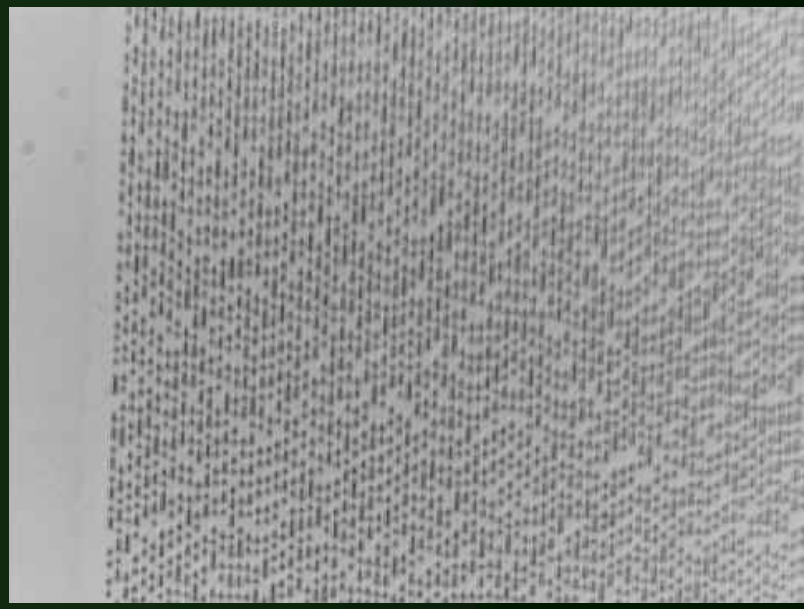

The two side spots are needed to "follow" the tracks.
They have a control role.

Diffraction - light deviation
Diffractin grating - deflection grid.
This “creates" the two side (secondary) light beams.

Collimator lens - collimator lens.
¼ wawe plate – turns the polarization plane of the light by 45°.
(The reflection by another 45°, i.e. the initial vertical polarization to horizontal).
The 2-axis-device is focusing.

Spot detectors: 4 pcs. for the main beam, 1-1 for the side beams, there are a total of 6 detectors. The detectors "read" the pit information.

http://www.usbyte.com/common/compact_disk.htm

Bit and channel coding

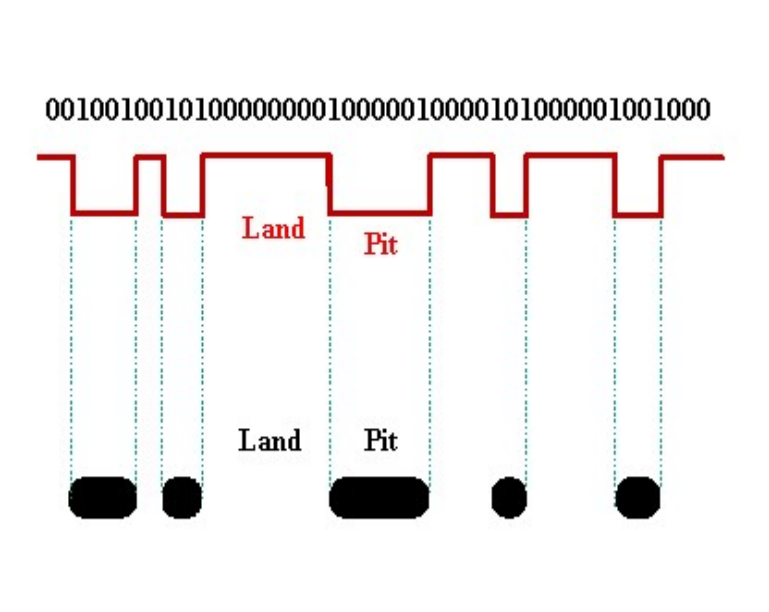

- The bit encoding: each "intensity change" (from pit to land and vice versa; see the red line in the figure) represents 1 bit.
- Channel coding is EFM (eight-to-fourteen modulation): a byte is converted into a 14-bit code (back to byte during read).
- cross-interleaved Reed–Solomon code (CIRC)

CIRC corrects error bursts up to 3,500 bits in sequence (2.4 mm in length as seen on CD surface) and compensates for error bursts up to 12,000 bits (8.5 mm) that may be caused by minor scratches

CD-ROM

- IEC-10149 standard
- On the “track" (approx. 270,000 units) sector... (approx. 4.5 Km)
- One sector (2352 bytes)
  - 12-byte synchronous field (00+10*FF+00)
  - 4 byte header field
    - 3-byte sector address (minutes:seconds:hundredths of a second)
    - Mode on 1 byte (mode 0, 1, 2)
  - 2048|2336 byte data field (the first for mode 1)
  - 288|0 byte EDC error correction code field (first for mode 1)
- EFM coding eight-to-fourteen
  ~ 8 bits are converted to 14 bits, the 0 Byte has also 1

CD-ROM

- *lead-in* track
  - silent audio
  - subcode contains repeated copies of the Table Of Contents (TOC)
  - index of the start positions of the tracks
  - in absolute timecode, relative to the start of the program area
- *program area*
- *lead-out* track

https://en.wikipedia.org/wiki/Compact_Disc_Digital_Audio

CD drive

- Parts of the drive
- CLV (Constant Line Velocity) constant peripheral speed: ~75 sectors/sec
- The angular velocity is therefore To be kept between 200 and 530 rpm...
- Of this, approx. 150 KB/sec channel speed …
- Today it can be multiple times:
  - 2X (double)
  - 4X (quadruple) etc.
  - (Constant Angular Velocity) at 12X (or higher) speed

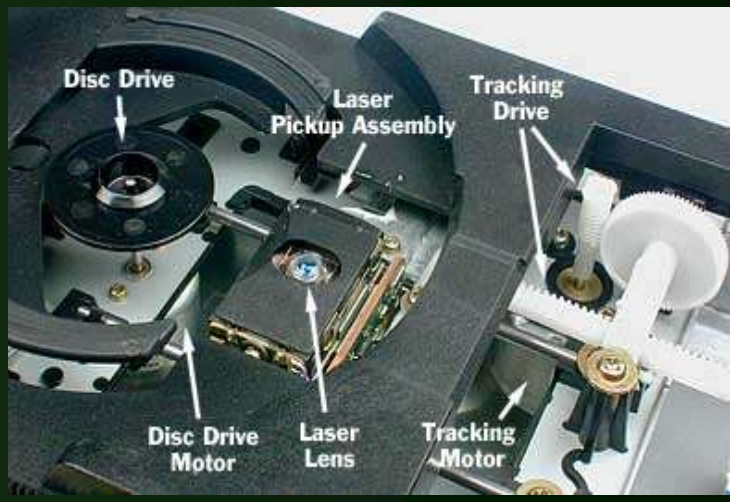

Writable, rewritable CD

- CD-R
  - Paint layer before the aluminum layer. In its basic state, the (weak) laser light passes through, it can be reflected on the aluminum layer
  - Stronger laser light creates a chemical change in the paint layer: it makes it impermeable to light (opaque)
- CD-RW
  - Phase change layer (compound layer) between 2 dielectric layers in front of the aluminum. This
    - Crystalline translucent (can be reflected from below)
    - In the amorphous state, it is not transmissive (does not reflect)
  - Reading laser, erasing laser, writing laser: increasingly "stronger"

Literature

- http://www.usbyte.com/common/compact_disk.htm
- I also recommend literature for the DVD: http://www.usbyte.com/common/dvd.htm

DVD

- Formerly: Digital Video Disc
- Today: Digital Versatile Disc

Features	DVD	CD-ROM
Substrate diameter / thickness (mm)	120 / 1.2	120 / 1.2
Sides	1 or 2	1
Layers per side	1 or 2	1
Capacity (GB)	4.7, 8.54, 9.4, or 17	~ 0.7
Track pitch (microns)	0.74	1.6
Min pit length (microns)	0.4 - 0.44	0.83
Linear velocity used for scanning (m/s)	3.5 - 3.84	1.3
Laser wavelength (nm)	635 or 650	780
Numerical aperture	0.6	0.45
Modulation	8 to 16	EFM (8 to 14)
Error correction code (ECC)	RSPC	CIRC
Durability and dust/scratch	same as that of CD	high

Various DVDs

Name	Media structure	Capacity (GB)
DVD-5	Single Side / Single Layer	4.7
DVD-9	Single Side / Dual Layer	8.54
DVD-10	Double Side / Single Layer	9.4
DVD-18	Double Side / Dual Layer	17.08
DVD-R	Single or Double Side / Single Layer	3.95 / 7.9
DVD-RAM	Single or Double Side / Single Layer	2.6 / 5.2

DVD-9 and DVD-10

DVD-9: Single Side, Double Layer Disk (8.54 GB)

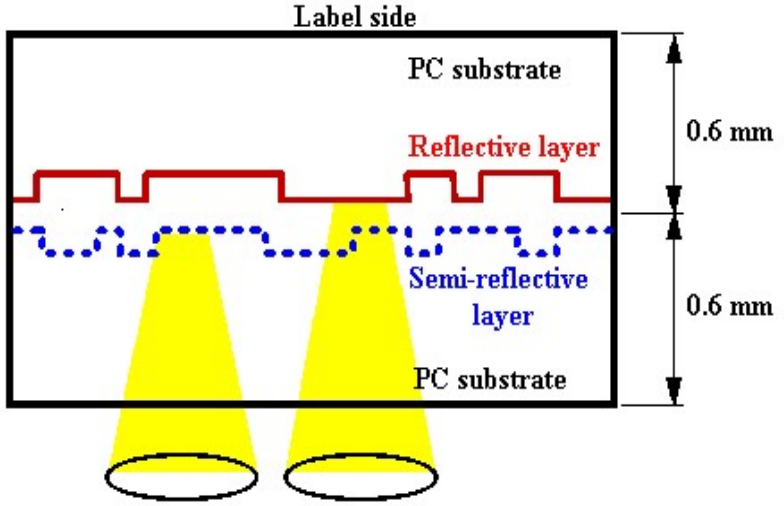

DVD-10: Double Side, Single Layer Disk (9.4 GB)

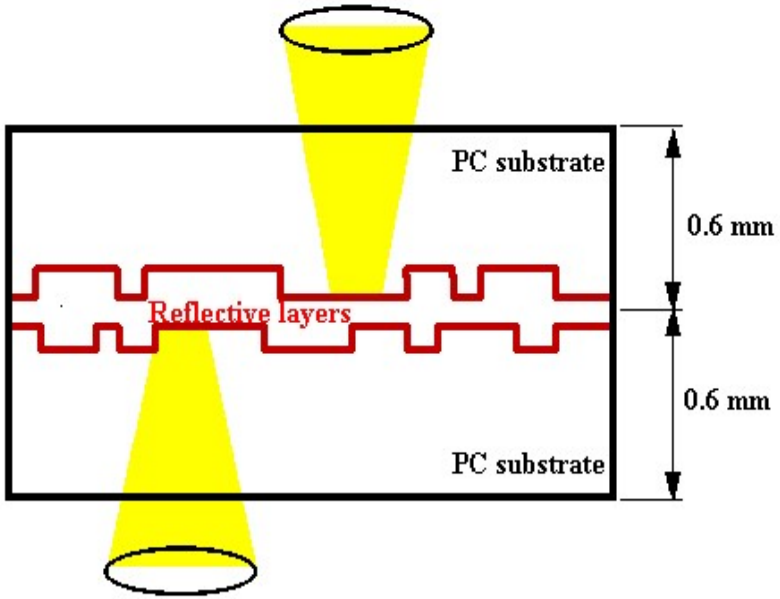

DVD-18

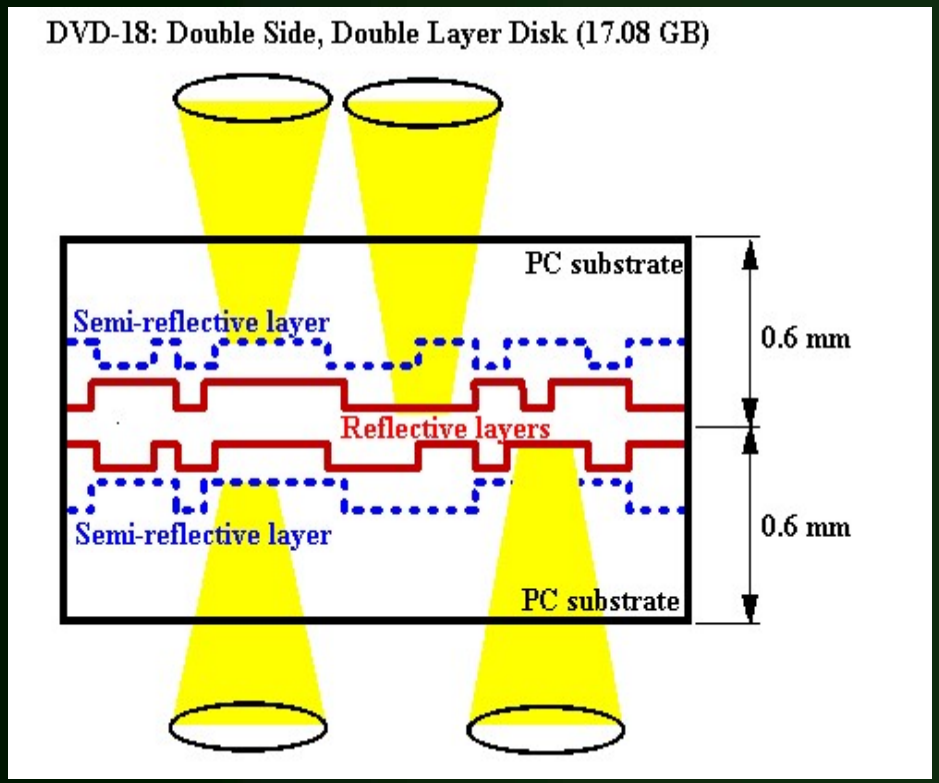

http://www.usbyte.com/common/dvd.htm

Blu-ray disc

- Single layer 25 GB
- Double layer 50 GB

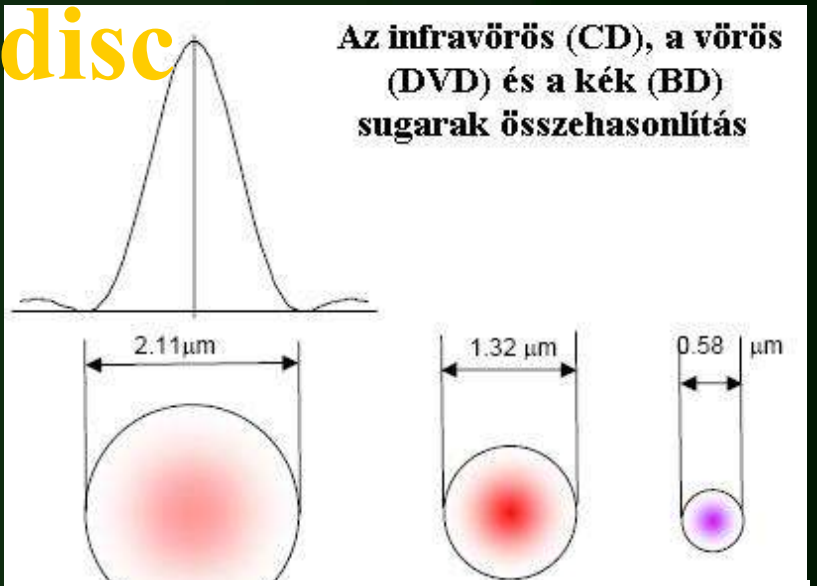

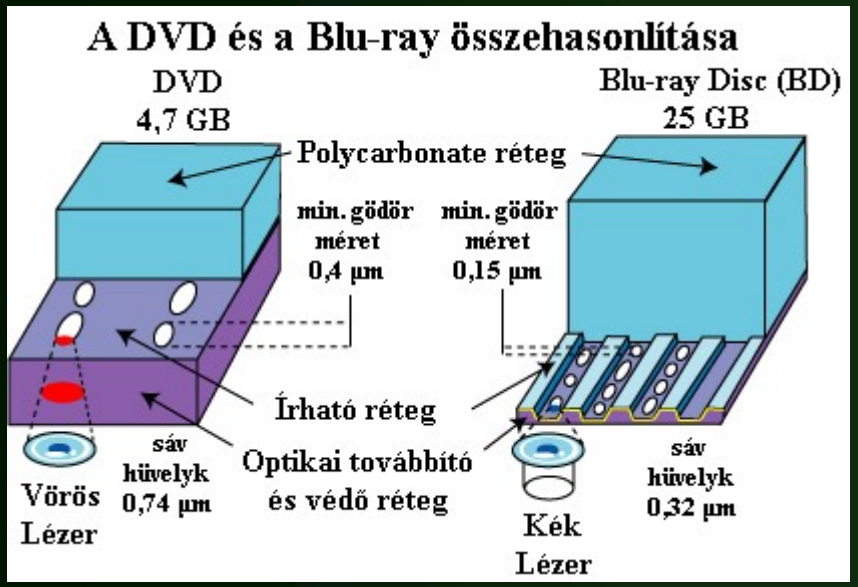

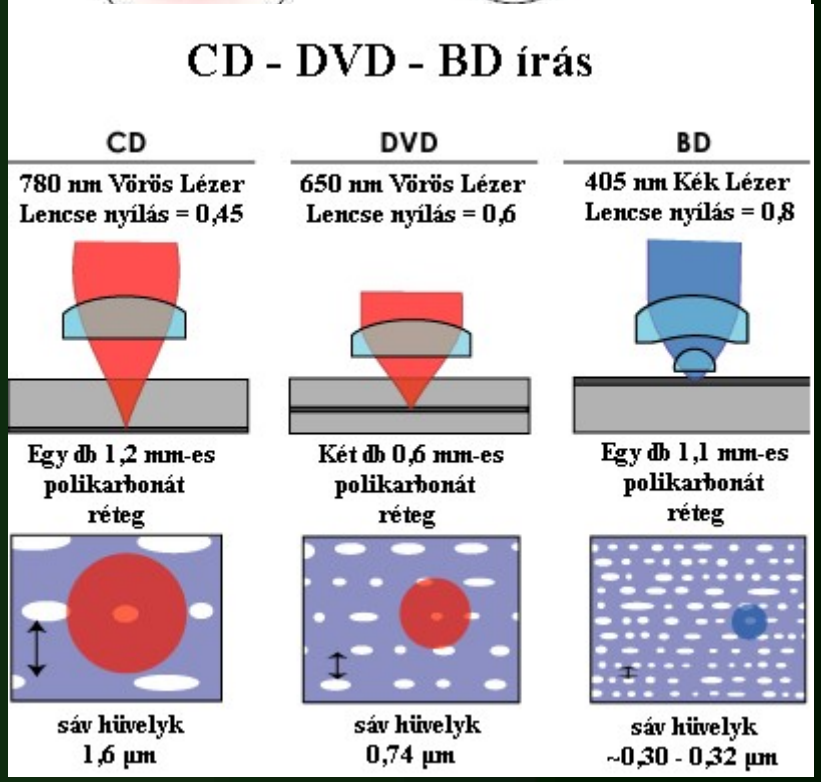

http://hu.wikipedia.org/wiki/Blu-ray_disc

COMPUTER ARCHITECTURES

Devices, device controls

End

Általános
INFORMATIKAI
Tanszékk