This is the third lecture of Chapter 7

Chapter 7

Input/Output Systems (C)

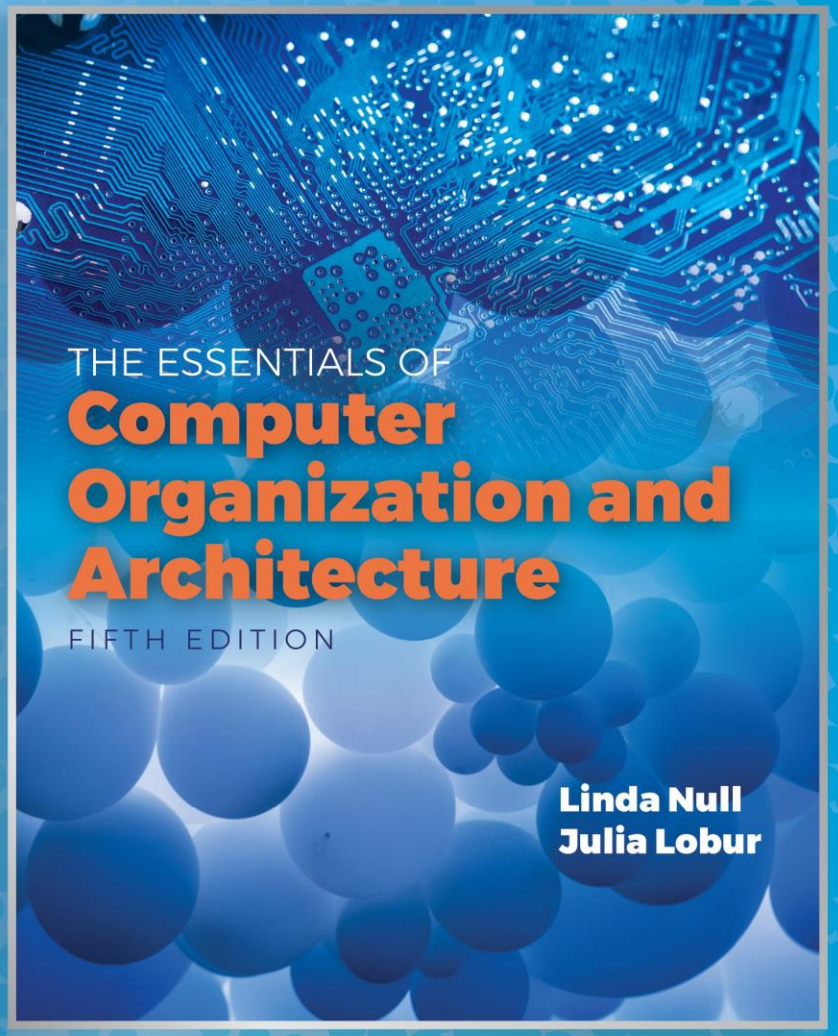

Quick review of last lecture (1)

- I/O Architectures
  - I/O control methods
    - Channel I/O
- Character I/O vs Block I/O
- I/O bus operations
  - Data bus, address bus, and control bus
  - Asynchronous (Use handshaking)
  - Configurations (Connections between IO bus and IO cards)
  - Bus arbitration
  - Bus timing diagrams

Quick review of last lecture (2)

- Data Transmission Modes
  - Parallel data transmission
  - serial data transmission
- Disk Technology
  - Magnetic disk organization: surfaces, tracks, sectors.
    - Disk Capacity
  - Disk performance
    - Seek time, rotational delay,
    - Average latency, access time, …
  - Role of Disk

$$\frac{\frac{60 \text{ seconds}}{\text{disk rotation speed}} \times \frac{1000 \text{ ms}}{\text{second}}}{2}$$

Example: Disk Capacity and Performance

- Suppose a disk drive has the following characteristics:
  - 5 surfaces
  - 1024 tracks per surface
  - 256 sectors per track
  - 512 bytes/sector
  - Track-to-track seek time of 8 milliseconds
  - Rotational speed of 7500 RPM.

a) What is the capacity of the drive? b) What is the access time?

- Disk Capacity C = N × S × T × P (Bytes)
  - N: bytes/sector, S: sectors/track, T: tracks/surface, P: platter surfaces
  - $512 \times 256 \times 1024 \times 5 = 640 \times 2^{20}$ Bytes = 640MB
- Average rotation delay: ((1/R)×60×1000)/2 (ms)
  - R: rotational speed
  - ((1/7500)×60 ×1000)/2 = 4ms
- Access time = Average Rotation delay + Seek Time
  - Access time = 4 ms (Rotational delay) + 8 ms (seek time) = 12 ms

More about Magnetic Disks

- Magnetic disks play two roles in computer systems
  - Long-term, nonvolatile storage for files, even when no program running
  - A level of the memory hierarchy below main memory used for as a backing store for virtual memory during program execution

- Outermost tracks have the lowest bit density per areal measure; hence less prone to bit error than innermost tracts
- The heads of a hard drive float just above the surface on a cushion of rapidly-moving air.

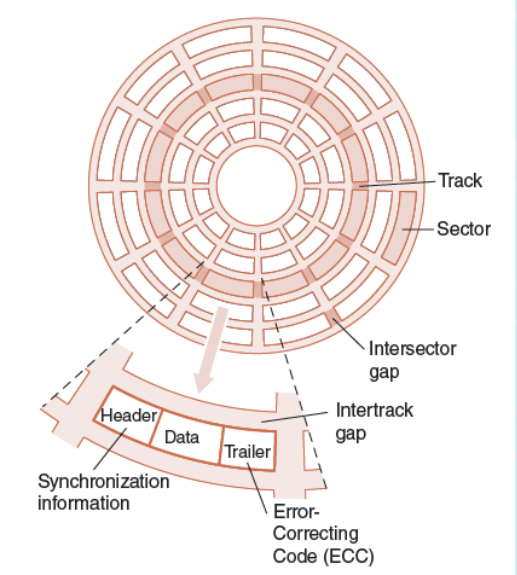

Master Control Block

is a reserved section of a disk that keeps track of the makeup of the rest of disk

A disk file is a collection of sectors that are linked together to form a single logical entity.

Disk Directory

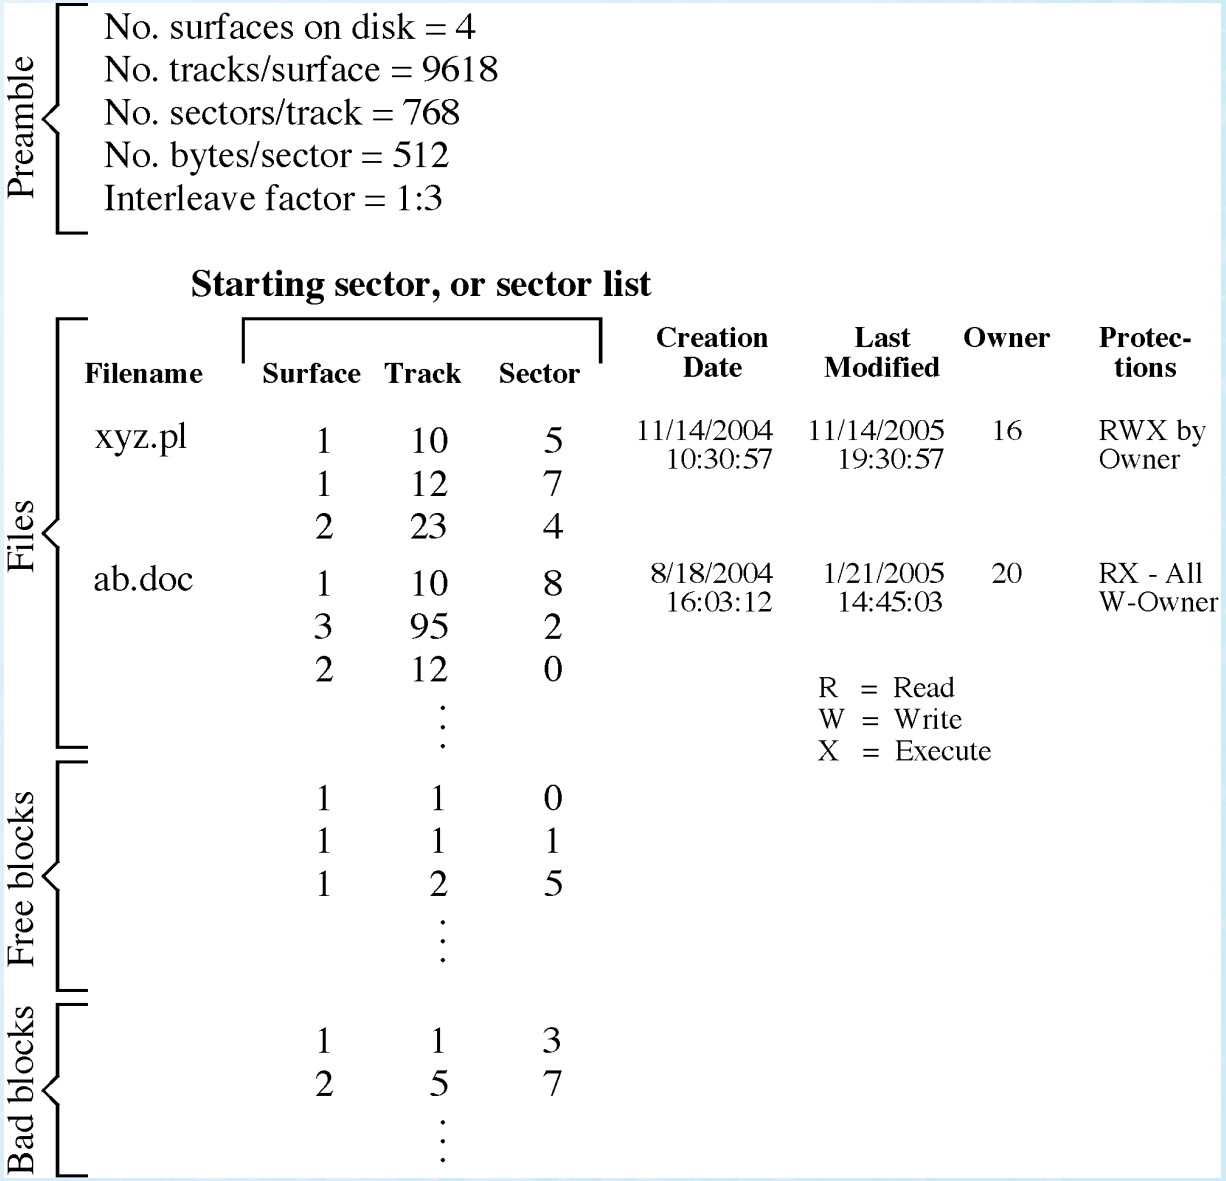

7.6.1 Rigid Disk Drives (6 of 6)

- Low cost is the major advantage of hard disks.
- But their limitations include:
  - Very slow speed compared to main memory
  - Fragility
  - Moving parts wear out
  - Power hungry
- Reductions in memory cost enable the widespread adoption of *solid state drives* (SSDs).
  - Computers "see" SSDs as just another disk drive, but they store data in non-volatile *flash* memory circuits.
  - Flash memory is also found in memory sticks and MP3 players.

7.6.2 Solid State Drives (1 of 3)

- SSD access time and transfer rates are *typically* 100 times faster than magnetic disk, but slower than onboard RAM by a factor of 100,000.
  - These numbers vary widely among manufacturers and interface methods.
- Unlike RAM, flash is block-addressable (like disk drives).
  - The duty cycle of flash is between 30,000 and 1,000,000 updates to a block.
  - Updates are spread over the entire medium through *wear leveling* to prolong the life of the SSD.

7.6.2 Solid State Drives (2 of 3)

- SSD specifications share many common metrics with HDDs.
  - Clearly, there is no need for any metrics that concern spinning platters, such as rotational delay.
  - Compare Figs 7.15 with 7.17 in your text.
- Enterprise SSDs must maintain the highest degree of performance and reliability.
  - Onboard cache memories are backed up by capacitors that briefly hold a charge during a power failure, giving time to commit pending writes.

CONFIGURATION:	
Capacity, GB	800
Integrated controller	SATA 3.0
Encryption	AES 256-bit
Cache size	1GB
Bytes per sector	512
PHYSICAL:	
Height	7mm
Length	100mm
Width	70mm
Weight	170g
Temperature (°C)	
Operating	0°C to 70°C
Nonoperating/storage	–55°C to 95°C
Relative humidity	5% – 95%
Acoustic noise	0dB

RELIABILITY AND MAINTENANCE:	
MTTF	2,000,000 hours
Endurance	450 TBW
Data retention	3 months
Data errors (UBER)	<1 per 10^{17} sector read
PERFORMANCE:	
Average latency	(Sequential)
Read	50ms
Write	65ms
I/O Operations/Sec (IOPS)	(Random)
8KB Reads	47,500 IOPS
8KB Writes	5,500 IOPS
Data transfer rate:	
Read	500MB/sec
Write	450MB/sec
Start time	
(0 to drive ready)	3 sec

Fig 7.17

7.6.2 Solid State Drives (3 of 3)

- The Joint Electron Devices Engineering Council (JEDEC) sets standards for SSD performance and reliability metrics. The most important are:
- Unrecoverable Bit Error Ratio (UBER) and terabytes written (TBW). TBW is a measure of disk endurance (or service life) and UBER is a measure of disk reliability.
  - UBER is calculated by dividing the number of data errors by the number of bits read using a simulated lifetime workload.
  - TBW is the number of terabytes that can be written to the disk before the disk fails to meet specifications for speed and error rates.

7.7 Optical Disks (1 of 7)

- Optical disks provide large storage capacities very inexpensively.
- They come in a number of varieties including
  - CD-ROM (compact disk – read only memory) ,
  - CD-R (CD recordable)
  - CD-RW (CD rewretable)
  - WORM (write once read many)
- Many large computer installations produce document output on optical disk rather than on paper. This idea is called COLD—*Computer Output Laser Disk*.
- It is estimated that optical disks can endure for a hundred years. Other media are good for only a decade—at best.

7.7 Optical Disks (2 of 7)

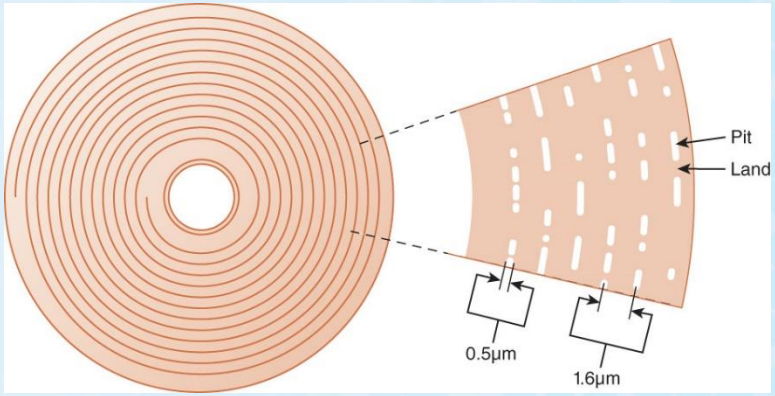

- CD-ROMs were designed by the music industry in the 1980s, and later adapted to data.
- This history is reflected by the fact that data is recorded in a single spiral track, starting from the center of the disk and spanning outward.
- Binary ones and zeros are delineated by bumps in the polycarbonate disk substrate. The transitions between pits and lands define binary ones.
- If you could unravel a full CD-ROM track, it would be nearly 5 miles long!

7.7 Optical Disks (3 of 7)

- The logical data format for a CD-ROM is much more complex than that of a magnetic disk. (See the text for details.)
- Different formats are provided for data and music.
- Two levels of error correction are provided for the data format.
- Because of this, a CD holds at most 650MB of data, but can contain as much as 742MB of music.

7.7 Optical Disks (4 of 7)

- DVDs can be thought of as quad-density CDs.
  - Varieties include single sided, single layer, single sided double layer, double sided double layer, and double sided double layer.
- Where a CD-ROM can hold at most 650MB of data, DVDs can hold as much as 17GB.
- One of the reasons for this is that DVD employs a laser that has a shorter wavelength than the CD's laser.
- This allows pits and lands to be closer together and the spiral track to be wound tighter.

7.7 Optical Disks (5 of 7)

- A shorter wavelength light can read and write bytes in greater densities than can be done by a longer wavelength laser.
- This is one reason that DVD's density is greater than that of CD.
- The 405 nm wavelength of blue-violet light is much shorter than either red (750 nm) or orange (650 nm).
- The manufacture of blue-violet lasers can now be done economically, bringing about the next generation of laser disks.

7.7 Optical Disks (6 of 7)

- The Blu-Ray disc format won market dominance over HD-CD owing mainly to the influence of Sony.
  - HD-CDs are backward compatible with DVD, but hold less data.
- Blu-Ray was developed by a consortium of nine companies that includes Sony, Samsung, and Pioneer.
  - Maximum capacity of a single layer Blu-Ray disk is 25GB.
  - Multiple layers can be “stacked” up to six deep.
  - Only double-layer disks are available for home use.

7.7 Optical Disks (7 of 7)

- Blue-violet laser disks are also used in the data center.
- The intention is to provide a means for long term data storage and retrieval.
- Two types are now dominant:
  - Sony's Professional Disk for Data (PDD) that can store 23GB on one disk
  - Plasmon's Ultra Density Optical (UDO) that can hold up to 30GB
- It is too soon to tell which of these technologies will emerge as the winner.