7.6 Magnetic Disk Technology

- Magnetic disks offer large amounts of durable storage that can be accessed quickly.
- Disk drives are called *random* (or *direct*) *access storage devices,* because blocks of data can be accessed according to their location on the disk.
  - This term was coined when all other durable storage (e.g., tape) was sequential.
- Magnetic disk organization is shown in Figure 7.13.

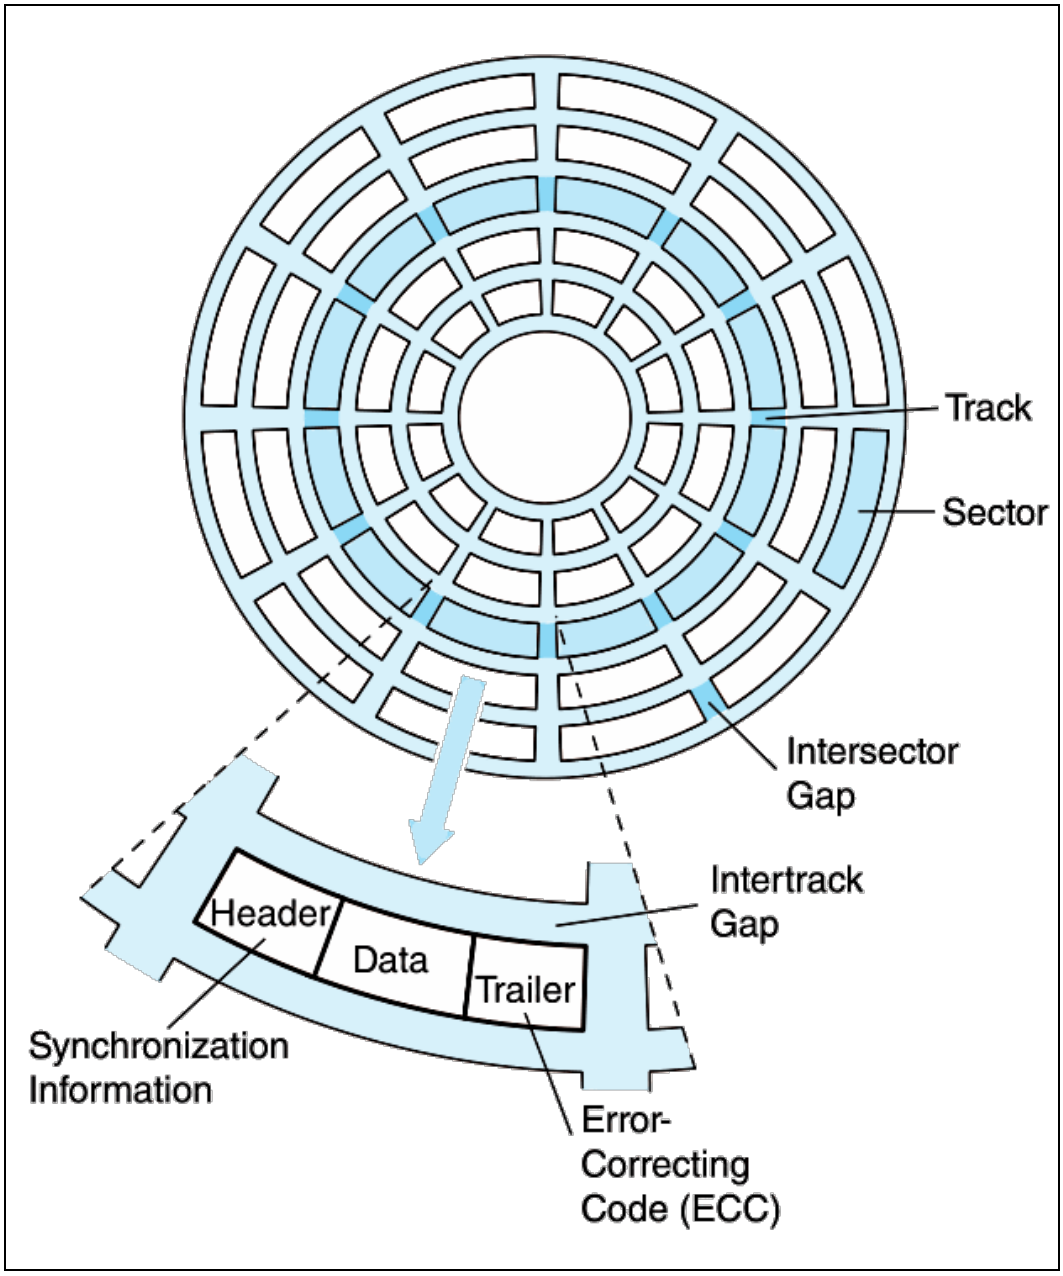

Figure 7.13 Disk Sectors Showing Intersector Gaps and Logical Sector Format

- Disk tracks are numbered from the outside edge, starting with zero.

7.6.1 Rigid Disk Drives

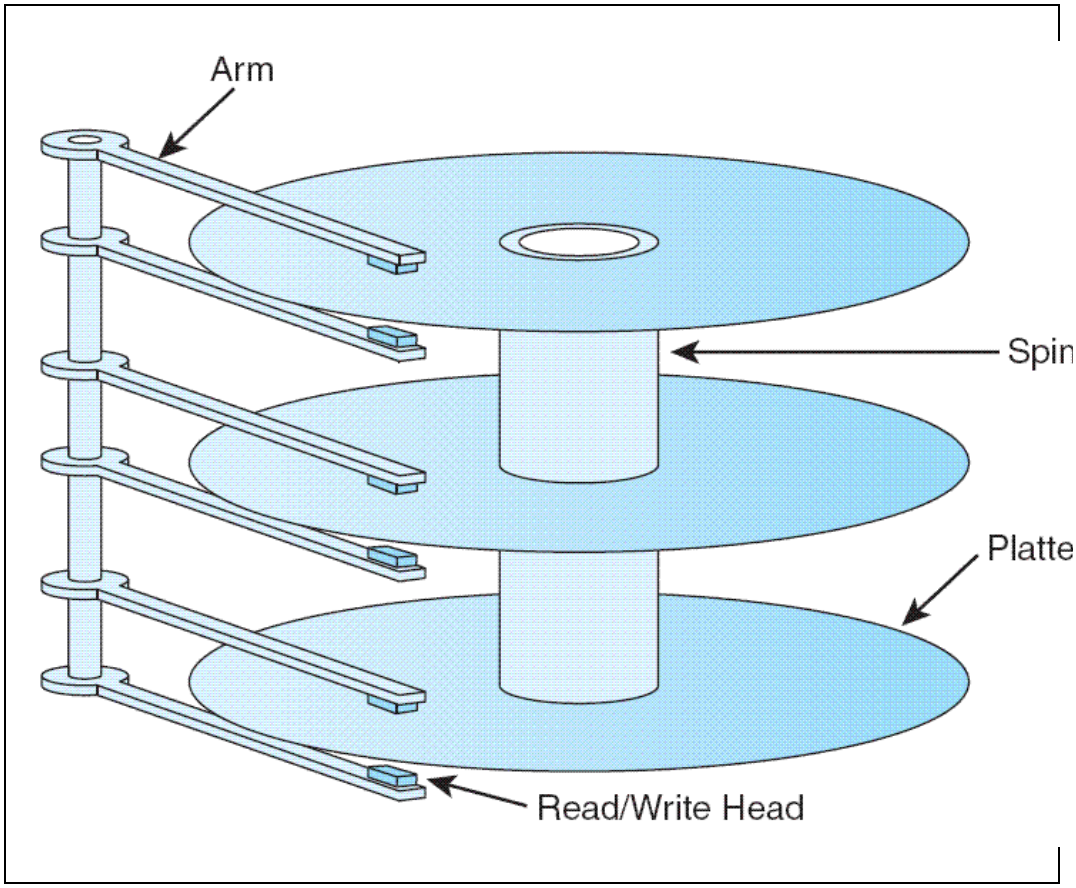

Figure 7.14 Rigid Disk Actuator (with Read/Write Heads) and Disk Platters

- Hard disk platters are mounted on spindles.
- Read/write heads are mounted on a comb that swings radially to read the disk.
- The rotating disk forms a logical cylinder beneath the read/write heads.
- Data blocks are addressed by their cylinder, surface, and sector.

- There are a number of electromechanical properties of hard disk drives that determine how fast its data can be accessed.
- Seek time is the time that it takes for a disk arm to move into position over the desired cylinder.
- Rotational delay is the time that it takes for the desired sector to move into position beneath the read/write head.
- Seek time + rotational delay = access time.
- *Transfer rate* gives us the rate at which data can be read from the disk.

- *Average latency* is a function of the rotational speed:

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

- *Mean Time To Failure (MTTF) is a statistically-determined value often calculated experimentally.*
  - *It usually doesn't tell us much about the actual expected life of the disk. Design life is usually more realistic.*
- *Low cost is the major advantage of hard disks.*
- *But their limitations include:*
- *Very slow compared to main memory*
- *Fragility*
- *Moving parts wear out*

7.6.2 Solid State Drives

- *Reductions in memory cost enable the widespread adoption of solid state drives, SSDs.*
  - *Computers "see" SSDs as jut another disk drive, but they store data in non-volatile flash memory circuits.*
  - *Flash memory is also found in memory sticks and MP3 players.*

- *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.*
  - *There 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.*
- 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.
- 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.