PART 6 External Memory

6.1 Magnetic Disk

A disk is a circular platter constructed of nonmagnetic material, called the substrate, coated with a magnetizable material. Traditionally, the substrate has been an aluminum or aluminum alloy material. More recently, glass substrates have been introduced.

6.1.1 Data Organization and Formatting

Data are recorded on and later retrieved from the disk via a conducting coil named the head; in many systems, there are two heads, a read head and a write head. During a read or write operation, the head is stationary while the platter rotates beneath it.

The head is a relatively small device capable of reading from or writing to a portion of the platter rotating beneath it. This gives rise to the organization of data on the platter in a concentric set of rings, called tracks. Each track is the same width as the head. There are thousands of tracks per surface.

Figure 6.1 depicts this data layout. Adjacent tracks are separated by gaps. This prevents, or at least minimizes, errors due to misalignment of the head or simply interference of magnetic fields.

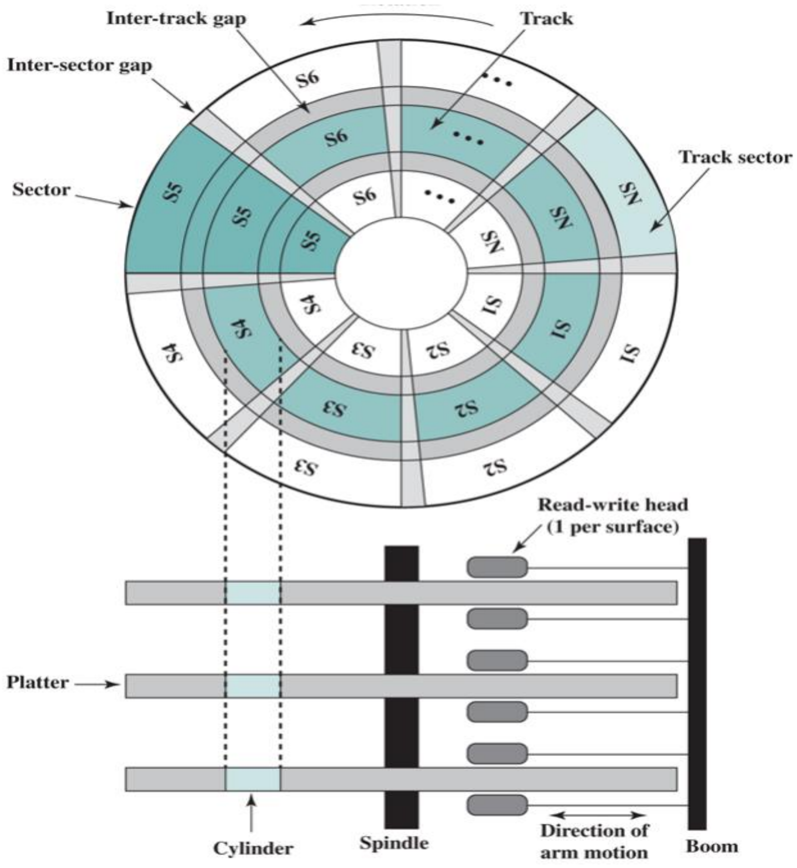

Figure 6.1 Disk Data Layout

Data are transferred to and from the disk in sectors. There are typically hundreds of sectors per track, and these may be of either fixed or variable length. To avoid imposing unreasonable precision requirements on the system, adjacent sectors are separated by intratrack (intersector) gaps.

A bit near the center of a rotating disk travels past a fixed point slower than a bit on the outside. Therefore, some way must be found to compensate for the variation in speed so that the head can read all the bits at the same rate:

constant angular velocity (CAV).This can be done by increasing the spacing between bits of information recorded in segments of the disk. The information can then be scanned at the same rate by rotating the disk at a fixed speed, as the Figure 6.2a shows the layout of a disk using CAV. The disk is divided into a number of pie-shaped sectors and into a series of concentric tracks.The advantage of using CAV is that individual blocks of data can be directly addressed by track and sector. To move the head from its current location to a specific address, it only takes a short movement of the head to a specific track and a short wait for the proper sector to spin under the head. The disadvantage of CAV is that the amount of data that can be stored on the long outer tracks is the only same as what can be stored on the short inner tracks

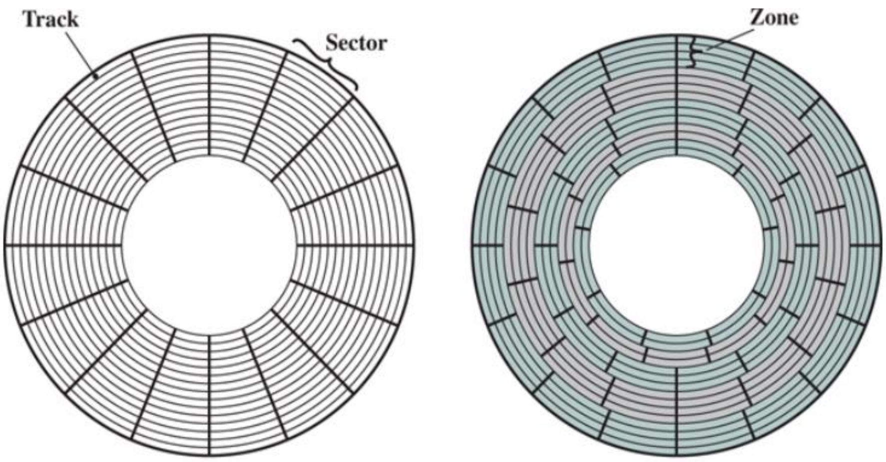

(a) Constant angular velocity (b) Multiple zone recording

Figure 6.2 Comparison of Disk Layout Method

MZR(multiple zone recording) :Figure 6.2b is a simplified layout, with 15 tracks organized into 5 zones. The innermost two zones have two tracks each, with each track having nine sectors; the next zone has 3 tracks, each with 12 sectors; and the outermost 2 zones have 4 tracks each, with each track having 16 sectors.

Figure 6.3. shows the common sector formats used in contemporary hard disk drives. The standard format used for many years divided the track into sectors, each containing 512 bytes of data.

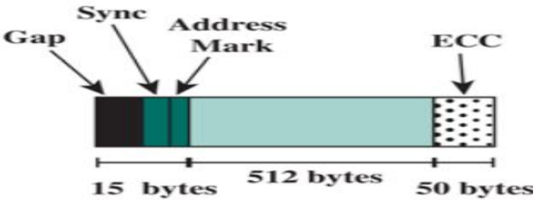

Figure 6.3 Legacy 512-byte sector

1. Gap: Separates sectors.
2. Sync: Indicates the beginning of the sector and provides timing alignment.

3. Address mark: Contains data to identify the sector's number and location. It also provides status about the sector itself.
4. Data: The 512 bytes of user data.
5. Error correction code (ECC): Used to correct data that might be damaged in the reading and writing process

6.2 RAID (Redundant Arrays of Independent Disks)

RAID, or "Redundant Arrays of Independent Disks" is a technique which makes use of a combination of multiple disks instead of using a single disk for increased performance, data redundancy or both.

Why data redundancy?data redundancy, although taking up extra space, adds to disk reliability. This means, in case of disk failure, if the same data is also backed up onto another disk, we can retrieve the data and go on with the operation. On the other hand, if the data is spread across just multiple disks without the RAID technique, the loss of a single disk can affect the entire data.

RAID 0:

- RAID 0 uses multiple disks and maps them as a single disk.This is mainly for performance and not for fault tolerance.
- If any drive in RAID 0 fails whole system becomes unusable.
- There is no duplication of data. Hence, a block once lost cannot be recovered.

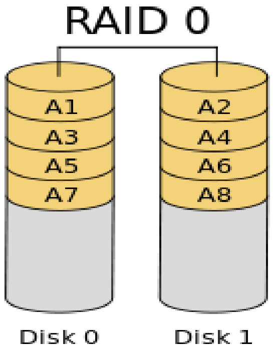

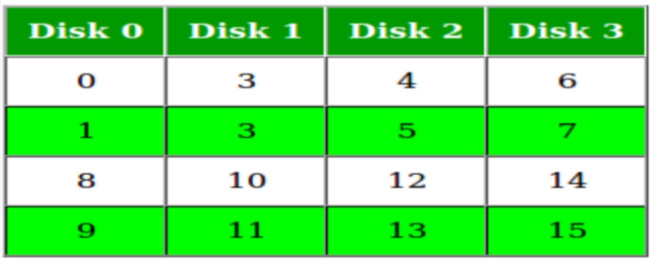

Disk 0	Disk 1	Disk 2	Disk 3
0	3	4	6
1	3	5	7
8	10	12	14
9	11	13	15

Figure 6.4 RAID-0

RAID 1:

- It's basically disk mirroring.it provides 100% redundancy as everything is stored in two disks and if one disk fails we can restore data from backup disk and whole raid can be recreated.
- RAID 0 was unable to tolerate any disk failure. But RAID 1 is capable of reliability.
- 1 disk failure can be handled for certain, because blocks of that disk would have duplicates on some other disk.

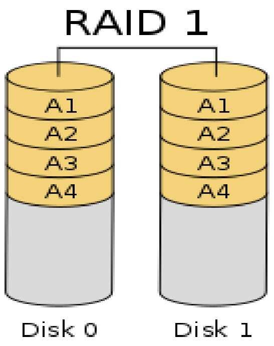

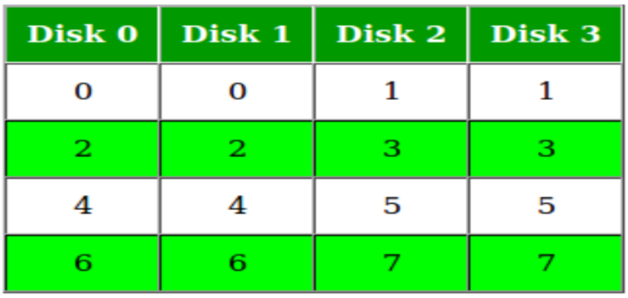

Disk 0	Disk 1	Disk 2	Disk 3
0	0	1	1
2	2	3	3
4	4	5	5
6	6	7	7

Figure 6.5 RAID-1

RAID 2:

- This configuration uses striping across disks, with some disks storing error checking and correcting information.
- It has no advantage over RAID 3 and is no longer used.
- Figure 2. In RAID 2, data is split at the bit level over a number of data and ECC(error-correcting code) disks. Every time data is written to the array, the Hamming codes are calculated and written to the ECC disks. When the data is read from the array, these ECC codes are read as well to confirm that no errors have occurred since the data was written. If a single-bit error occurs, it can be corrected immediately.

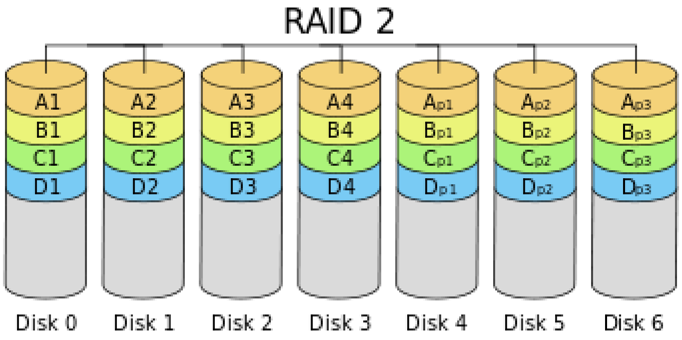

Figure 6.6 RAID-2

RAID 3:

RAID 3 is organized in a similar fashion to RAID 2.

- It uses striping at the byte level and stores dedicated parity bits on a separate disk drive.
- Instead of striping data blocks into different disks, RAID 3 stripes the bits, which are stored on different disk drives. This configuration is used less commonly than other RAID levels.
- The difference is that RAID 3 requires only a single redundant disk, no matter how large the disk array.
- RAID 3 uses the more effective XOR algorithm to generate parity

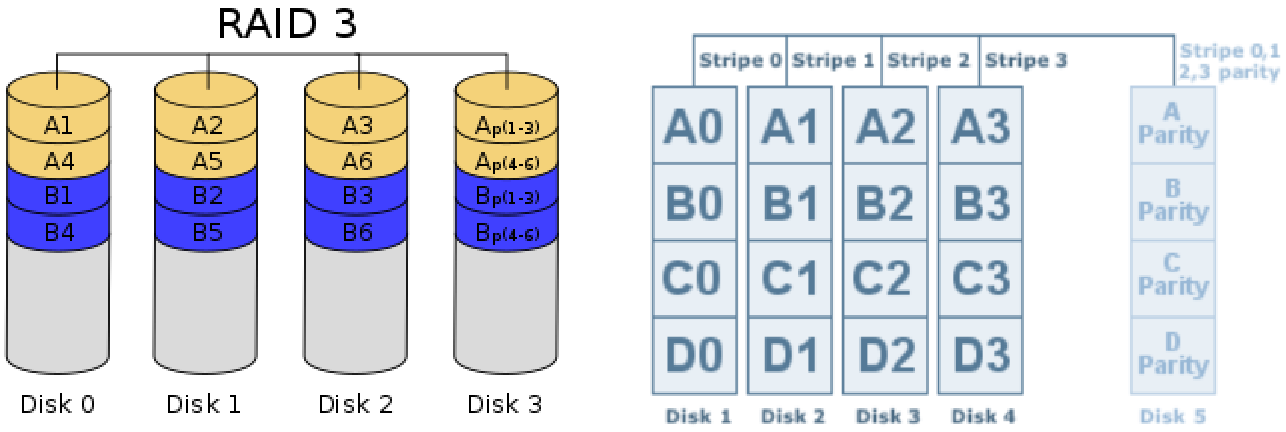

Figure 6.7 RAID-3

Data reconstruction is simple. Consider an array of five drives in which X0 through X3 contain data and X4 is the parity disk. The parity for the *i*th bit is calculated as follows:

$$X4(i) = X3(i) \oplus X2(i) \oplus X1(i) \oplus X0(i)$$

where $\oplus$ is exclusive-OR function.

Suppose that drive X1 has failed. If we add $X4(i) \oplus X1(i)$ to both sides of the preceding equation, get

$$X1(i) = X4(i) \oplus X3(i) \oplus X2(i) \oplus X0(i)$$

RAID 4:

stripes the data across multiple disks just like RAID 0. In addition to that, it also stores parity information of all the disks in a separate dedicated disk to achieve redundancy. In the diagram below, Disk 4 serves as the parity disk having parity blocks A_p, B_p, C_p and D_p. So, if one of the disks fails, the data can be reconstructed using the parity information of that disk.

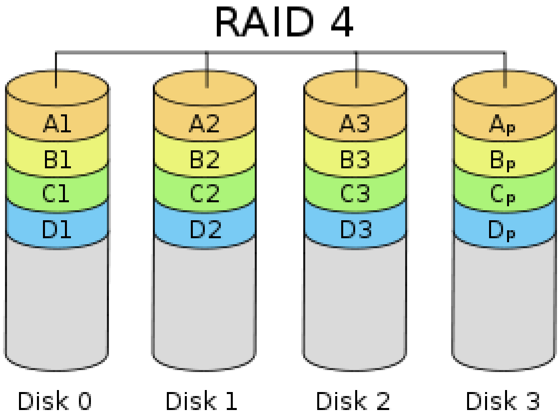

Figure 6.8 RAID-4

Consider an array of five drives in which X0 through X3 contain data and X4 is the parity disk. Suppose that a write is performed that only involves a strip on disk X1. Initially, for each bit i, we have the following relationship:

$X4(i) = X3(i) \oplus X2(i) \oplus X1(i) \oplus X0(i)$

After the update, with potentially altered bits indicated by a prime symbol:

$X4'(i) = X4(i) \oplus X1(i) \oplus X1'(i)$

To calculate the new parity, the array management software must read the old user strip and the old parity strip. Then it can update these two strips with the new data and the newly calculated parity. Thus, each strip write involves two reads and two writes

RAID 5:

RAID 5 is a data backup technology for hard disk drives that uses both disk striping and parity, ALSO can be used for solid-state drives. It is one of the levels of RAID. RAID 5 doesn't use disk mirroring, however; it combines disk striping and a data-checking technique called parity. RAID 5 parity is spread among each drive, unlike RAID 4, which stores it all on one disk. This provides added security (not all parity data is lost if that drive fails). Also, a general note about disk drives is that the exact amount of room available on a disk may not be what's advertised, because formatting RAID configurations takes some of that space.

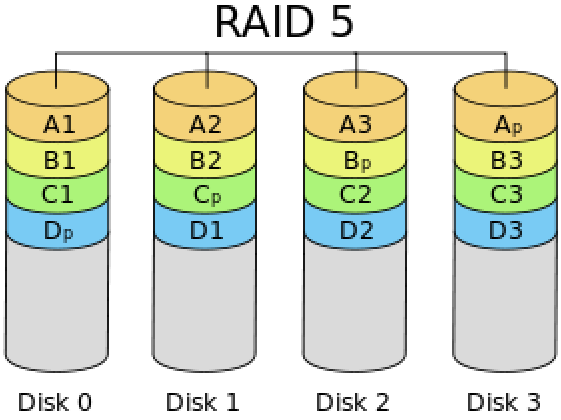

Figure 6.9 RAID-5

RAID LEVEL 6

In the RAID 6 scheme, two different parity calculations are carried out and stored in separate blocks on different disks. Thus, a RAID 6 array whose user data require N disks consists of N+2 disks. Figure 6.10 illustrates the scheme. *P and Q* are two different data check algorithms. One of the two is the exclusive-OR calculation used in RAID 4 and 5. But the other is an independent data check algorithm. This makes it possible to regenerate data even if two disks containing user data fail.

that RAID 6 arrays will need to dedicate approximately two disks' worth of data to parity, because the amount of data stored in parity is about the storage size of two disks.

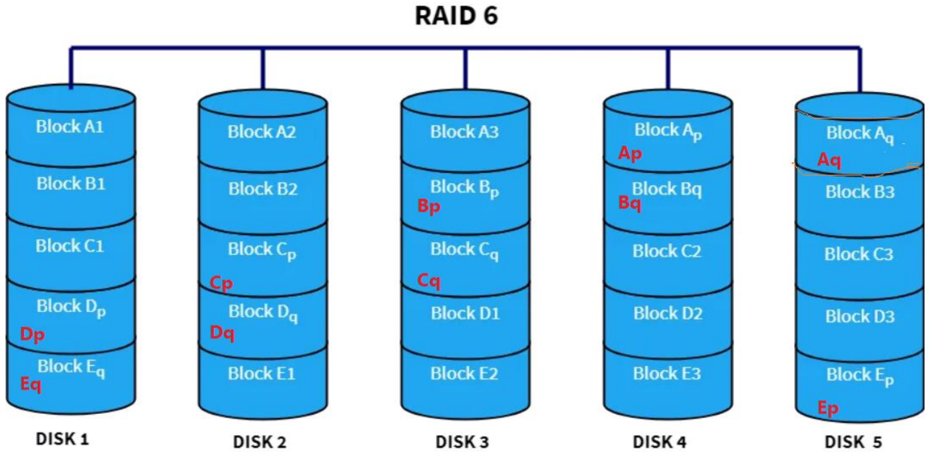

Figure 6.10 RAID-6

6.3 Solid State Drives

One of the most significant developments in computer architecture in recent years is the increasing use of solid state drives (SSDs) to complement or even replace hard disk drives (HDDs), both as internal and external secondary memory. The term solid state refers to electronic circuitry built with semiconductors. An SSD is a memory device made with solid state components that can be used as a replacement to a hard disk drive. The SSDs now on the market and coming on line use NAND flash memory.

SSD Compared to HDD As the cost of flash-based SSDs has dropped and the performance and bit density increased, SSDs have become increasingly competitive with HDDs.The SSDs have the following advantages over HDDs:

- High-performance input/output operations per second (IOPS)
- Durability: Less susceptible to physical shock and vibration.
- Longer lifespan
- Lower power consumption
- Quieter and cooler running capabilities
- lower energy costs, and a greener enterprise

SSD Organization

Figure 6.11 illustrates a general view of the common architectural system component associated with any SSD system.

- On the host system, the operating system invokes file system software to access data on the disk.
- The file system, in turn, invokes I/O driver software.
- The I/O driver software provides host access to the particular SSD product.

The interface component in Figure 6.11 refers to the physical and electrical interface between the host processor and the SSD peripheral device. If the device is an internal hard drive, a common interface is PCIe. For external devices, one common interface is USB.

PCIe (peripheral component interconnect express) is an **interface** standard for connecting high-speed components. Every desktop PC motherboard has a number of **PCIe** slots you can use to add GPUs (aka video cards aka graphics cards), RAID cards, Wi-Fi cards or SSD (solid-state drive) add-on cards.

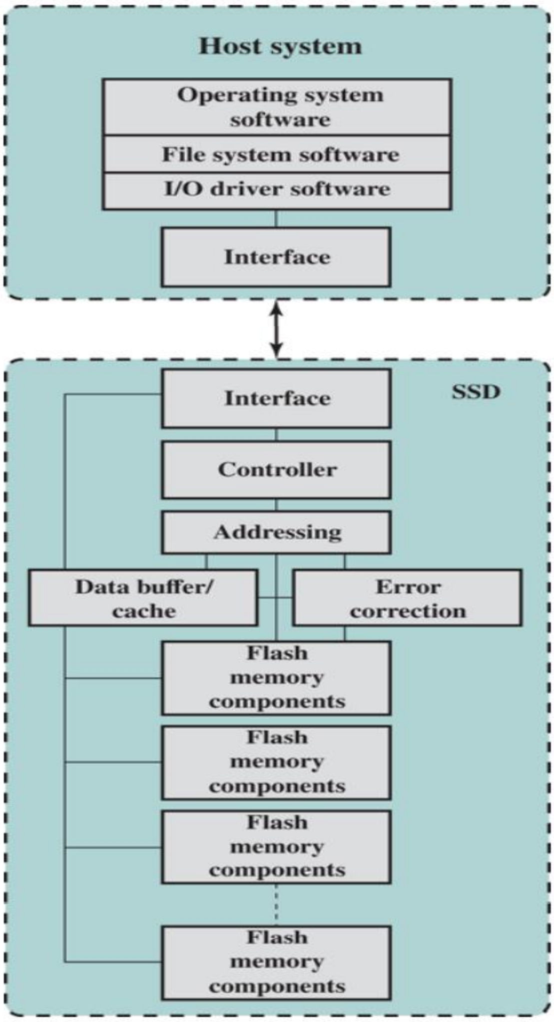

Figure 6.11 Solid State Drive Architecture

6.4 Optical Memory

Optical memory is an electronic storage medium that uses a laser beam to store and retrieve the data. If we classify the memory system then **optical memory** comes under the external memory in the computer system. **Optical memory** can be classified into many types.

6.4.1 Optical Disk Products

CD

CDs or Compact Disks are *opticL readable media*. The main material of the CD is plastic. The shape of the plastic is circular and one side of the circular plastic is coated with the reflecting metal coating, usually aluminum. CDs can store many types of data, like audio, video, games, any documents, etc. The data are scanned by a laser beam with a CD driver to visualize the data. The storage capacity of a CD is 700 MB only.

The standard diameter of a CD is 120mm and thickness is 1.2mm. The main body of the CD is made of hard plastic known as **polycarbonate** with a reflective metallic layer. This layer is coated with metallic **acrylic plastic**. This metal is aluminum. In this layer, the data are stored using the laser light which reflects the coated layer for reading and writing the data. The data is read in the form of **pits**, each pit is of **0.83-micrometre** and the data is arranged as **spiral track** from the disc's inner hole to its outer edge, because the CD is of circular shape.

CD-ROM

Both the audio CD and the CD-ROM (compact disk read-only memory) share a similar technology. The main difference is that CD-ROM players are more rugged and have error correction devices to ensure that data are properly transferred from disk to computer. Both types of disk are made the same way.

The disk is formed from a resin, such as polycarbonate. Digitally recorded information (either music or computer data) is imprinted as a series of microscopic pits on the surface of the polycarbonate., the disk contains a single spiral track, beginning near the center and spiraling out to the outer edge of the disk. Sectors near the outside of the disk are the same length as those near the inside. Thus, information is packed evenly across the disk in segments of the same size and these are scanned at the same rate by rotating the disk at a variable speed. The pits are then read by the laser at a constant linear velocity (CLV). The disk rotates more slowly for accesses near the outer edge than for those near the center. Thus, the capacity of a track and the rotational delay both increase for positions nearer the outer edge of the disk. The data capacity for a CD-ROM is about 680 MB. Data on the CD-ROM are organized as a sequence of blocks. A typical block format is shown in Figure 6.12 It consists of the following fields

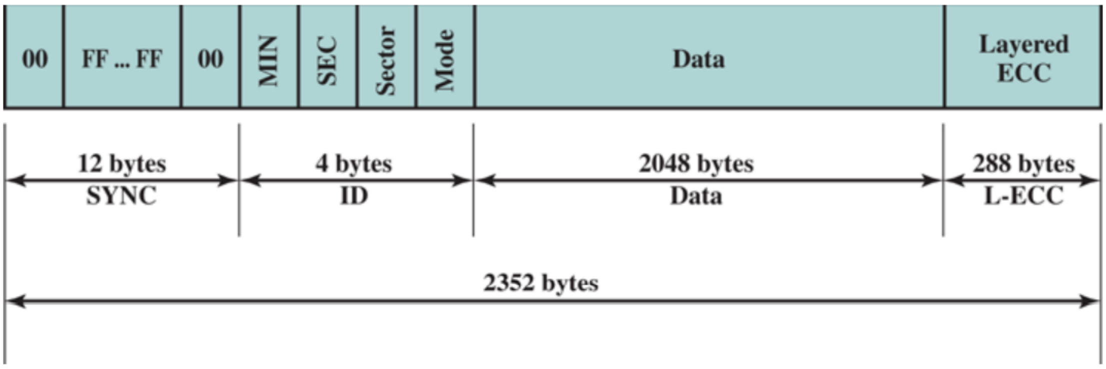

Figure 6.12 CD-ROM Block Format

- Sync: The sync field identifies the beginning of a block.
- Header: The header contains the block address and the mode byte. Mode 0 specifies a blank data field; mode 1 specifies the use of an error-correcting code and 2048 bytes of data; mode 2 specifies 2336 bytes of user data with no error-correcting code
- Data: User data.
- Auxiliary: Additional user data in mode 2. In mode 1, this is a 288-byte error-correcting code.

CD RECORDABLE

To accommodate applications in which only one or a small number of copies of a set of data is needed, the write-once read-many CD, known as the CD recordable (CD-R), has been developed. For CD-R, a disk is prepared in such a way that it can be subsequently written once with a laser beam of modest intensity. Thus, with a somewhat more expensive disk controller than for CD-ROM, the customer can write once as well as read the disk. The CD-R medium is similar but not identical to that of a CD or CD-ROM. For CDs and CD-ROMs, information is recorded by the pitting of the surface of the medium, which changes reflectivity. For a CD-R, the medium includes a dye layer. The dye is used to change reflectivity and is activated by a high-intensity laser. The resulting disk can be read on a CD-R drive or a CD-ROM drive

CD REWRITABLE

The CD-RW optical disk can be repeatedly written and overwritten, as with a magnetic disk. Although a number of approaches have been tried, the only pure optical approach that has proved attractive is called *phase change*. The phase change disk uses a material that has two significantly different reflectivity's in two different phase states. A beam of laser light can change the material from one phase to the other. The primary disadvantage of phase change optical disks is that the material eventually and permanently loses its desirable properties. Current materials can be used for between 500,000 and 1,000,000 erase cycles. The CD-RW has the obvious advantage over CD-ROM and CD-R that it can be rewritten and thus used as a true secondary storage.

DIGITAL VERSATILE DISK

The DVD takes video into the digital age. It delivers movies with impressive picture quality, and it can be randomly accessed like audio CDs, which DVD machines can also play. Vast volumes of data can be crammed onto the disk, currently seven times as much as a CD-ROM. With DVD’s huge storage capacity and vivid quality, PC games have become more realistic and educational software incorporates more video. The DVD’s greater capacity is due to three differences from CDs (Figure 6.13):

1. Bits are packed more closely on a DVD.
2. The DVD employs a second layer of pits and lands on top of the first layer.
3. The DVD-ROM can be two sided, whereas data are recorded on only one side of a CD. This brings total capacity up to 17 GB

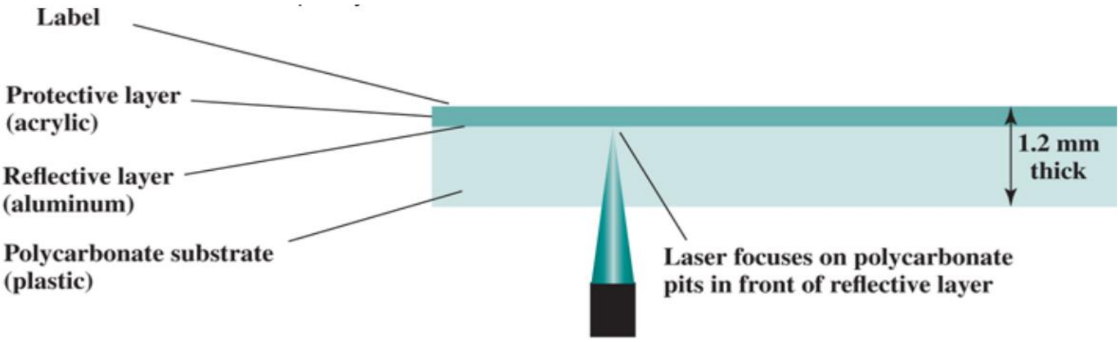

(a) CD-ROM–Capacity 682 MB

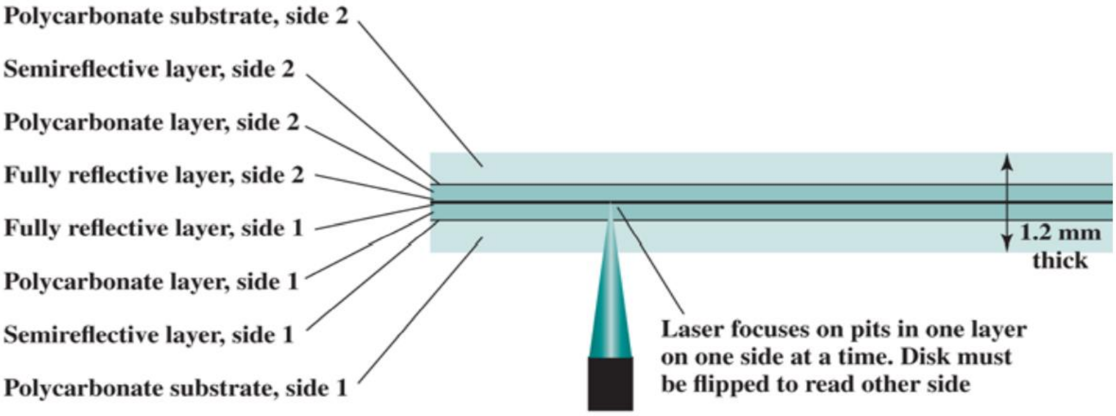

(b) DVD-ROM, double-sided, dual-layer–Capacity 17 GB